



System Manager

SANtricity 11.6

NetApp
February 12, 2024

Table of Contents

- System Manager 1
 - Main interface..... 1
 - Storage..... 15
 - Mirroring..... 214
 - Hardware 256
 - Settings 302
 - Support..... 408

System Manager

Main interface

Home page

The Home page provides a dashboard view for the daily management of your storage array. When you log into System Manager, the Home page is the first screen displayed.

The dashboard view comprises four summary areas that contain key information about the state and health of your storage array. You can find more information from the summary area.

- **Notifications**

The Notifications area displays problem notifications that indicate the status of the storage array and its components. In addition, this portlet displays automated alerts that can help you troubleshoot issues before it affects other areas of your storage environment.

- **Performance**

The Performance area allows you to compare and contrast resource usage over time. You can view a storage array's performance metrics for response time (IOPS), transfer rates (MiB/s), and the amount of processing capacity being used (CPU).

- **Capacity**

The Capacity area displays a chart view of the allocated capacity, free storage capacity, and unassigned storage capacity in your storage array.

- **Storage Hierarchy**

The Storage Hierarchy area provides an organized view of the various hardware components and storage objects managed by your storage array. Click the drop-down arrow to perform a certain action on that hardware component or storage object.

What's new in SANtricity OS 11.6

The following tables describe new features in SANtricity OS 11.6.

New features in Version 11.62

New feature	Description
Downloadable CLI	System Manager for the E5700, EF570, E2800, and EF280 arrays now includes the ability to download and install the SANtricity command line interface (CLI) via a link in the Settings > System > Add-ons page. This is the https-based version of the CLI (also referred to as "Secure CLI"). This capability was previously released with the EF600 array.

New feature	Description
Mirroring configuration changes in System Manager and Unified Manager	The tasks for configuring synchronous and asynchronous mirrored pairs have moved from System Manager to Unified Manager. All other tasks for managing mirrored pairs remain in System Manager.
New 200Gb-capable HIC (EF600 arrays only)	This release adds a new 200Gb-capable HIC for EF600 storage arrays. Interfaces supported are NVMe/IB, NVMe/RoCE, and iSER/IB. Additionally, 100Gb SRP/IB is supported.
Additional options on 100Gb HIC (EF600 arrays only)	On the existing 100Gb HIC, iSER/IB and SRP/IB interfaces are now supported for EF600 storage arrays. (These interfaces are already supported for EF570 and E5700 arrays.)
Delete mail server in System Manager	System Manager allowed a mail server to be configured, but did not have an easy mechanism to remove it. With this release, the mail server configuration in System Manager can now be removed from Alerts, so that alerts are no longer sent to the email addresses associated with this mail server.
Optimization capacity adjustments for pools and volume groups (SSD drives only) in System Manager	For SSD drives, a new optimization capacity slider is available in System Manager for the Pool settings and Volume Group settings. The slider enables you to adjust the balance of available capacity versus SSD write performance and drive wear life.
New host types in System Manager	When you create new hosts in System Manager, the presented host options are now organized into three categories to provide better guidance: Common, Uncommon, and Use only if directed.

New features in Version 11.61

New feature	Description
Fibre Channel support for the EF600	This release adds Fibre Channel host support for the EF600 storage system. This is the first SCSI host supported by the EF600, which initially released with all NVMe over Fabrics host protocols. A single controller for the EF600 can be viewed and managed in System Manager. Multiple EF600 storage systems can be viewed and managed in Unified Manager.
Password requirements for admin user	For first-time login in Unified Manager, you must now enter a password for the administrator user. There is no longer a default "admin" password.

New features in Version 11.60

New feature	Description
New storage system model – EF600	<p>This release offers a new EF600 all-flash storage system. The EF600 includes NVMe-oF host interfaces and NVMe SSDs.</p> <p>The EF600 significantly increases throughput and reduces latency. The supported host interfaces include NVMe over IB, NVMe over RoCE, and NVMe over FC, which can be configured in System Manager. Multiple EF600 storage systems can be viewed and managed in Unified Manager.</p>
Downloadable CLI	<p>System Manager now includes the ability to download and install the SANtricity command line interface (CLI) via a link in the Settings > System > Add-ons page. This is the https- based version of the CLI. The legacy SANtricity Storage Manager package continues to include the CLI as well.</p>

Setup wizard

Concepts

Setup wizard overview

Use the Setup wizard to configure your storage array, including hardware, hosts, applications, workloads, pools, alerts, and AutoSupport.

When you open System Manager for the first time, the Setup wizard launches.

The Setup wizard prompts you to perform basic configuration tasks, such as naming your storage array, configuring your hosts, selecting applications, and creating pools of storage.

If you cancel the wizard, you cannot manually relaunch it.

The wizard automatically relaunches when you open System Manager or refresh your browser and *at least one* of the following conditions is met:

- No pools and volume groups are detected.
- No workloads are detected.
- No notifications are configured.

Setup wizard terminology

Learn how the Setup wizard terms apply to your storage array.

Term	Description
Application	An application is a software program, such as Microsoft SQL Server or Microsoft Exchange.
Alert	Alerts notify administrators about important events that occur on the storage arrays. Alerts can be sent via email, SNMP traps, or syslog.

Term	Description
AutoSupport	The AutoSupport feature monitors the health of a storage array and sends automatic dispatches to technical support.
Hardware	The storage system hardware includes storage arrays, controllers, and drives.
Host	A host is a server that sends I/O to a volume on a storage array.
Object	An object is any logical or physical storage component. Logical objects include volume groups, pools, and volumes. Physical objects include the storage array, array controllers, hosts, and drives.
Pool	A pool is a set of drives that is logically grouped. You can use a pool to create one or more volumes accessible to a host. (You create volumes from either a pool or a volume group.)
Volume	<p>A volume is a container in which applications, databases, and file systems store data. It is the logical component created for the host to access storage on the storage array.</p> <p>A volume is created from the capacity available in a pool or a volume group. A volume has a defined capacity. Although a volume might consist of more than one drive, a volume appears as one logical component to the host.</p>
Volume group	A volume group is a container for volumes with shared characteristics. A volume group has a defined capacity and RAID level. You can use a volume group to create one or more volumes accessible to a host. (You create volumes from either a volume group or a pool.)
Workload	A workload is a storage object that supports an application. You can define one or more workloads, or instances, per application. For some applications, System Manager configures the workload to contain volumes with similar underlying volume characteristics. These volume characteristics are optimized based on the type of application the workload supports. For example, if you create a workload that supports a Microsoft SQL Server application and then subsequently create volumes for that workload, the underlying volume characteristics are optimized to support Microsoft SQL Server.

FAQs

What if I don't see all of my hardware components?

If you do not see all your hardware components on the **Verify hardware** dialog box, it could mean that a drive shelf is not connected correctly, or that an incompatible shelf is installed in the storage array.

Verify that all drive shelves are connected correctly. If you are uncertain about which drive shelves are compatible, contact technical support.

What if I don't see all of my hosts?

If you do not see your connected hosts, then automatic detection has failed, the hosts are improperly connected, or no hosts are currently connected.

You can configure hosts later, once you are done with the setup. You can create hosts either automatically or manually as follows:

- If you installed the Host Context Agent (HCA) on your hosts, the HCA pushes the host configuration information to the storage array. System Manager automatically configures these hosts and displays them in the Initial Setup wizard. (HCA does not apply to NVMe over Fabrics hosts.)
- You can manually create hosts and associate the appropriate host port identifiers by going to **Storage > Hosts**. Hosts that have been created manually also display in the Initial Setup wizard.
- The target and host must be configured for the host port type (for example, iSCSI or NVMe over RoCE), and a session to the storage established before automatic detection will work.

How does identifying applications help me manage my storage array?

When you identify applications, System Manager automatically recommends a volume configuration that optimizes storage based on application type.

Optimizing volumes by application can make data storage operations more efficient. Characteristics such as I/O type, segment size, controller ownership, and read and write cache are included in the volume configuration. In addition, you can view performance data by application and by workload to assess the latency, IOPS, and MiB/s of applications and their associated workloads.

What is a workload?

For some applications in your network, such as SQL Server or Exchange, you can define a workload that optimizes storage for that application.

A workload is a storage object that supports an application. You can define one or more workloads, or instances, per application. For some applications, System Manager configures the workload to contain volumes with similar underlying volume characteristics. These volume characteristics are optimized based on the type of application the workload supports. For example, if you create a workload that supports a Microsoft SQL Server application and then subsequently create volumes for that workload, the underlying volume characteristics are optimized to support Microsoft SQL Server.

During volume creation, System Manager prompts you to answer questions about a workload's use. For example, if you are creating volumes for Microsoft Exchange, you are asked how many mailboxes you need, what your average mailbox capacity requirements are, and how many copies of the database you want. System Manager uses this information to create an optimal volume configuration for you, which can be edited as needed.

How do I configure the delivery method for AutoSupport?

To access configuration tasks for AutoSupport delivery methods, go to **Support > Support Center**, and then click the **AutoSupport** tab.

The following protocols are supported: HTTPS, HTTP, and SMTP.

How do I know if I should accept the recommended pool configuration?

Whether you accept the recommended pool configuration depends on a few factors.

Determine the type of storage that is best for your requirements by answering the following questions:

- Do you prefer multiple pools of smaller capacities, rather than the largest pools possible?
- Do you prefer RAID volume groups over pools?
- Do you prefer to manually provision your drives, rather than having a configuration recommended for you?

If you answered Yes to any of these questions, consider rejecting the recommended pool configuration.

System Manager has not detected any hosts. What do I do?

If you do not see your connected hosts, then automatic detection has failed, the hosts are improperly connected, or no hosts are currently connected.

You can configure hosts later, once you are done with the setup. You can create hosts either automatically or manually as follows:

- If you installed the Host Context Agent (HCA) on your hosts, the HCA pushes the host configuration information to the storage array. System Manager automatically configures these hosts and displays them in the Initial Setup wizard. (HCA does not apply to NVMe over Fabrics hosts.)
- You can manually create hosts and associate the appropriate host port identifiers by going to **Storage > Hosts**. Hosts that have been created manually also display in the Initial Setup wizard.
- The target and host must be configured for the host port type (for example, iSCSI or NVMe over RoCE), and a session to the storage established before automatic detection will work.

Storage array administration

Concepts


Storage array overview

A storage array is a storage entity that is managed by the System Manager software. A storage array consists of a collection of both physical components and logical components.

Physical components

The physical components of a storage array are described in this table.

Component	Description
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the System Manager functions.

Component	Description
Shelf	<p>A shelf is an enclosure installed in a cabinet or rack. It contains the hardware components for the storage array. There are two types of shelves: a controller shelf and a drive shelf. A controller shelf includes controllers and drives. A drive shelf includes input/output modules (IOMs) and drives.</p> <div>  <p>If your storage array contains different media types or different interface types, a drive shelf for each drive type appears.</p> </div>
Drive	A drive is an electromagnetic mechanical device or solid state memory device that provides the physical storage media for data.
Host	A host is a server that sends I/O to a volume on a storage array.
Host bus adapter (HBA)	A host bus adapter (HBA) is a board that resides in a host and contains one or more host ports.
Host port	A host port is a port on a host bus adapter (HBA) that provides the physical connection to a controller and is used for I/O operations.
Management client	A management client is the computer where a browser is installed for accessing System Manager.

Logical components

The drives in the storage array provide the physical storage capacity for data. Use System Manager to configure the physical capacity into logical components, such as pools, volume groups, and volumes. These components are the tools that you use to configure, store, maintain, and preserve data on the storage array. The logical components of a storage array are described in this table.

Component	Description
Pool	A pool is a set of drives that is logically grouped. You can use a pool to create one or more volumes accessible to a host. (You create volumes from either a pool or a volume group.)
Volume group	A volume group is a container for volumes with shared characteristics. A volume group has a defined capacity and RAID level. You can use a volume group to create one or more volumes accessible to a host. (You create volumes from either a volume group or a pool.)
Volume	A volume is a container in which applications, databases, and file systems store data. It is the logical component created for the host to access storage on the storage array.

Component	Description
Logical unit number (LUN)	<p>A logical unit number (LUN) is the number assigned to the address space that a host uses to access a volume. The volume is presented to the host as capacity in the form of a LUN.</p> <p>Each host has its own LUN address space. Therefore, the same LUN can be used by different hosts to access different volumes.</p>

Administrator password protection

You must configure the storage array with an administrator password to protect it from unauthorized access.

Setting an administrator password

Setting an administrator password for your storage array protects it from users who unknowingly or maliciously run destructive commands. You are required to set an administrator password when you start System Manager for the first time.

There is one administrator password that is shared among all users. Any user who has this password can make configuration changes to the storage array, such as adding, changing, or removing objects or settings.

Entering your password

The system prompts you for the password only once during a single management session. A session times out after 20 minutes of inactivity, at which time, you must enter the password again.

If another user managing the same storage array from another management client changes the password while your session is in progress, you are prompted for a password the next time you attempt a configuration operation or a view operation.

For security reasons, you can attempt to enter a password only five times before the storage array enters a "lockout" state. In this state, the storage array will reject subsequent password attempts. You must wait 10 minutes for the storage array to reset to a "normal" state before you try to enter a password again.

Removing drives with password protection

The password is stored on a reserved area of each drive on the storage array. If you remove all drives from a storage array, its password will no longer work. To correct this condition, re-install one of the original drives to the storage array.

Removing password protection

If you no longer want to have commands password-protected, enter the current administrator password, and leave the new password text boxes blank.







Running configuration commands on a storage array can cause serious damage, including data loss. For this reason, you should always set an administrator password for your storage array. Use a long administrator password with at least 15 alphanumeric characters to increase security.

Problem notification

System Manager uses icons and several other methods to notify you that problems exist with the storage array.

Icons

System Manager uses these icons to indicate the status of the storage array and its components.

Icon	Description
	Optimal
	Non-optimal or failed
	Needs attention or fixing
	Caution

System Manager displays these icons in various locations.

- The Notifications area on the Home page displays the failed icon and a message.
- The Home page icon in the navigation area displays the failed icon.
- On the Components page, the graphics for drives and controllers display the failed icon.

Alerts and LEDs

In addition, System Manager notifies you of problems in other ways.

- System Manager sends SNMP notifications or email error messages.
- The Service Action Required LEDs on the hardware come on.

When you receive notification of a problem, use the Recovery Guru to help you fix the problem. Where necessary, use the hardware documentation with the recovery steps to replace failed components.

Operations in progress

The Operations in Progress page allows you to view a list of long-running operations that are currently running or are in a suspended, stopped, or failed state. You can take action on certain types of operations.

Long-running operations can affect system performance. For some operations, you have the option to stop the operation or change its priority so that more important operations complete first.

Operations that appear on the Operations in Progress page include the following:

Operation	Possible status of the operation	Actions you can take
Volume copy	Completed	Clear

Operation	Possible status of the operation	Actions you can take
Volume copy	In progress	<ul style="list-style-type: none"> • Change priority • Stop
Volume copy	Pending	Clear
Volume copy	Failed	<ul style="list-style-type: none"> • Clear • Re-copy
Volume copy	Stopped	<ul style="list-style-type: none"> • Clear • Re-copy
Volume create (thick pool volumes larger than 64TiB only)	In progress	<i>none</i>
Volume delete (thick pool volumes larger than 64TiB only)	In progress	<i>none</i>
Asynchronous mirror group initial synchronization	In progress	Suspend
Asynchronous mirror group initial synchronization	Suspended	Resume
Synchronous mirroring	In progress	Suspend
Synchronous mirroring	Suspended	Resume
Snapshot image rollback	In progress	Cancel
Snapshot image rollback	Pending	Cancel
Snapshot image rollback	Paused	<ul style="list-style-type: none"> • Cancel • Resume
Drive evacuation	In progress	Cancel (depends on the drive evacuation type)
Add capacity to pool or volume group	In progress	<i>none</i>
Change a RAID level for a volume	In progress	<i>none</i>
Reduce capacity for a pool	In progress	<i>none</i>

Operation	Possible status of the operation	Actions you can take
Thin volume reclamation	In progress	<i>none</i>
Check the time remaining on an instant availability format (IAF) operation for pool volumes	In progress	<i>none</i>
Check the data redundancy of a volume group	In progress	<i>none</i>
Defragment a volume group	In progress	<i>none</i>
Initialize a volume	In progress	<i>none</i>
Increase capacity for a volume	In progress	<i>none</i>
Change segment size for a volume	In progress	<i>none</i>
Drive copy	In progress	<i>none</i>
Data reconstruction	In progress	<i>none</i>
Copyback	In progress	<i>none</i>

How to

Set default units for capacity values

System Manager can display capacity values in either gibibytes (GiB) or tebibytes (TiB).

About this task

Abbreviation	Value
GiB	1,024 ³ bytes
TiB	1,024 ⁴ bytes

Preferences are stored in the browser's local storage so all users can have their own settings.

Steps

1. Select **Preferences** > **Set preferences**.
2. Click the radio button for either **Gibibytes** or **Tebibytes**, and confirm that you want to perform the operation.

Set default time frame for performance graphs

You can change the default time frame shown by the performance graphs.

About this task

Performance graphs shown on the Home page and on the Performance page initially show a time frame of 1 hour. Preferences are stored in the browser's local storage so all users can have their own settings.

Steps

1. Select **Preferences** › **Set preferences**.
2. In the drop-down list, select either **5 minutes**, **1 hour**, **8 hours**, **1 day**, or **7 days**, and confirm that you want to perform the operation.

Recover from problems using Recovery Guru

The Recovery Guru is a component of System Manager that diagnoses storage array problems and recommends recovery procedures to fix the problems.

Steps

1. Select **Home**.
2. Click the link labeled **Recover from n problems** in the center-top of the window.

The Recovery Guru dialog appears.

3. Select the first problem shown in the summary list, and then follow the instructions in the recovery procedure to correct the problem. Where necessary, use the replacement instructions to replace failed components. Repeat this step for each listed problem.

Multiple problems within a storage array can be related. In this case, the order in which the problems are corrected can affect the outcome. Select and correct the problems in the order that they are listed in the summary list.

Multiple failures for a power-supply canister are grouped and listed as one problem in the summary list. Multiple failures for a fan canister are also listed as one problem.

4. To make sure that the recovery procedure was successful, click **Recheck**.

If you selected a problem for an asynchronous mirror group or a member of an asynchronous mirror group, click **Clear** first to clear the fault from the controller, and then click **Recheck** to remove the event from the Recovery Guru.

If all of the problems have been corrected, the storage array icon eventually transitions from Needs Attention to Optimal. For some problems, a Fixing icon appears while an operation, such as reconstruction, is in progress.

5. **Optional:** To save the Recovery Guru information to a file, click the **Save** icon.

The file is saved in the Downloads folder for your browser with the name `recovery-guru-failure-yyyy-mm-dd-hh-mm-ss-mmm.html`.

6. **Optional:** To print the Recovery Guru information, click the **Print** icon.

View and act on operations in progress

To view and take action on long-running operations, use the Operations in Progress page.

About this task

For each operation listed on the Operations in Progress page, a percentage of completion and estimated time remaining to complete the operation are shown. In some cases, you can stop an operation or place it at a higher or lower priority. You can also clear a completed Volume copy operation from the list.

Steps

1. On the **Home** page, select **Show operations in progress**.

The **Operations in Progress** page appears.

2. If desired, use the links in the **Actions** column to stop or change priority for an operation.



Read all cautionary text provided in the dialog boxes, particularly when stopping an operation.

You can stop a volume copy operation or change its priority.

3. Once a volume copy operation is complete, you can select **Clear** to remove it from the list.

At the top of the **Home** page, an informational message and yellow wrench icon appear when an operation is complete. This message includes a link that allows you to clear the operation from the **Operations in Progress** page.

FAQs

Where can I find more information about display preferences?

You can find information about the available display options in the following help topics.

- To read more about the default units for displaying capacity values, see [Set default units for capacity values](#).
- To read more about the default time frame for displaying performance graphs, see [Set default time frame for performance graphs](#).

What are the supported browsers?

System Manager supports these browser versions.

Browser	Minimum version
Google Chrome	47
Microsoft Internet Explorer	11
Microsoft Edge	EdgeHTML 12

Browser	Minimum version
Mozilla Firefox	31
Safari	9

What are the keyboard shortcuts?

You can navigate around System Manager using the keyboard alone.

Overall navigation

Action	Keyboard shortcut
Move to the next item.	Tab
Move to the previous item.	Shift + Tab
Select an item.	Enter
Drop-down list—Move to the next or previous item.	Down arrow or up arrow
Check box—Select an item.	Spacebar
Radio buttons—Toggle between items.	Down arrow or up arrow
Expandable text—Expand or contract item.	Enter

Table navigation

Action	Keyboard shortcut
Select a row.	Tab to select a row, then press Enter
Scroll up or down.	Down arrow/up arrow or Page Down/Page Up
Change the sort order of a column.	Tab to select a column heading, then press Enter

Calendar navigation

Action	Keyboard shortcut
Move to the previous month.	Page Up
Move to the next month.	Page Down

Action	Keyboard shortcut
Move to the previous year.	Control + Page Up
Move to the next year.	Control + Page Down
Open the date picker if closed.	Control + Home
Move to the current month.	Control / Command + Home
Move to the previous day.	Control / Command + Left
Move to the next day.	Control / Command + Right
Move to the previous week.	Control / Command + Up
Move to the next week.	Control / Command + Down
Select the focused date.	Enter
Close the date picker and erase the date.	Control / Command + End
Close the date picker without selection.	Escape

Storage

Pools and volume groups

Concepts


How pools and volume groups work

To provision storage, you create either a pool or volume group that will contain the Hard Disk Drives (HDD) or Solid State Disk (SSD) drives that you want to use in your storage array.

Physical hardware is provisioned into logical components so that data can be organized and easily retrieved. There are two types of groupings supported:

- Pools
- RAID volume groups

The pools and volume groups are the top-level units of storage in a storage array: they divide the capacity of drives into manageable divisions. Within these logical divisions are the individual volumes or LUNs where data is stored. The following figure illustrates this concept.



¹ Host LUNs; ² Volumes; ³ Volume groups or pools; ⁴ HDD or SSD drives

When a storage system is deployed, the first step is to present the available drive capacity to the various hosts by:

- Creating pools or volume groups with sufficient capacity
- Adding the number of drives required to meet performance requirements to the pool or volume group
- Selecting the desired level of RAID protection (if using volume groups) to meet specific business requirements

You can have pools or volume groups on the same storage system, but a drive cannot be part of more than one pool or volume group. Volumes that are presented to hosts for I/O are then created, using the space on the pool or volume group.

Pools

Pools are designed to aggregate physical hard disk drives into a large storage space and to provide enhanced RAID protection for it. A pool creates many virtual RAID sets from the total number of drives assigned to the pool, and it spreads the data out evenly among all participating drives. If a drive is lost or added, System Manager dynamically re-balances the data across all the active drives.

Pools function as another RAID level, virtualizing the underlying RAID architecture to optimize performance and flexibility when performing tasks such as rebuilding, drive expansion, and handling drive loss. The RAID level is automatically set at 6 in an 8+2 configuration (eight data disks plus two parity disks).

Drive matching

You can choose from either HDD or SSDs for use in pools; however, as with volume groups, all drives in the pool must use the same technology. The controllers automatically select which drives to include, so you must make sure that you have a sufficient number of drives for the technology you choose.

Managing failed drives

Pools have a minimum capacity of 11 drives; however, one drive's worth of capacity is reserved for spare capacity in the event of a drive failure. This spare capacity is called "preservation capacity."

When pools are created, a certain amount of capacity is preserved for emergency use. This capacity is expressed in terms of a number of drives in System Manager, but the actual implementation is spread across the entire pool of drives. The default amount of capacity that is preserved is based on the number of drives in the pool.

After the pool is created, you can change the preservation capacity value to more or less capacity, or even set it to no preservation capacity (0 drive's worth). The maximum amount of capacity that can be preserved (expressed as a number of drives) is 10, but the capacity that is available might be less, based on the total number of drives in the pool.

Volume groups

Volume groups define how capacity is allotted in the storage system to volumes. Disk drives are organized into RAID groups and volumes reside across the drives in a RAID group. Therefore, volume group configuration settings identify which drives are part of the group and what RAID level is used.

When you create a volume group, controllers automatically select the drives to include in the group. You must manually choose the RAID level for the group. The capacity of the volume group is the total of the number of drives that you select, multiplied by their capacity.

Drive matching

You must match the drives in the volume group for size and performance. If there are smaller and larger drives in the volume group, all drives are recognized as the smallest capacity size. If there are slower and faster drives in the volume group, all drives are recognized at the slowest speed. These factors affect the performance and overall capacity of the storage system.

You cannot mix different drive technologies (HDD and SSD drives). RAID 3, 5, and 6 are limited to a maximum of 30 drives. RAID 1 and RAID 10 uses mirroring, so these volume groups must have an even number of disks.

Managing failed drives

Volume groups use hot spare drives as a standby in case a drive fails in RAID 1/10, RAID 3, RAID 5, or RAID 6 volumes contained in a volume group. A hot spare drive contains no data and adds another level of redundancy to your storage array.

If a drive fails in the storage array, the hot spare drive is automatically substituted for the failed drive without requiring a physical swap. If the hot spare drive is available when a drive fails, the controller uses redundancy data to reconstruct the data from the failed drive to the hot spare drive.

Capacity terminology

Learn how the capacity terms apply to your storage array.

Storage objects

The following terminology describes the different types of storage objects that can interact with your storage array.


Storage object	Description
Host	A host is a server that sends I/O to a volume on a storage array.

Storage object	Description
LUN	<p>A logical unit number (LUN) is the number assigned to the address space that a host uses to access a volume. The volume is presented to the host as capacity in the form of a LUN.</p> <p>Each host has its own LUN address space. Therefore, the same LUN can be used by different hosts to access different volumes.</p>
Mirror consistency group	A mirror consistency group is a container for one or more mirrored pairs. For asynchronous mirroring operations, you must create a mirror consistency group.
Mirrored volume pair	A mirrored pair is comprised of two volumes, a primary volume and a secondary volume.
Pool	A pool is a set of drives that is logically grouped. You can use a pool to create one or more volumes accessible to a host. (You create volumes from either a pool or a volume group.)
Snapshot consistency group	A snapshot consistency group is a collection of volumes that are treated as a single entity when a snapshot image is created. Each of these volumes has its own snapshot image, but all the images are created at the same point in time.
Snapshot group	A snapshot group is a collection of snapshot images from a single base volume.
Snapshot volume	A snapshot volume allows the host to access data in the snapshot image. The snapshot volume contains its own reserved capacity, which saves any modifications to the base volume without affecting the original snapshot image.
Volume	A volume is a container in which applications, databases, and file systems store data. It is the logical component created for the host to access storage on the storage array.
Volume group	A volume group is a container for volumes with shared characteristics. A volume group has a defined capacity and RAID level. You can use a volume group to create one or more volumes accessible to a host. (You create volumes from either a volume group or a pool.)

Storage capacity

The following terminology describes the different types of capacity used on your storage array.

Capacity type	Description
Allocated capacity	<p>Allocated capacity is the physical capacity allocated from the drives in a pool or volume group.</p> <p>You use allocated capacity to create volumes and for copy services operations.</p>

Capacity type	Description
Free capacity	Free capacity is the capacity available in a pool or volume group that has not yet been allocated to volume creation or copy services operations and storage objects.
Pool or volume group capacity	Pool, volume, or volume group capacity is the capacity in a storage array that has been assigned to a pool or volume group. This capacity is used to create volumes and service the various capacity needs of copy services operations and storage objects.
Pool unusable capacity	Pool unusable capacity is the space in a pool that cannot be used due to mismatched drive sizes.
Preservation capacity	Preservation capacity is the amount of capacity (number of drives) that is reserved in a pool to support potential drive failures.
Reported capacity	Reported capacity is the capacity that is reported to the host and can be accessed by the host.
Reserved capacity	Reserved capacity is the physical allocated capacity that is used for any copy service operation and storage object. It is not directly readable by the host.
SSD Cache	<div>  <p>This feature is not available on the EF600 storage system.</p> </div> <p>SSD Cache is a set of Solid-State Disk (SSD) drives that you logically group together in your storage array. The SSD Cache feature caches the most frequently accessed data ("hot" data) onto lower latency SSD drives to dynamically accelerate application workloads.</p>
Unassigned capacity	Unassigned capacity is the space in a storage array that has not been assigned to a pool or volume group.
Written capacity	Written capacity is the amount of capacity that has been written from the reserved capacity allocated for thin volumes.

How reserved capacity works

Reserved capacity is automatically created when copy service operations, such as snapshots or asynchronous mirroring operations, are provided for your volumes. The purpose of reserved capacity is to store data changes on these volumes, should something go wrong. Like volumes, reserved capacity is created from pools or volume groups.

Copy service objects that use reserved capacity

Reserved capacity is the underlying storage mechanism used by these copy service objects:

- Snapshot groups
- Read/Write snapshot volumes
- Consistency group member volumes
- Mirrored pair volumes

When creating or expanding these copy service objects, you must create new reserved capacity from either a pool or volume group. Reserved capacity is typically 40 percent of the base volume for snapshot operations and 20 percent of the base volume for asynchronous mirroring operations. Reserved capacity, however, varies depending on the number of changes to the original data.

Thin volumes and reserved capacity

For a thin volume, if the maximum reported capacity of 256 TiB has been reached, you cannot increase its capacity. Make sure the thin volume's reserved capacity is set to a size larger than the maximum reported capacity. (A thin volume is always thinly-provisioned, which means that the capacity is allocated as the data is being written to the volume.)

If you create reserved capacity using a thin volume in a pool, review the following actions and results on reserved capacity:

- If a thin volume's reserved capacity fails, the thin volume itself will not automatically transition to the Failed state. However, because all I/O operations on a thin volume require access to the reserved capacity volume, I/O operations will always result in a Check Condition being returned to the requesting host. If the underlying problem with the reserved capacity volume can be resolved, the reserved capacity volume is returned to an Optimal state and the thin volume will become functional again.
- If you use an existing thin volume to complete an asynchronous mirrored pair, that thin volume is re-initialized with a new reserved capacity volume. Only provisioned blocks on the primary side are transferred during the initial synchronization process.

Capacity alerts

The copy service object has a configurable capacity warning and alert threshold, as well as a configurable response when reserved capacity is full.

When the reserved capacity of a copy service object volume is nearing the fill point, an alert is issued to the user. By default, this alert is issued when the reserved capacity volume is 75 percent full; however, you can adjust this alert point up or down as needed. If you receive this alert, you can increase the capacity of the reserved capacity volume at that time. Each copy service object can be configured independently in this regard.

Orphaned reserved capacity volumes

An orphaned reserved capacity volume is a volume that is no longer storing data for copy service operations because its associated copy service object was deleted. When the copy service object was deleted, its reserved capacity volume should have been deleted as well. However, the reserved capacity volume failed to delete.

Because orphaned reserved capacity volumes are not accessed by any host, they are candidates for reclamation. Manually delete the orphaned reserved capacity volume so you can use its capacity for other operations.

System Manager alerts you of orphaned reserved capacity volumes with a **Reclaim unused capacity** message in the Notifications area on the Home page. You can click **Reclaim unused capacity** to display the

Reclaim Unused Capacity dialog box, where you can delete the orphaned reserved capacity volume.

Characteristics of reserved capacity

- Capacity allocated to reserved capacity needs to be considered during the volume creation to retain sufficient free capacity.
- Reserved capacity can be smaller than the base volume (the minimum size is 8 MiB).
- Some space is consumed by metadata, but it is very little (192 KiB), so it does not need to be taken into account when determining the size of reserved capacity volume.
- Reserved capacity is not directly readable or writeable from a host.
- Reserved capacity exists for each read/write snapshot volume, snapshot group, consistency group member volume, and mirrored pair volume.

How SSD Cache works

The SSD Cache feature is a controller-based solution that caches the most frequently accessed data ("hot" data) onto lower latency Solid State Drives (SSDs) to dynamically accelerate system performance. SSD Cache is used exclusively for host reads.



This feature is not available on the EF600 storage system.

SSD Cache versus primary cache

SSD Cache is a secondary cache for use with the primary cache in the controller's dynamic random-access memory (DRAM).

SSD Cache operates differently than primary cache:

- For primary cache, each I/O operation must stage data through the cache to perform the operation.

In primary cache, the data is stored in DRAM after a host read.

- SSD Cache is used only when System Manager determines that it is beneficial to place the data in cache to improve overall system performance.

In SSD Cache, the data is copied from volumes and stored on two internal RAID volumes (one per controller) that are automatically created when you create an SSD Cache.

The internal RAID volumes are used for internal cache processing purposes. These volumes are not accessible or displayed in the user interface. However, these two volumes do count against the total number of volumes allowed in the storage array.

How SSD Cache is used

Intelligent caching places data in a lower-latency drive so responses to future requests for that data can occur much faster. If a program requests data that is in the cache (called a "cache hit"), then the lower-latency drive can service that transaction. Otherwise, a "cache miss" occurs and the data must be accessed from the original, slower drive. As more cache hits occur, overall performance improves.

When a host program accesses the storage array's drives, the data is stored in the SSD Cache. When the same data is accessed by the host program again, it is read from the SSD Cache instead of the hard drives.

The commonly accessed data is stored in the SSD Cache. The hard drives are only accessed when the data cannot be read from the SSD Cache.

SSD Cache is used only when System Manager determines that it is beneficial to place the data in cache to improve overall system performance.

When the CPU needs to process read data, it follows the steps below:

1. Check DRAM cache.
2. If not found in DRAM cache, then check SSD Cache.
3. If not found in SSD Cache, then get from hard drive. If data is deemed worthwhile to cache, then copy to SSD Cache.

Improved performance

Copying the most accessed data (hot spot) to SSD Cache allows for more efficient hard disk operation, reduced latency, and accelerated read and write speeds. Using high performance SSDs to cache data from HDD volumes improves I/O performance and response times.

Simple volume I/O mechanisms are used to move data to and from the SSD Cache. After data is cached and stored on the SSDs, subsequent reads of that data are performed on the SSD Cache, thereby eliminating the need to access the HDD volume.

SSD Cache and the Drive Security feature

To use SSD Cache on a volume that is also using Drive Security (is secure-enabled), the Drive Security capabilities of the volume and the SSD Cache must match. If they do not match, the volume will not be secure-enabled.

Implement SSD Cache

To implement SSD Cache, do the following:

1. Create the SSD Cache.
2. Associate the SSD Cache with the volumes for which you want to implement SSD read caching.



Any volume assigned to use a controller's SSD Cache is not eligible for an automatic load balance transfer.

SSD Cache restrictions

Learn about the restrictions when using SSD Cache on your storage array.

Restrictions

- Any volume assigned to use a controller's SSD Cache is not eligible for an automatic load balance transfer.
- Currently, only one SSD Cache is supported per storage array.
- The maximum usable SSD Cache capacity on a storage array depends on the controller's primary cache capacity.
- SSD Cache is not supported on snapshot images.
- SSD Cache is not available on the EF600 storage system.

- If you import or export volumes that are SSD Cache enabled or disabled, the cached data is not imported or exported.
- You cannot remove the last drive in an SSD Cache without first deleting the SSD Cache.

Restrictions with Drive Security

- You can enable security on SSD Cache only when you create the SSD Cache. You cannot enable security later as you can on a volume.
- If you mix drives that are secure-capable with drives that are not secure-capable in SSD Cache, you cannot enable Drive Security for these drives.
- Secure-enabled volumes must have an SSD Cache that is secure enabled.

Decide whether to use a pool or a volume group

You can create volumes using either a pool or a volume group. The best selection depends primarily on the key storage requirements such as the expected I/O workload, the performance requirements, and the data protection requirements.

Reasons to choose a pool or volume group

Choose a pool

- If you need faster drive rebuilds and simplified storage administration, require thin volumes, and/or have a highly random workload.
- If you want to distribute the data for each volume randomly across a set of drives that comprise the pool.

You cannot set or change the RAID level of pools or the volumes in the pools. Pools use RAID level 6.

Choose a volume group

- If you need maximum system bandwidth, the ability to tune storage settings, and a highly sequential workload.
- If you want to distribute the data across the drives based on a RAID level. You can specify the RAID level when you create the volume group.
- If you want to write the data for each volume sequentially across the set of drives that comprise the volume group.



Because pools can co-exist with volume groups, a storage array can contain both pools and volume groups.

Feature differences between pools and volume groups

The following table provides a feature comparison between volume groups and pools.

Use	Pool	Volume group
Workload random	Better	Good
Workload sequential	Good	Better

Use	Pool	Volume group
Drive rebuild time	Faster	Slower
Performance (optimal mode)	Good: Best for small block, random workload.	Good: Best for large block, sequential workloads
Performance (drive rebuild mode)	Better: Usually better than RAID 6	Degraded: Up to 40% drop in performance
Multiple drive failures	Greater data protection: Faster, prioritized rebuilds	Less data protection: Slow rebuilds, greater risk of data loss
Adding drives	Faster: Add to pool on the fly	Slower: Requires Dynamic Capacity Expansion operation
Thin volumes support	Yes	No
Solid State Disk (SSD) support	Yes	Yes
Simplified administration	Yes: No hot spares or RAID settings to configure	No: Must allocate hot spares, configure RAID
Tunable performance	No	Yes

Functional comparison of pools and volume groups

The function and purpose of a pool and a volume group are the same. Both objects are a set of drives logically grouped together in a storage array and are used to create volumes that a host can access.

The following table helps you decide whether a pool or volume group best suits your storage needs.

Function	Pool	Volume Group
Different RAID level supported	No. Always RAID 6.	Yes. RAID 0, 1, 10, 5, and 6 available.
Thin volumes supported	Yes	No
Full disk encryption (FDE) supported	Yes	Yes
Data Assurance (DA) supported	Yes	Yes
Shelf loss protection supported	Yes	Yes

Function	Pool	Volume Group
Drawer loss protection supported	Yes	Yes
Mixed drive speeds supported	Recommended to be the same, but not required. Slowest drive determines speed for all drives.	Recommended to be the same, but not required. Slowest drive determines speed for all drives.
Mixed drive capacity supported	Recommended to be the same, but not required. Smallest drive determines capacity for all drives.	Recommended to be the same, but not required. Smallest drive determines capacity for all drives.
Minimum number of drives	11	Depends on RAID level. RAID 0 needs 1. RAID 1 or 10 needs 2 (requires an even number). RAID 5 minimum is 3. RAID 6 minimum is 5.
Maximum number of drives	Up to the maximum limit for the storage array	RAID 1 and 10—up to the maximum limit of the storage array RAID 5, 6—30 drives
Can choose individual drives when creating a volume	No	Yes
Can specify segment size when creating a volume	Yes. 128K supported.	Yes
Can specify I/O characteristics when creating a volume	No	Yes. File system, database, multimedia, and custom supported.
Drive failure protection	Uses preservation capacity on each drive in the pool making reconstruction faster.	Uses a hot spare drive. Reconstruction is limited by the IOPs of the drive.
Warning when reaching capacity limit	Yes. Can set an alert when used capacity reaches a percentage of the maximum capacity.	No
Migration to a different storage array supported	No. Requires that you migrate to a volume group first.	Yes
Dynamic Segment Size (DSS)	No	Yes
Can change RAID level	No	Yes
Volume expansion (increase capacity)	Yes	Yes

Function	Pool	Volume Group
Capacity expansion (add capacity)	Yes	Yes
Capacity reduction	Yes	No



Mixed drive types (HDD, SSD) are not supported for either pools or volume groups.

Automatic versus manual pool creation

You create pools automatically or manually to allow physical storage to be grouped, and then dynamically allocated as needed. When a pool is created, you can add physical drives.

Automatic creation

Automatic pool creation is initiated when System Manager detects unassigned capacity in a storage array. When unassigned capacity is detected, System Manager automatically prompts you to create one or more pools, or add the unassigned capacity to an existing pool, or both.

Automatic pool creation occurs when one of these conditions is true:

- Pools do not exist in the storage array, and there are enough similar drives to create a new pool.
- New drives are added to a storage array that has at least one pool.

Each drive in a pool must be of the same drive type (HDD or SSD) and have similar capacity. System Manager will prompt you to complete the following tasks:

- Create a single pool if there are a sufficient number of drives of those types.
- Create multiple pools if the unassigned capacity consists of different drive types.
- Add the drives to the existing pool if a pool is already defined in the storage array, and add new drives of the same drive type to the pool.
- Add the drives of the same drive type to the existing pool, and use the other drive types to create different pools if the new drives are of different drive types.

Manual creation

You might want to create a pool manually when automatic creation cannot determine the best configuration. This situation can occur for one of the following reasons:

- The new drives could potentially be added to more than one pool.
- One or more of the new pool candidates can use shelf loss protection or drawer loss protection.
- One or more of the current pool candidates cannot maintain their shelf loss protection or drawer loss protection status.

You might also want to create a pool manually if you have multiple applications on your storage array and do not want them competing for the same drive resources. In this case, you might consider manually creating a smaller pool for one or more of the applications. You can assign just one or two volumes instead of assigning the workload to a large pool that has many volumes across which to distribute the data. Manually creating a

separate pool that is dedicated to the workload of a specific application can allow storage array operations to perform more rapidly, with less contention.

How to

Create pools and volume groups

Create pool automatically

Pool creation is initiated automatically when System Manager detects unassigned drives in the storage array. You can use automatic pool creation to easily configure all unassigned drives in the storage array into one pool and to add drives into existing pools.

Before you begin

You can launch the **Pool Auto-Configuration** dialog box when one of these conditions are true:

- At least one unassigned drive has been detected that can be added to an existing pool with similar drive types.
- Eleven (11) or more unassigned drives have been detected that can be used to create a new pool (if they cannot be added to an existing pool due to dissimilar drive types).

About this task

Keep in mind the following:

- When you add drives to a storage array, System Manager automatically detects the drives and prompts you to create a single pool or multiple pools based on the drive type and the current configuration.
- If pools were previously defined, System Manager automatically prompts you with the option of adding the compatible drives to an existing pool. When new drives are added to an existing pool, System Manager automatically redistributes the data across the new capacity, which now includes the new drives that you added.
- When configuring an EF600 storage array, make sure each controller has access to an equal number of drives in the first 12 slots and an equal number of drives in the last 12 slots. This configuration helps the controllers use both drive-side PCIe buses more effectively. Currently System Manager allows for drive selection under the Advanced feature when creating a volume group. For pool creation, it is recommended to use all drives in the storage array.

You can launch the **Pool Auto-Configuration** dialog box using any of the following methods:

- When unassigned capacity is detected, the Pool Auto-Configuration recommendation appears on the Home page in the Notification area. Click **View Pool Auto-Configuration** to launch the dialog box.
- You can also launch the **Pool Auto-Configuration** dialog box from the Pools and Volume Groups page as described in the following task.

Steps

1. Select **Storage > Pools & Volume Groups**.
2. Select **More > Launch pool auto-configuration**.

The results table lists new pools, existing pools with drives added, or both. A new pool is named with a sequential number by default.

Notice that System Manager will do the following:

- Create a single pool if there are a sufficient number of drives with the same drive type (HDD or SSD) and have similar capacity.
 - Create multiple pools if the unassigned capacity consists of different drive types.
 - Add the drives to an existing pool if a pool is already defined in the storage array, and you add new drives of the same drive type to the pool.
 - Add the drives of the same drive type to the existing pool, and use the other drive types to create different pools if the new drives are of different drive types.
3. To change the name of a new pool, click the **Edit** icon (the pencil).
 4. To view additional characteristics of the pool, position the cursor over or touch the **Details** icon (the page).

Information about the drive type, security capability, data assurance (DA) capability, shelf loss protection, and drawer loss protection appears.

5. Click **Accept**.

Create pool manually

You can create a pool manually (from a set of candidates) if the Pool Auto Configuration feature does not provide a pool that meets your needs. A pool provides the logical storage capacity necessary from which you can create individual volumes that can then be used to host your applications.

Before you begin

- You must have a minimum of 11 drives with the same drive type (HDD or SSD).
- Shelf loss protection requires that the drives comprising the pool are located in at least six different drive shelves and there are no more than two drives in a single drive shelf.
- Drawer loss protection requires that the drives comprising the pool are located in at least five different drawers and the pool includes an equal number of drive shelves from each drawer.
- When configuring an EF600 storage array, make sure each controller has access to an equal number of drives in the first 12 slots and an equal number of drives in the last 12 slots. This configuration helps the controllers use both drive-side PCIe buses more effectively. Currently System Manager allows for drive selection under the Advanced feature when creating a volume group. For pool creation, it is recommended to use all drives in the storage array.

Steps


1. Select **Storage › Pools & Volume Groups**.
2. Click **Create › Pool**.

The **Create Pool** dialog box appears.

3. Type a name for the pool.
4. **Optional:** If you have more than one type of drive in your storage array, select the drive type that you want to use.

The results table lists all the possible pools that you can create.

5. Select the pool candidate that you want to use based on the following characteristics, and then click **Create**.

Characteristic	Use
Free Capacity	<p>Shows the free capacity of the pool candidate in GiB. Select a pool candidate with the capacity for your application's storage needs.</p> <p>Preservation (spare) capacity is also distributed throughout the pool and is not part of the free capacity amount.</p>
Total Drives	<p>Shows the number of drives available in the pool candidate.</p> <p>System Manager automatically reserves as many drives as possible for preservation capacity (for every six drives in a pool, System Manager reserves one drive for preservation capacity).</p> <p>When a drive failure occurs, the preservation capacity is used to hold the reconstructed data.</p>
Secure-Capable	<p>Indicates whether this pool candidate is comprised entirely of secure-capable drives, which can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.</p> <ul style="list-style-type: none"> • You can protect your pool with Drive Security, but all drives must be secure-capable to use this feature. • If you want to create an FDE-only pool, look for Yes - FDE in the Secure-Capable column. If you want to create a FIPS-only pool, look for Yes - FIPS in the Secure-Capable column. • You can create a pool comprised of drives that may or may not be secure-capable or are a mix of security levels. If the drives in the pool include drives that are not secure-capable, you cannot make the pool secure.
Enable Security?	<p>Provides the option for enabling the Drive Security feature with secure-capable drives. If the pool is secure-capable and you have created a security key, you can enable security by selecting the check box.</p> <div>  <p>The only way to remove Drive Security after it is enabled is to delete the pool and erase the drives.</p> </div>

Characteristic	Use
DA Capable	<p>Indicates if Data Assurance (DA) is available for this pool candidate. DA checks for and corrects errors that might occur as data is transferred through the controllers down to the drives.</p> <p>If you want to use DA, select a pool that is DA capable. This option is available only when the DA feature has been enabled.</p> <p>A pool can contain drives that are DA-capable or not DA-capable, but all drives must be DA capable for you to use this feature.</p>
Shelf Loss Protection	<p>Shows if shelf loss protection is available.</p> <p>Shelf loss protection guarantees accessibility to the data on the volumes in a pool if a total loss of communication occurs with a single drive shelf.</p>
Drawer Loss Protection	<p>Shows if drawer loss protection is available, which is provided only if you are using a drive shelf that contains drawers.</p> <p>Drawer loss protection guarantees accessibility to the data on the volumes in a pool if a total loss of communication occurs with a single drawer in a drive shelf.</p>

Create a volume group

You use a volume group to create one or more volumes that are accessible to the host. A volume group is a container for volumes with shared characteristics such as RAID level and capacity.

About this task

With larger capacity drives and the ability to distribute volumes across controllers, creating more than one volume per volume group is a good way to make use of your storage capacity and to protect your data.

Follow these guidelines when you create a volume group.

- You need at least one unassigned drive.
- Limits exist as to how much drive capacity you can have in a single volume group. These limits vary according to your host type.
- To enable shelf/drawer loss protection, you must create a volume group that uses drives located in at least three shelves or drawers, unless you are using RAID 1, where two shelves/drawers is the minimum.
- When configuring an EF600 storage array, make sure each controller has access to an equal number of drives in the first 12 slots and an equal number of drives in the last 12 slots. This configuration helps the controllers use both drive-side PCIe buses more effectively. Currently System Manager allows for drive selection under the Advanced feature when creating a volume group. For pool creation, it is recommended

to use all drives in the storage array.

Review how your choice of RAID level affects the resulting capacity of the volume group.

- If you select RAID 1, you must add two drives at a time to make sure that a mirrored pair is selected. Mirroring and striping (known as RAID 10 or RAID 1+0) is achieved when four or more drives are selected.
- If you select RAID 5, you must add a minimum of three drives to create the volume group.
- If you select RAID 6, you must add a minimum of five drives to create the volume group.

Steps

1. Select **Storage > Pools & Volume Groups**.
2. Click **Create > Volume group**.

The **Create Volume Group** dialog box appears.

3. Type a name for the volume group.
4. Select the RAID level that best meets your requirements for data storage and protection.

The volume group candidate table appears and displays only the candidates that support the selected RAID level.

5. **Optional:** If you have more than one type of drive in your storage array, select the drive type that you want to use.

The volume group candidate table appears and displays only the candidates that support the selected drive type and RAID level.

6. **Optional:** You can select either the automatic method or manual method to define which drives to use in the volume group. The Automatic method is the default selection.

To select drives manually, click the **Manually select drives (advanced)** link. When clicked, it changes to **Automatically select drives (advanced)**.

The Manual method lets you select which specific drives comprise the volume group. You can select specific unassigned drives to obtain the capacity that you require. If the storage array contains drives with different media types or different interface types, you can choose only the unconfigured capacity for a single drive type to create the new volume group.




Only experts who understand drive redundancy and optimal drive configurations should use the Manual method.

7. Based on the displayed drive characteristics, select the drives you want to use in the volume group, and then click **Create**.

The drive characteristics displayed depend on whether you selected the automatic method or manual method.

Automatic method drive characteristics

Characteristic	Use
Free Capacity	Shows the available capacity in GiB. Select a volume group candidate with the capacity for your application's storage needs.
Total Drives	Shows the number of drives available for this volume group. Select a volume group candidate with the number of drives that you want. The more drives that a volume group contains, the less likely it is that multiple drive failures will cause a critical drive failure in a volume group.
Secure-Capable	<p>Indicates whether this volume group candidate is comprised entirely of secure-capable drives, which can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.</p> <ul style="list-style-type: none"> • You can protect your volume group with Drive Security, but all drives must be secure-capable to use this feature. • If you want to create an FDE-only volume group, look for Yes - FDE in the Secure-Capable column. If you want to create a FIPS-only volume group, look for Yes - FIPS in the Secure-Capable column. • You can create a volume group comprised of drives that might or might not be secure-capable or are a mix of security levels. If the drives in the volume group include drives that are not secure-capable, you cannot make the volume group secure.
Enable Security?	<p>Provides the option for enabling the Drive Security feature with secure-capable drives. If the volume group is secure-capable and you have set up a security key, you can enable Drive Security by selecting the check box.</p> <div>  <p>The only way to remove Drive Security after it is enabled is to delete the volume group and erase the drives.</p> </div>

Characteristic	Use
DA Capable	<p>Indicates if Data Assurance (DA) is available for this group. Data Assurance (DA) checks for and corrects errors that might occur as data is transferred through the controllers down to the drives.</p> <p>If you want to use DA, select a volume group that is DA capable. This option is available only when the DA feature has been enabled.</p> <p>A volume group can contain drives that are DA-capable or not DA-capable, but all drives must be DA capable for you to use this feature.</p>
Shelf Loss Protection	<p>Shows if shelf loss protection is available. Shelf loss protection guarantees accessibility to the data on the volumes in a volume group if a total loss of communication to a shelf occurs.</p>
Drawer Loss Protection	<p>Shows if drawer loss protection is available, which is provided only if you are using a drive shelf that contains drawers. Drawer loss protection guarantees accessibility to the data on the volumes in a volume group if a total loss of communication occurs with a single drawer in a drive shelf.</p>

Manual method drive characteristics

Characteristic	Use
Media Type	<p>Indicates the media type. The following media types are supported:</p> <ul style="list-style-type: none">• Hard drive• Solid State Disk (SSD) 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.
Drive Capacity	<p>Indicates the drive capacity.</p> <ul style="list-style-type: none">• Whenever possible, select drives that have a capacity equal to the capacities of the current drives in the volume group.• If you must add unassigned drives with a smaller capacity, be aware that the usable capacity of each drive currently in the volume group is reduced. Therefore, the drive capacity is the same across the volume group.• If you must add unassigned drives with a larger capacity, be aware that the usable capacity of the unassigned drives that you add is reduced so that they match the current capacities of the drives in the volume group.
Tray	Indicates the tray location of the drive.
Slot	Indicates the slot location of the drive.
Speed (rpm)	Indicates the speed of the drive.
Logical sector size	Indicates the sector size and format.

Characteristic	Use
Secure-Capable	<p>Indicates whether this volume group candidate is comprised entirely of secure-capable drives, which can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.</p> <ul style="list-style-type: none"> • You can protect your volume group with Drive Security, but all drives must be secure-capable to use this feature. • If you want to create an FDE-only volume group, look for Yes - FDE in the Secure-Capable column. If you want to create a FIPS-only volume group, look for Yes - FIPS in the Secure-Capable column. • You can create a volume group comprised of drives that might or might not be secure-capable or are a mix of security levels. If the drives in the volume group include drives that are not secure-capable, you cannot make the volume group secure.
DA Capable	<p>Indicates if Data Assurance (DA) is available for this group. Data Assurance (DA) checks for and corrects errors that might occur as data is communicated through the controllers down to the drives.</p> <p>If you want to use DA, select a volume group that is DA capable. This option is available only when the DA feature has been enabled.</p> <p>A volume group can contain drives that are DA-capable or not DA-capable, but all drives must be DA capable for you to use this feature.</p>

Create SSD Cache

To dynamically accelerate system performance, you can use the SSD Cache feature to cache the most frequently accessed data ("hot" data) onto lower latency Solid State Drives (SSDs). SSD Cache is used exclusively for host reads.

Before you begin

Your storage array must contain some SSD drives.



SSD Cache is not available on the EF600 storage system.

About this task

When you create SSD Cache, you can use a single drive or multiple drives. Because the read cache is in the storage array, caching is shared across all applications using the storage array. You select the volumes that you want to cache, and then caching is automatic and dynamic.

Follow these guidelines when you create SSD Cache.

- You can enable security on the SSD Cache only when you are creating it, not later.
- Only one SSD Cache is supported per storage array.
- The maximum usable SSD Cache capacity on a storage array is dependent on the controller’s primary cache capacity.
- SSD Cache is not supported on snapshot images.
- If you import or export volumes that are SSD Cache enabled or disabled, the cached data is not imported or exported.
- Any volume assigned to use a controller’s SSD Cache is not eligible for an automatic load balance transfer.
- If the associated volumes are secure-enabled, create a secure-enabled SSD Cache.


Steps

1. Select **Storage > Pools & Volume Groups**.
2. Click **Create > SSD Cache**.

The Create **SSD Cache** dialog box appears.

3. Type a name for the SSD Cache.
4. Select the SSD Cache candidate that you want to use based on the following characteristics.

Characteristic	Use
Capacity	<p>Shows the available capacity in GiB. Select the capacity for your application’s storage needs.</p> <p>The maximum capacity for SSD Cache depends on the controller’s primary cache capacity. If you allocate more than the maximum amount to SSD Cache, then any extra capacity is unusable.</p> <p>SSD Cache capacity counts towards your overall allocated capacity.</p>
Total drives	<p>Shows the number of drives available for this SSD cache. Select the SSD candidate with the number of drives that you want.</p>

Characteristic	Use
Secure-capable	<p>Indicates whether the SSD Cache candidate is comprised entirely of secure-capable drives, which can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.</p> <p>If you want to create a secure-enabled SSD Cache, look for Yes - FDE or Yes - FIPS in the Secure-capable column.</p>
Enable security?	<p>Provides the option for enabling the Drive Security feature with secure-capable drives. If you want to create a secure-enabled SSD Cache, select the Enable Security check box.</p> <div>  <p>Once enabled, security cannot be disabled. You can enable security on the SSD Cache only when you are creating it, not later.</p> </div>
DA capable	<p>Indicates if Data Assurance (DA) is available for this SSD Cache candidate. Data Assurance (DA) checks for and corrects errors that might occur as data is transferred through the controllers down to the drives.</p> <p>If you want to use DA, select an SSD Cache candidate that is DA capable. This option is available only when the DA feature has been enabled.</p> <p>SSD Cache can contain both DA-capable and non-DA-capable drives, but all drives must be DA-capable for you to use DA.</p>

- Associate the SSD Cache with the volumes for which you want to implement SSD read caching. To enable SSD Cache on compatible volumes immediately, select the **Enable SSD Cache on existing compatible volumes that are mapped to hosts** check box.

Volumes are compatible if they share the same Drive Security and DA capabilities.

- Click **Create**.

Add capacity to a pool or volume group

You can add drives to expand the free capacity in an existing pool or volume group. The expansion causes additional free capacity to be included in the pool or volume group. You can use this free capacity to create additional volumes. The data in the volumes remains accessible during this operation.

Before you begin

- Drives must be in an Optimal status.
- Drives must have the same drive type (HDD or SSD).
- The pool or volume group must be in an Optimal status.
- If the pool or volume group contains all secure-capable drives, add only drives that are secure-capable to continue to use the encryption abilities of the secure-capable drives.

Secure-capable drives can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.

About this task

For pools, you can add a maximum of 60 drives at a time or up to 60 drives through multiples of 5. For volume groups, you can add a maximum of two drives at a time. If you need to add more than the maximum number of drives, repeat the procedure. (A pool cannot contain more drives than the maximum limit for a storage array.)



With the addition of drives, your preservation capacity may need to be increased. You should consider increasing your reserved capacity after an expansion operation.



Avoid using drives that are Data Assurance (DA) capable to add capacity to a pool or volume group that is not DA capable. The pool or volume group cannot take advantage of the capabilities of the DA-capable drive. Consider using drives that are not DA capable in this situation.

Steps

1. Select **Storage > Pools & Volume Groups**.
2. Select the pool or volume group to which you want to add drives, and then click **Add Capacity**.

The **Add Capacity** dialog box appears. Only the unassigned drives that are compatible with the pool or volume group appear.

3. Under **Select drives to add capacity...**, select one or more drives that you want to add to the existing pool or volume group.

The controller firmware arranges the unassigned drives with the best options listed at the top. The total free capacity that is added to the pool or volume group appears below the list in **Total capacity selected**.

Field Details

Field	Description
Shelf	Indicates the shelf location of the drive.
Bay	Indicates the bay location of the drive.
Capacity (GiB)	<p>Indicates the drive capacity.</p> <ul style="list-style-type: none">• Whenever possible, select drives that have a capacity equal to the capacities of the current drives in the pool or volume group.• If you must add unassigned drives with a smaller capacity, be aware that the usable capacity of each drive currently in the pool or volume group is reduced. Therefore, the drive capacity is the same across the pool or volume group.• If you must add unassigned drives with a larger capacity, be aware that the usable capacity of the unassigned drives that you add is reduced so that they match the current capacities of the drives in the pool or volume group.
Secure-Capable	<p>Indicates whether the drive is secure-capable.</p> <ul style="list-style-type: none">• You can protect your pool or volume group with the Drive Security feature, but all drives must be secure-capable to use this feature.• You can mix secure-capable and non-secure-capable drives, but the encryption abilities of the secure-capable drives cannot be used.• Secure-capable drives can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.
DA Capable	<p>Indicates whether the drive is Data Assurance (DA) capable.</p> <ul style="list-style-type: none">• Using drives that are not Data Assurance (DA) capable to add capacity to a DA-capable pool or volume group is not recommended. The pool or volume group no longer has DA capabilities, and you no longer have the option to enable DA on newly created volumes within the pool or volume group.• Using drives that are Data Assurance (DA) capable to add capacity to a pool or volume group that is non DA-capable is not recommended, because that pool or volume group cannot take advantage of the capabilities of the DA-capable drive (the drive attributes do not match). Consider using drives that are not DA-capable in this situation.
DULBE capable	<p>Indicates whether the drive has the option for Deallocated or Unwritten Logical Block Error (DULBE). DULBE is an option on NVMe drives that allows an EF600 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.</p>

4. Click **Add**.

If you are adding drives to a pool or volume group, a confirmation dialog box appears if you selected a drive that causes the pool or volume group to no longer have one or more of the following attributes:

- Shelf loss protection
- Drawer loss protection
- Full Disk Encryption capability
- Data Assurance capability
- DULBE capability

To continue, click **Yes**; otherwise click **Cancel**.

Results

After you add the unassigned drives to a pool or volume group, the data in each volume of the pool or volume group is redistributed to include the additional drives.

Manage pools, volume groups, and SSD Cache

Change configuration settings for a pool

You can edit the settings for a pool, including its name, capacity alerts settings, modification priorities, and preservation capacity.

About this task

This task describes how to change configuration settings for a pool.



You cannot change the RAID level of a pool. System Manager automatically configures pools as RAID 6.

Steps

1. Select **Storage > Pools & Volume Groups**.
2. Select the pool that you want to edit, and then click **View/Edit Settings**.

The **Pool Settings** dialog box appears.

3. Select the **Settings** tab, and then edit the pool settings as appropriate.

Field Details

Setting	Description
Name	You can change the user-supplied name of the pool. Specifying a name for a pool is required.
Capacity alerts	<p>You can send alert notifications when the free capacity in a pool reaches or exceeds a specified threshold. When the data stored in the pool exceeds the specified threshold, System Manager sends a message, allowing you time to add more storage space or to delete unnecessary objects.</p> <p>Alerts are shown in the Notifications area on the Dashboard and can be sent from the server to administrators by email and SNMP trap messages.</p> <p>You can define the following capacity alerts:</p> <ul style="list-style-type: none"> • Critical alert — This critical alert notifies you when the free capacity in the pool reaches or exceeds the specified threshold. Use the spinner controls to adjust the threshold percentage. Select the check box to disable this notification. • Early alert — This early alert notifies you when the free capacity in a pool is reaching a specified threshold. Use the spinner controls to adjust the threshold percentage. Select the check box to disable this notification.
Modification priorities	<p>You can specify the priority levels for modification operations in a pool relative to system performance. A higher priority for modification operations in a pool causes an operation to complete faster, but can slow the host I/O performance. A lower priority causes operations to take longer, but host I/O performance is less affected.</p> <p>You can choose from five priority levels: lowest, low, medium, high, and highest. The higher the priority level, the larger is the impact on host I/O and system performance.</p> <ul style="list-style-type: none"> • Critical reconstruction priority — This slider bar determines the priority of a data reconstruction operation when multiple drive failures result in a condition where some data has no redundancy and an additional drive failure might result in loss of data. • Degraded reconstruction priority — This slider bar determines the priority of the data reconstruction operation when a drive failure has occurred, but the data still has redundancy and an additional drive failure does not result in loss of data. • Background operation priority — This slider bar determines the priority of the pool background operations that occur while the pool is in an optimal state. These operations include Dynamic Volume Expansion (DVE), Instant Availability Format (IAF), and migrating data to a replaced or added drive.

Setting	Description
Preservation capacity ("Reserve capacity" for the EF600)	<p>Preservation capacity: You can define the number of drives to determine the capacity that is reserved on the pool to support potential drive failures. When a drive failure occurs, the preservation capacity is used to hold the reconstructed data. Pools use preservation capacity during the data reconstruction process instead of hot spare drives, which are used in volume groups.</p> <p>Use the spinner controls to adjust the number of drives. Based on the number of drives, the preservation capacity in the pool appears next to the spinner box.</p> <p>Keep the following information in mind about preservation capacity.</p> <ul style="list-style-type: none"> • Because preservation capacity is subtracted from the total free capacity of a pool, the amount of capacity that you reserve affects how much free capacity is available to create volumes. If you specify 0 for the preservation capacity, all of the free capacity on the pool is used for volume creation. • If you decrease the preservation capacity, you increase the capacity that can be used for pool volumes. <p>Additional optimization capacity (EF600 arrays only): When a pool is created, a recommended optimization capacity is generated that provides a balance of available capacity versus performance and drive wear life. You can adjust this balance by moving the slider to the right for better performance and drive wear life at the expense of increased available capacity, or by moving it to the left for increased available capacity at the expense of better performance and drive wear life.</p> <p>SSD drives will have longer life and better maximum write performance when a portion of their capacity is unallocated. For drives associated with a pool, unallocated capacity is comprised of a pool's preservation capacity, the free capacity (capacity not used by volumes), and a portion of the usable capacity set aside as additional optimization capacity. The additional optimization capacity ensures a minimum level of optimization capacity by reducing the usable capacity, and as such, is not available for volume creation.</p>

4. Click **Save**.

Change SSD Cache settings

You can edit the name of the SSD Cache and view its status, maximum and current capacity, Drive Security and Data Assurance status, and its associated volumes and drives.

Steps

1. Select **Storage > Pools & Volume Groups**.
2. Select the SSD Cache that you want to edit, and then click **View/Edit Settings**.

The **SSD Cache Settings** dialog box appears.

3. Review or edit the SSD Cache settings as appropriate.

Field Details

Setting	Description
Name	Displays the name of the SSD Cache, which you can change. A name for the SSD Cache is required.
Characteristics	Shows the status for the SSD Cache. Possible statuses include: <ul style="list-style-type: none">• Optimal• Unknown• Degraded• Failed (A failed state results in a critical MEL event.)• Suspended
Capacities	Shows the current capacity and maximum capacity allowed for the SSD Cache. The maximum capacity allowed for the SSD Cache depends on the controller's primary cache size: <ul style="list-style-type: none">• Up to 1 GiB• 1 GiB to 2 GiB• 2 GiB to 4 GiB• More than 4 GiB
Security and DA	Shows the Drive Security and Data Assurance status for the SSD Cache. <ul style="list-style-type: none">• Secure-capable — Indicates whether the SSD Cache is comprised entirely of secure-capable drives. A secure-capable drive is a self-encrypting drive that can protect its data from unauthorized access.• Secure-enabled — Indicates whether security is enabled on the SSD Cache.• DA capable — Indicates whether the SSD Cache is comprised entirely of DA-capable drives. A DA-capable drive can check for and correct errors that might occur as data is communicated between the host and storage array.
Associated objects	Shows the volumes and drives associated with the SSD Cache.

4. Click **Save**.

Change configuration settings for a volume group

You can edit the settings for a volume group, including its name and RAID level.

Before you begin

If you are changing the RAID level to accommodate the performance needs of the applications that are accessing the volume group, be sure to meet the following prerequisites:

- The volume group must be in Optimal status.
- You must have enough capacity in the volume group to convert to the new RAID level.

Steps

1. Select **Storage > Pools & Volume Groups**.
2. Select the volume group that you want to edit, and then click **View/Edit Settings**.

The **Volume Group Settings** dialog box appears.

3. Select the **Settings** tab, and then edit the volume group settings as appropriate.

Field Details

Setting	Description
Name	You can change the user-supplied name of the volume group. Specifying a name for a volume group is required.
RAID level	<p>Select the new RAID level from the drop-down menu.</p> <ul style="list-style-type: none"> • RAID 0 striping. Offers high performance, but does not provide any data redundancy. If a single drive fails in the volume group, all of the associated volumes fail, and all data is lost. A striping RAID group combines two or more drives into one large, logical drive. • RAID 1 mirroring. Offers high performance and the best data availability, and is suitable for storing sensitive data on a corporate or personal level. Protects your data by automatically mirroring the contents of one drive to the second drive in the mirrored pair. It provides protection in the event of a single drive failure. • RAID 10 striping/mirroring. Provides a combination of RAID 0 (striping) and RAID 1 (mirroring), and is achieved when four or more drives are selected. RAID 10 is suitable for high volume transaction applications, such as a database, that require high performance and fault tolerance. • RAID 5. Optimal for multi-user environments (such as database or file system storage) where typical I/O size is small and there is a high proportion of read activity. • RAID 6. Optimal for environments requiring redundancy protection beyond RAID 5, but not requiring high write performance. <p>RAID 3 can be assigned only to volume groups using the command line interface (CLI).</p> <p>When you change the RAID level, you cannot cancel this operation after it begins. During the change, your data remains available.</p>
Optimization capacity (EF600 arrays only)	<p>When a volume group is created, a recommended optimization capacity is generated that provides a balance of available capacity versus performance and drive wear life. You can adjust this balance by moving the slider to the right for better performance and drive wear life at the expense of increased available capacity, or by moving it to the left for increased available capacity at the expense of better performance and drive wear life.</p> <p>SSD drives will have longer life and better maximum write performance when a portion of their capacity is unallocated. For drives associated with a volume group, unallocated capacity is comprised of a group's free capacity (capacity not used by volumes) and a portion of the usable capacity set aside as additional optimization capacity. The additional optimization capacity ensures a minimum level of optimization capacity by reducing the usable capacity, and as such, is not available for volume creation.</p>

4. Click **Save**.

A confirmation dialog box appears if capacity is reduced, volume redundancy is lost, or shelf/drawer loss protection is lost as a result of the RAID level change. Select **Yes** to continue; otherwise click **No**.

Results

If you change the RAID level for a volume group, System Manager changes the RAID levels of every volume that comprises the volume group. Performance might be slightly affected during the operation.

View SSD Cache statistics

You can view statistics for the SSD Cache, such as reads, writes, cache hits, cache allocation percentage, and cache utilization percentage.

About this task



This feature is not available on the EF600 storage system.

The nominal statistics, which are a subset of the detailed statistics, are shown on the **View SSD Cache Statistics** dialog box. You can view detailed statistics for the SSD Cache only when you export all SSD statistics to a `.csv` file.

As you review and interpret the statistics, keep in mind that some interpretations are derived by looking at a combination of statistics.

Steps

1. Select **Storage > Pools & Volume Groups**.
2. Select the SSD Cache for which you want to view statistics, and then click **More > View SSD Cache statistics**.

The **View SSD Cache Statistics** dialog box appears and displays the nominal statistics for the selected SSD cache.

Field Details

Settings	Description
Reads	Shows the total number of host reads from the SSD Cache-enabled volumes. The greater the ratio of Reads to Writes, the better is the operation of the cache.
Writes	The total number of host writes to the SSD Cache-enabled volumes. The greater the ratio of Reads to Writes, the better is the operation of the cache.
Cache hits	Shows the number of cache hits.
Cache hits %	Shows the percentage of cache hits. This number is derived from Cache Hits / (reads + writes). The cache hit percentage should be greater than 50 percent for effective SSD Cache operation.
Cache allocation %	Shows the percentage of SSD Cache storage that is allocated, expressed as a percentage of the SSD Cache storage that is available to this controller and is derived from allocated bytes / available bytes.
Cache utilization %	Shows the percentage of SSD Cache storage that contains data from enabled volumes, expressed as a percentage of SSD Cache storage that is allocated. This amount represents the utilization or density of the SSD Cache. Derived from allocated bytes / available bytes.
Export All	Exports all SSD Cache statistics to a CSV format. The exported file contains all available statistics for the SSD Cache (both nominal and detailed).

3. Click **Cancel** to close the dialog box.

Check volume redundancy

Under the guidance of technical support or as instructed by the Recovery Guru, you can check the redundancy on a volume in a pool or volume group to determine whether the data on that volume is consistent. Redundancy data is used to quickly reconstruct information on a replacement drive if one of the drives in the pool or volume group fails.

Before you begin

- The status of the pool or volume group must be Optimal.
- The pool or volume group must have no volume modification operations in progress.
- You can check redundancy on any RAID level except on RAID 0, because RAID 0 has no data redundancy. (Pools are configured only as RAID 6.)



Check volume redundancy only when instructed to do so by the Recovery Guru and under the guidance of technical support.

About this task

You can perform this check only on one pool or volume group at a time. A volume redundancy check performs the following actions:

- Scans the data blocks in a RAID 3 volume, a RAID 5 volume, or a RAID 6 volume, and checks the redundancy information for each block. (RAID 3 can only be assigned to volume groups using the command line interface.)
- Compares the data blocks on RAID 1 mirrored drives.
- Returns redundancy errors if the controller firmware determines that the data is inconsistent.



Immediately running a redundancy check on the same pool or volume group might cause an error. To avoid this problem, wait one to two minutes before running another redundancy check on the same pool or volume group.

Steps

1. Select **Storage > Pools & Volume Groups**.
2. Select **Uncommon Tasks > Check volume redundancy**.

The **Check Redundancy** dialog box appears.

3. Select the volumes you want to check, and then type `check` to confirm you want to perform this operation.
4. Click **Check**.

The check volume redundancy operation starts. The volumes in the pool or volume group are sequentially scanned, starting from the top of the table in the dialog box. These actions occur as each volume is scanned:

- The volume is selected in the volume table.
- The status of the redundancy check is shown in the Status column.
- The check stops on any media or parity error encountered, and then reports the error.

More about the status of the redundancy check

Status	Description
Pending	This is the first volume to be scanned, and you have not clicked Start to start the redundancy check. or The redundancy check operation is being performed on other volumes in the pool or volume group.
Checking	The volume is undergoing the redundancy check.
Passed	The volume passed the redundancy check. No inconsistencies were detected in the redundancy information.
Failed	The volume failed the redundancy check. Inconsistencies were detected in the redundancy information.
Media error	The drive media is defective and is unreadable. Follow the instructions displayed in the Recovery Guru.
Parity error	The parity is not what it should be for a given portion of the data. A parity error is potentially serious and could cause a permanent loss of data.

5. Click **Done** after the last volume in the pool or volume group has been checked.

Delete pool or volume group

You can delete a pool or volume group to create more unassigned capacity, which you can reconfigure to meet your application storage needs.

Before you begin

- You must have backed up the data on all of the volumes in the pool or volume group.
- You must have stopped all input/output (I/O).
- You must unmount any file systems on the volumes.
- You must have deleted any mirror relationships in the pool or volume group.
- You must have stopped any volume copy operation in progress for the pool or volume group.
- The pool or volume group must not be participating in an asynchronous mirroring operation.
- The drives in the volume group must not have a persistent reservation.

Steps

1. Select **Storage › Pools & Volume Groups**.
2. Select one pool or volume group from the list.

You can select only one pool or volume group at a time. Scroll down the list to see additional pools or volume groups.

3. Select **Uncommon Tasks** > **Delete** and confirm.

Results

System Manager performs the following actions:

- Deletes all of the data in the pool or volume group.
- Deletes all the drives associated with the pool or volume group.
- Unassigns the associated drives, which allows you to reuse them in new or existing pools or volume groups.

Consolidate free capacity for a volume group

Use the Consolidate Free Capacity option to consolidate existing free extents on a selected volume group. By performing this action, you can create additional volumes from the maximum amount of free capacity in a volume group.

Before you begin

- The volume group must contain at least one free capacity area.
- All of the volumes in the volume group must be online and in Optimal status.
- Volume modification operations must not be in progress, such as changing the segment size of a volume.

About this task

You cannot cancel the operation after it begins. Your data remains accessible during the consolidation operation.

You can launch the **Consolidate Free Capacity** dialog box using any of the following methods:

- When at least one free capacity area is detected for a volume group, the **Consolidate free capacity** recommendation appears on the Home page in the Notification area. Click the **Consolidate free capacity** link to launch the dialog box.
- You can also launch the Consolidate Free Capacity dialog box from the Pools & Volume Groups page as described in the following task.

More about free capacity areas

A free capacity area is the free capacity that can result from deleting a volume or from not using all available free capacity during volume creation. When you create a volume in a volume group that has one or more free capacity areas, the volume's capacity is limited to the largest free capacity area in that volume group. For example, if a volume group has a total of 15 GiB free capacity, and the largest free capacity area is 10 GiB, the largest volume you can create is 10 GiB.

You consolidate free capacity on a volume group to improve write performance. Your volume group's free capacity will become fragmented over time as the host writes, modifies, and deletes files. Eventually, the available capacity will not be located in a single contiguous block, but will be scattered in small fragments across the volume group. This causes further file fragmentation, since the host must write new files as fragments to fit them into the available ranges of free clusters.

By consolidating free capacity on a selected volume group, you will notice improved file system performance whenever the host writes new files. The consolidation process will also help prevent new files from being fragmented in the future.

Steps

1. Select **Storage › Pools & Volume Groups**.
2. Select the volume group with free capacity that you want to consolidate, and then select **Uncommon Tasks › Consolidate volume group free capacity**.

The **Consolidate Free Capacity** dialog box appears.

3. Type `consolidate` to confirm you want to perform this operation.
4. Click **Consolidate**.

Results

System Manager begins consolidating (defragmenting) the volume group's free capacity areas into one contiguous amount for subsequent storage configuration tasks.

After you finish

Select **Home › View Operations in Progress** to view the progress of the Consolidate Free Capacity operation. This operation can be lengthy and could affect system performance.

Export/Import volume groups

Volume group migration lets you export a volume group so that you can import the volume group to a different storage array.

The Export/Import function is not supported in the SANtricity System Manager user interface. You must use the Command Line Interface (CLI) to export/import a volume group to a different storage array.

Manage drives

Turn on locator lights in a pool, volume group, or SSD Cache

You can locate drives to physically identify all of the drives that comprise a selected pool, volume group, or SSD Cache. An LED indicator lights up on each drive in the selected

pool, volume group, or SSD Cache.

Steps

1. Select **Storage > Pools & Volume Groups**.
2. Select the pool, volume group, or SSD Cache you want to locate, and then click **More > Turn on locator lights**.

A dialog box appears that indicates the lights on the drives comprising the selected pool, volume group, or SSD Cache are turned on.

3. After you successfully locate the drives, click **Turn Off**.

Remove capacity from a pool or SSD Cache

You can remove drives to decrease the capacity of an existing pool or SSD Cache. After you remove drives, the data in each volume of the pool or SSD Cache is redistributed to the remaining drives. The removed drives become unassigned and their capacity becomes part of the total free capacity of the storage array.

About this task

Follow these guidelines when you remove capacity:

- You cannot remove the last drive in an SSD Cache without first deleting the SSD Cache.
- You cannot reduce the number of drives in a pool to be less than 11 drives.
- You can remove a maximum of 12 drives at a time. If you need to remove more than 12 drives, repeat the procedure.
- You cannot remove drives if there is not enough free capacity in the pool or SSD Cache to contain the data, when that data is redistributed to the remaining drives in the pool or SSD Cache.

Read about potential performance impacts

- Removing drives from a pool or SSD Cache might result in reduced volume performance.
- The preservation capacity is not consumed when you remove capacity from a pool or SSD Cache. However, the preservation capacity might decrease based on the number of drives remaining in the pool or SSD Cache.

Read about impacts to secure-capable drives

- If you remove the last drive that is not secure-capable, the pool is left with all secure-capable drives. In this situation, you are given the option to enable security for the pool.
- If you remove the last drive that is not Data Assurance (DA) capable, the pool is left with all DA-capable drives.

Any new volumes that you create on the pool will be DA-capable. If you want existing volumes to be DA-capable, you need to delete and then re-create the volume.

Steps

1. Select **Storage > Pools & Volume Groups**.

2. Select the pool or SSD Cache, and then click **More › Remove capacity**.

The **Remove Capacity** dialog box appears.

3. Select one or more drives in the list.

As you select or de-select drives in the list, the **Total capacity selected** field updates. This field shows the total capacity of the pool or SSD Cache that results after you remove the selected drives.

4. Click **Remove**, and then confirm you want to remove the drives.

Results

The newly reduced capacity of the pool or SSD Cache is reflected in the Pools and Volume Groups view.

Enable security for a pool or volume group

You can enable Drive Security for a pool or volume group to prevent unauthorized access to the data on the drives contained in the pool or volume group. Read and write access for the drives is only available through a controller that is configured with a security key.

Before you begin

- The Drive Security feature must be enabled.
- A security key must be created.
- The pool or volume group must be in an Optimal state.
- All of the drives in the pool or volume group must be secure-capable drives.

About this task

If you want to use Drive Security, select a pool or volume group that is secure-capable. A pool or volume group can contain both secure-capable and non-secure-capable drives, but all drives must be secure-capable to use their encryption capabilities.

After enabling security, you can only remove it by deleting the pool or volume group, and then erasing the drives.

Steps

1. Select **Storage › Pools & Volume Groups**.
2. Select the pool or volume group on which you want to enable security, and then click **More › Enable security**.

The **Confirm Enable Security** dialog box appears.

3. Confirm that you want to enable security for the selected pool or volume group, and then click **Enable**.

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.

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.

Manage reserved capacity

Increase reserved capacity

You can increase reserved capacity, which is the physically allocated capacity used for any copy service operation on a storage object. For snapshot operations, it is typically 40 percent of the base volume; for asynchronous mirroring operations, it is typically 20 percent of the base volume. Typically, you increase reserved capacity when you receive a warning that the storage object's reserved capacity is becoming full.

Before you begin

- The volume in the pool or volume group must have an Optimal status and must not be in any state of modification.
- Free capacity must exist in the pool or volume group that you want to use to increase capacity.

If no free capacity exists on any pool or volume group, you can add unassigned capacity in the form of unused drives to a pool or volume group.

About this task

You can increase reserved capacity only in increments of 8 GiB for the following storage objects:

- Snapshot group
- Snapshot volume
- Consistency group member volume
- Mirrored pair volume

Use a high percentage if you believe the primary volume will undergo many changes or if the lifespan of a particular copy service operation will be very long.



You cannot increase reserved capacity for a snapshot volume that is read-only. Only snapshot volumes that are read-write require reserved capacity.

Steps

1. Select **Storage > Pools & Volume Groups**.
2. Select the **Reserved Capacity** tab.
3. Select the storage object for which you want to increase reserved capacity, and then click **Increase Capacity**.

The **Increase Reserved Capacity** dialog box appears.

4. Use the spinner box to adjust the capacity percentage.

If free capacity does not exist on the pool or volume group that contains the storage object you selected, and the storage array has Unassigned Capacity, you can create a new pool or volume group. You can then retry this operation using the new free capacity on that pool or volume group.

5. Click **Increase**.

Results

System Manager performs the following actions:

- Increases the reserved capacity for the storage object.
- Displays the newly-added reserved capacity.

Decrease reserved capacity

You use the Decrease Capacity option to decrease the reserved capacity for the following storage objects: snapshot group, snapshot volume, and consistency group member volume. You can decrease reserved capacity only by the amount(s) you used to increase it.

Before you begin

- The storage object must contain more than one reserved capacity volume.
- The storage object must not be a mirrored pair volume.
- If the storage object is a snapshot volume, then it must be a disabled snapshot volume.
- If the storage object is a snapshot group, then it must not contain any associated snapshot images.

About this task

Review the following guidelines:

- You can remove reserved capacity volumes only in the reverse order that they were added.
- You cannot decrease the reserved capacity for a snapshot volume that is read-only because it does not have any associated reserved capacity. Only snapshot volumes that are read-write require reserved capacity.

Steps

1. Select **Storage > Pools & Volume Groups**.
2. Click the **Reserved Capacity** tab.
3. Select the storage object for which you want to decrease reserved capacity, and then click **Decrease Capacity**.

The **Decrease Reserved Capacity** dialog box appears.

4. Select the capacity amount by which you want to decrease reserved capacity, and then click **Decrease**.

Results

System Manager performs the following actions:

- Updates the capacity for the storage object.
- Displays the newly updated reserved capacity for the storage object.
- When you decrease capacity for a snapshot volume, System Manager automatically transitions the snapshot volume to a Disabled state. Disabled means that the snapshot volume is not currently associated with a snapshot image, and therefore, cannot be assigned to a host for I/O.

Change the reserved capacity settings for a snapshot group

You can change the settings for a snapshot group to change its name, auto-delete

settings, the maximum number of allowed snapshot images, the percentage point at which System Manager sends a reserved capacity alert notification, or the policy to use when the reserved capacity reaches its maximum defined percentage.

Before you begin

During the creation of a snapshot group, reserved capacity is created to store the data for all the snapshot images contained in the group.

Steps

1. Select **Storage > Pools & Volume Groups**.
2. Click the **Reserved Capacity** tab.
3. Select the snapshot group that you want to edit, and then click **View/Edit Settings**.

The **Snapshot Group Settings** dialog box appears.

4. Change the settings for the snapshot group as appropriate.

Field Details

Setting	Description
Snapshot group settings	
Name	The name of the snapshot group. Specifying a name for the snapshot group is required.
Auto-deletion	A setting that keeps the total number of snapshot images in the group at or below a user-defined maximum. When this option is enabled, System Manager automatically deletes the oldest snapshot image in the group any time a new snapshot is created, to comply with the maximum number of snapshot images allowed for the group.
Snapshot image limit	A configurable value that specifies the maximum number of snapshot images allowed for a snapshot group.
Snapshot schedule	If Yes, a schedule is set for automatically creating snapshots.
Reserved capacity settings	
Alert me when...	<p>Use the spinner box to adjust the percentage point at which System Manager sends an alert notification when the reserved capacity for a snapshot group is nearing full.</p> <p>When the reserved capacity for the snapshot group exceeds the specified threshold, System Manager sends an alert, allowing you time to increase reserved capacity or to delete unnecessary objects.</p>
Policy for full reserved capacity	<p>You can choose one of the following policies:</p> <ul style="list-style-type: none"> • Purge oldest snapshot image — System Manager automatically purges the oldest snapshot image in the snapshot group, which releases the snapshot image reserved capacity for reuse within the group. • Reject writes to base volume — When the reserved capacity reaches its maximum defined percentage, System Manager rejects any I/O write request to the base volume that triggered the reserved capacity access.
Associated objects	
Base volume	The name of the base volume used for the group. A base volume is the source from which a snapshot image is created. It can be a thick or thin volume and is typically assigned to a host. The base volume can reside in either a volume group or disk pool.

Setting	Description
Snapshot images	The number of images created from this group. A snapshot image is a logical copy of volume data, captured at a particular point-in-time. Like a restore point, snapshot images allow you to roll back to a known good data set. Although the host can access the snapshot image, it cannot directly read or write to it.

- Click **Save** to apply your changes to the snapshot group settings.

Change the reserved capacity settings for a snapshot volume

You can change the settings for a snapshot volume to adjust the percentage point at which the system sends an alert notification when the reserved capacity for a snapshot volume is nearing full.

Steps

- Select **Storage > Pools & Volume Groups**.
- Click the **Reserved Capacity** tab.
- Select the snapshot volume that you want to edit, and then click **View/Edit Settings**.

The **Snapshot Volume Reserved Capacity Settings** dialog box appears.

- Change the reserved capacity settings for the snapshot volume as appropriate.

Field Details

Setting	Description
Alert me when...	<p>Use the spinner box to adjust the percentage point at which the system sends an alert notification when the reserved capacity for a member volume is nearing full.</p> <p>When the reserved capacity for the snapshot volume exceeds the specified threshold, the system sends an alert, allowing you time to increase reserved capacity or to delete unnecessary objects.</p>

- Click **Save** to apply your changes to the snapshot volume reserved capacity settings.

Change the reserved capacity settings for a consistency group member volume

You can change the settings for a consistency group member volume to adjust the percentage point at which System Manager sends an alert notification when the reserved capacity for a member volume is nearing full and to change the policy to use when the reserved capacity reaches its maximum defined percentage.

About this task

Changing the reserved capacity settings for an individual member volume also changes the reserved capacity

settings for all member volumes associated with a consistency group.


Steps

1. Select **Storage › Pools & Volume Groups**.
2. Click the **Reserved Capacity** tab.
3. Select the consistency group member volume that you want to edit, and then click **View/Edit Settings**.

The **Member Volume Reserved Capacity Settings** dialog box appears.

4. Change the reserved capacity settings for the member volume as appropriate.

Field Details

Setting	Description
Alert me when...	<p>Use the spinner box to adjust the percentage point at which System Manager sends an alert notification when the reserved capacity for a member volume is nearing full.</p> <p>When the reserved capacity for the member volume exceeds the specified threshold, System Manager sends an alert, allowing you time to increase reserved capacity or to delete unnecessary objects.</p> <div> Changing the Alert setting for one member volume will change it for <i>all</i> member volumes that belong to the same consistency group.</div>
Policy for full reserved capacity	<p>You can choose one of the following policies:</p> <ul style="list-style-type: none">• Purge oldest snapshot image — System Manager automatically purges the oldest snapshot image in the consistency group, which releases the member's reserved capacity for reuse within the group.• Reject writes to base volume — When the reserved capacity reaches its maximum defined percentage, System Manager rejects any I/O write request to the base volume that triggered the reserved capacity access.

5. Click **Save** to apply your changes.

Results

System Manager changes the reserved capacity settings for the member volume, as well as the reserved capacity settings for all member volumes in the consistency group.

Change the reserved capacity settings for a mirrored pair volume

You can change the settings for a mirrored pair volume to adjust the percentage point at which System Manager sends an alert notification when the reserved capacity for a mirrored pair volume is nearing full.


Steps

1. Select **Storage › Pools & Volume Groups**.
2. Select the **Reserved Capacity** tab.
3. Select the mirrored pair volume that you want to edit, and then click **View/Edit Settings**.

The **Mirrored Pair Volume Reserved Capacity Settings** dialog box appears.

4. Change the reserved capacity settings for the mirrored pair volume as appropriate.

Field Details

Setting	Description
Alert me when...	<p>Use the spinner box to adjust the percentage point at which System Manager sends an alert notification when the reserved capacity for a mirrored pair is nearing full.</p> <p>When the reserved capacity for the mirrored pair exceeds the specified threshold, System Manager sends an alert, allowing you time to increase reserved capacity.</p> <div> Changing the Alert setting for one mirrored pair changes the Alert setting for all mirrored pairs that belong to the same mirror consistency group.</div>
Policy for full reserved capacity	<p>You can choose one of the following policies:</p> <ul style="list-style-type: none">• Purge oldest snapshot image – System Manager automatically purges the oldest snapshot image in the consistency group, which releases the member's reserved capacity for reuse within the group.• Reject writes to base volume – When the reserved capacity reaches its maximum defined percentage, System Manager rejects any I/O write request to the base volume that triggered the reserved capacity access.

5. Click **Save** to apply your changes.

Result

System Manager changes the reserved capacity settings for the member volume, as well as the reserved capacity settings for all member volumes in the consistency group.

Cancel pending snapshot image

You can cancel a pending snapshot image before it completes. Snapshots occur asynchronously, and the status of the snapshot is pending until the snapshot is complete. The snapshot image completes as soon as the synchronization operation is complete.

About this task

A snapshot image is in a Pending state due to the following concurrent conditions:

- The base volume for a snapshot group or one or more member volumes of a consistency group that contains this snapshot image is a member of an asynchronous mirror group.
- The volume or volumes are currently in an asynchronous mirroring synchronizing operation.

Steps

1. Select **Storage › Pools & Volume Groups**.
2. Click the **Reserved Capacity** tab.
3. Select the snapshot group for which you want to cancel a pending snapshot image, and then click **Uncommon Tasks › Cancel pending snapshot image**.
4. Click **Yes** to confirm that you want to cancel the pending snapshot image.

Delete snapshot group

You delete a snapshot group when you want to permanently delete its data and remove it from the system. Deleting a snapshot group reclaims reserved capacity for reuse in the pool or volume group.

About this task

When a snapshot group is deleted, all snapshot images in the group also are deleted.

Steps

1. Select **Storage › Pools & Volume Groups**.
2. Click the **Reserved Capacity** tab.
3. Select the snapshot group that you want to delete, and then click **Uncommon Tasks › Delete snapshot group**.

The **Confirm Delete Snapshot Group** dialog box appears.

4. Type `delete` to confirm.

Results

System Manager performs the following actions:

- Deletes all snapshot images associated with the snapshot group.
- Disables any snapshot volumes associated with the snapshot group's images.
- Deletes the reserved capacity that exists for the snapshot group.

FAQs

What is a hot spare drive?

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.

If a drive fails in the storage array, the hot spare drive is automatically substituted for the failed drive without requiring a physical swap. If the hot spare drive is available when a drive fails, the controller uses redundancy data to reconstruct the data from the failed drive to the hot spare drive.

A hot spare drive is not dedicated to a specific volume group. Instead, you can use a hot spare drive for any failed drive in the storage array with the same capacity or smaller capacity. A hot spare drive must be of the same media type (HDD or SSD) as the drives that it is protecting.



Hot spare drives are not supported with pools. Instead of hot spare drives, pools use the preservation capacity within each drive that comprises the pool.

What is a volume group?

A volume group is a container for volumes with shared characteristics. A volume group has a defined capacity and RAID level. You can use a volume group to create one or more volumes accessible to a host. (You create volumes from either a volume group or a pool.)

What is a pool?

A pool is a set of drives that is logically grouped. You can use a pool to create one or more volumes accessible to a host. (You create volumes from either a pool or a volume group.)

Pools can eliminate the need for administrators to monitor usage on each host to determine when they are likely to run out of storage space and avoid conventional disk resizing outages. When a pool nears depletion, additional drives can be added to the pool non-disruptively and capacity growth is transparent to the host.

With pools, data is automatically re-distributed to maintain equilibrium. By distributing parity information and spare capacity throughout the pool, every drive in the pool can be used to rebuild a failed drive. This approach does not use dedicated hot spare drives; instead, preservation (spare) capacity is reserved throughout the pool. Upon drive failure, segments on other drives are read to recreate the data. A new drive is then chosen to write each segment that was on a failed drive so that data distribution across drives is maintained.

What is reserved capacity?

Reserved capacity is the physically allocated capacity that stores data for copy service objects such as snapshot images, consistency group member volumes, and mirrored pair volumes.

The reserved capacity volume that is associated with a copy service operation resides in a pool or a volume group. You create reserved capacity from either a pool or volume group.

What is FDE/FIPS security?

FDE/FIPS security refers to secure-capable drives that encrypt data during writes and decrypt data during reads using a unique encryption key. These secure-capable drives prevent unauthorized access to the data on a drive that is physically removed from the storage array.

Secure-capable drives can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives. FIPS drives have undergone certification testing.



For volumes that require FIPS support, use only FIPS drives. Mixing FIPS and FDE drives in a volume group or pool will result in all drives being treated as FDE drives. Also, an FDE drive cannot be added to or used as a spare in an all-FIPS volume group or pool.

What is redundancy check?

A redundancy check determines whether the data on a volume in a pool or volume group is consistent. Redundancy data is used to quickly reconstruct information on a replacement drive if one of the drives in the pool or volume group fails.

You can perform this check only on one pool or volume group at a time. A volume redundancy check performs the following actions:

- Scans the data blocks in a RAID 3 volume, a RAID 5 volume, or a RAID 6 volume, and then checks the redundancy information for each block. (RAID 3 can only be assigned to volume groups using the command line interface.)
- Compares the data blocks on RAID 1 mirrored drives.
- Returns redundancy errors if the data is determined to be inconsistent by the controller firmware.



Immediately running a redundancy check on the same pool or volume group might cause an error. To avoid this problem, wait one to two minutes before running another redundancy check on the same pool or volume group.

What are the differences between pools and volume groups?

A pool is similar to a volume group, with the following differences.

- The data in a pool is stored randomly on all drives in the pool, unlike data in a volume group, which is stored on the same set of drives.
- A pool has less performance degradation when a drive fails, and takes less time to reconstruct.
- A pool has built-in preservation capacity; therefore, it does not require dedicated hot spare drives.
- A pool allows a large number of drives to be grouped.
- A pool does not need a specified RAID level.

Why would I want to manually configure a pool?

The following examples describe why you would want to manually configure a pool.

- If you have multiple applications on your storage array and do not want them competing for the same drive resources, you might consider manually creating a smaller pool for one or more of the applications.

You can assign just one or two volumes instead of assigning the workload to a large pool that has many volumes across which to distribute the data. Manually creating a separate pool that is dedicated to the workload of a specific application can allow storage array operations to perform more rapidly, with less contention.

To manually create a pool: Select **Storage**, and then select **Pools & Volume Groups**. From the **All Capacity** tab, click **Create > Pool**.

- If there are multiple pools of the same drive type, a message appears indicating that System Manager

cannot recommend the drives for a pool automatically. However, you can manually add the drives to an existing pool.

To manually add drives to an existing pool: From the **Pools & Volume Groups** page, select the pool, and then click **Add Capacity**.

Why are capacity alerts important?

Capacity alerts indicate when to add drives to a pool. A pool needs sufficient free capacity to successfully perform storage array operations. You can prevent interruptions to these operations by configuring System Manager to send alerts when the free capacity of a pool reaches or exceeds a specified percentage.

You set this percentage when you create a pool using either the **Pool auto-configuration** option or the **Create pool** option. If you choose the automatic option, default settings automatically determine when you receive alert notifications. If you choose to manually create the pool, you can determine the alert notification settings; or if you prefer, you can accept the default settings. You can adjust these settings later in **Settings > Alerts**.



When the free capacity in the pool reaches the specified percentage, an alert notification is sent using the method you specified in the alert configuration.

Why can't I increase my preservation capacity?

If you have created volumes on all available usable capacity, you might not be able to increase preservation capacity.

Preservation capacity is the amount of capacity (number of drives) that is reserved on 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. If you have created volumes on all available usable capacity, you cannot increase preservation capacity without adding capacity to the pool by either adding drives or deleting volumes.

You can change the preservation capacity by selecting **Storage**, and then the **Pools & Volume Groups** tile. Select the pool that you want to edit. Click **View/Edit Settings**, and then select the **Settings** tab.



Preservation capacity is specified as a number of drives, even though the actual preservation capacity is distributed across the drives in the pool.

Is there a limit on the number of drives I can remove from a pool?

System Manager sets limits for how many drives you can remove from a pool.

- You cannot reduce the number of drives in a pool to be less than 11 drives.
- You cannot remove drives if there is not enough free capacity in the pool to contain the data from the removed drives when that data is redistributed to the remaining drives in the pool.
- You can remove a maximum of 60 drives at a time. If you select more than 60 drives, the Remove Drives option is disabled. If you need to remove more than 60 drives, repeat the Remove Drives operation.

What media types are supported for a drive?

The following media types are supported: Hard Disk Drive (HDD) and Solid State Disk

(SSD).

Why are some drives not showing up?

In the **Add Capacity** dialog, not all drives are available for adding capacity to an existing pool or volume group.

Drives are not eligible for any of the following reasons:

- A drive must be unassigned and not secure-enabled. Drives already part of another pool, another volume group, or configured as a hot spare are not eligible. If a drive is unassigned but is secure-enabled, you must manually erase that drive for it to become eligible.
- A drive that is in a non-optimal state is not eligible.
- If the capacity of a drive is too small, it is not eligible.
- The drive media type must match within a pool or volume group. You cannot mix Hard Disk Drives (HDDs) with Solid State Disks (SSDs).
- If a pool or volume group contains all secure-capable drives, non-secure-capable drives are not listed.
- If a pool or volume group contains all Federal Information Processing Standards (FIPS) drives, non-FIPS drives are not listed.
- If a pool or volume group contains all Data Assurance (DA)-capable drives and there is at least one DA-enabled volume in the pool or volume group, a drive that is not DA capable is not eligible, so it cannot be added to that pool or volume group. However, if there is no DA-enabled volume in the pool or volume group, a drive that is not DA capable can be added to that pool or volume group. If you decide to mix these drives, keep in mind that you cannot create any DA-enabled volumes.



Capacity can be increased in your storage array by adding new drives or by deleting pools or volume groups.

How do I maintain shelf/drawer loss protection?

To maintain shelf/drawer loss protection for a pool or volume group, use the criteria specified in the following table.

Level	Criteria for shelf/drawer loss protection	Minimum number of shelves/drawers required
Pool	For shelves, the pool must contain no more than two drives in a single shelf. For drawers, the pool must include an equal number of drives from each drawer.	6 for shelves 5 for drawers
RAID 6	The volume group contains no more than two drives in a single shelf or drawer.	3

Level	Criteria for shelf/drawer loss protection	Minimum number of shelves/drawers required
RAID 3 or RAID 5	Each drive in the volume group is located in a separate shelf or drawer.	3
RAID 1	Each drive in a mirrored pair must be located in a separate shelf or drawer.	2
RAID 0	Cannot achieve shelf/drawer loss protection.	Not applicable



Shelf/drawer loss protection is not maintained if a drive has already failed in the pool or volume group. In this situation, losing access to a drive shelf or drawer, and consequently another drive in the pool or volume group, causes loss of data.

What RAID level is best for my application?

To maximize the performance of a volume group, you must select the appropriate RAID level. You can determine the appropriate RAID level by knowing the read and write percentages for the applications that are accessing the volume group. Use the Performance page to obtain these percentages.

RAID levels and application performance

RAID relies on a series of configurations, called *levels*, to determine how user and redundancy data is written and retrieved from the drives. Each RAID level provides different performance features. Applications with a high read percentage perform well using RAID 5 volumes or RAID 6 volumes because of the outstanding read performance of the RAID 5 and RAID 6 configurations.

Applications with a low read percentage (write-intensive) do not perform as well on RAID 5 volumes or RAID 6 volumes. The degraded performance is the result of the way that a controller writes data and redundancy data to the drives in a RAID 5 volume group or a RAID 6 volume group.

Select a RAID level based on the following information.

RAID 0

- **Description**

- Non-redundant, striping mode.

- **How it works**

- RAID 0 stripes data across all of the drives in the volume group.

- **Data protection features**

- RAID 0 is not recommended for high availability needs. RAID 0 is better for non-critical data.
- If a single drive fails in the volume group, all of the associated volumes fail, and all data is lost.

- **Drive number requirements**

- A minimum of one drive is required for RAID Level 0.
- RAID 0 volume groups can have more than 30 drives.
- You can create a volume group that includes all of the drives in the storage array.

RAID 1 or RAID 10

- **Description**

- Striping/mirror mode.

- **How it works**

- RAID 1 uses disk mirroring to write data to two duplicate disks simultaneously.
- RAID 10 uses drive striping to stripe data across a set of mirrored drive pairs.

- **Data protection features**

- RAID 1 and RAID 10 offer high performance and the best data availability.
- RAID 1 and RAID 10 use drive mirroring to make an exact copy from one drive to another drive.
- If one of the drives in a drive pair fails, the storage array can instantly switch to the other drive without any loss of data or service.
- A single drive failure causes associated volumes to become degraded. The mirror drive allows access to the data.
- A drive-pair failure in a volume group causes all of the associated volumes to fail, and data loss could occur.

- **Drive number requirements**

- A minimum of two drives is required for RAID 1: one drive for the user data, and one drive for the mirrored data.
- If you select four or more drives, RAID 10 is automatically configured across the volume group: two drives for user data, and two drives for the mirrored data.
- You must have an even number of drives in the volume group. If you do not have an even number of drives and you have some remaining unassigned drives, select **Storage > Pools & Volume Groups** to add additional drives to the volume group, and retry the operation.
- RAID 1 and RAID 10 volume groups can have more than 30 drives. A volume group can be created that includes all of the drives in the storage array.

RAID 5

- **Description**

- High I/O mode.

- **How it works**

- User data and redundant information (parity) are striped across the drives.
- The equivalent capacity of one drive is used for redundant information.

- **Data protection features**

- If a single drive fails in a RAID 5 volume group, all of the associated volumes become degraded. The redundant information allows the data to still be accessed.
- If two or more drives fail in a RAID 5 volume group, all of the associated volumes fail, and all data is lost.

- **Drive number requirements**

- You must have a minimum of three drives in the volume group.
- Typically, you are limited to a maximum of 30 drives in the volume group.

RAID 6

- **Description**

- High I/O mode.

- **How it works**

- User data and redundant information (dual parity) are striped across the drives.
- The equivalent capacity of two drives is used for redundant information.

- **Data protection features**

- If one or two drives fail in a RAID 6 volume group, all of the associated volumes become degraded, but the redundant information allows the data to still be accessed.
- If three or more drives fail in a RAID 6 volume group, all of the associated volumes fail, and all data is lost.

- **Drive number requirements**

- You must have a minimum of five drives in the volume group.
- Typically, you are limited to a maximum of 30 drives in the volume group.



You cannot change the RAID level of a pool. System Manager automatically configures pools as RAID 6.

RAID levels and data protection

RAID 1, RAID 5, and RAID 6 write redundancy data to the drive media for fault tolerance. The redundancy data might be a copy of the data (mirrored) or an error-correcting code derived from the data. You can use the redundancy data to quickly reconstruct information on a replacement drive if a drive fails.

You configure a single RAID level across a single volume group. All redundancy data for that volume group is stored within the volume group. The capacity of the volume group is the aggregate capacity of the member drives minus the capacity reserved for redundancy data. The amount of capacity needed for redundancy depends on the RAID level used.

What is Data Assurance?

Data Assurance (DA) implements the T10 Protection Information (PI) standard, which increases data integrity by checking for and correcting errors that might occur as data is transferred along the I/O path.

The typical use of the Data Assurance feature will check the portion of the I/O path between the controllers and drives. DA capabilities are presented at the pool and volume group level in System Manager.

When this feature is enabled, the storage array appends error-checking codes (also known as cyclic redundancy checks or CRCs) to each block of data in the volume. After a data block is moved, the storage array uses these CRC codes to determine if any errors occurred during transmission. Potentially corrupted data is neither written to disk nor returned to the host. If you want to use the DA feature, select a pool or volume group that is DA capable when you create a new volume (look for "Yes" next to "DA" in the pool and

volume group candidates table).

Make sure you assign these DA-enabled volumes to a host using an I/O interface that is capable of DA. I/O interfaces that are capable of DA include Fibre Channel, SAS, iSCSI over TCP/IP, NVMe/FC, NVMe/IB, NVMe/RoCE and iSER over InfiniBand (iSCSI Extensions for RDMA/IB). DA is not supported by SRP over InfiniBand.

What is secure-capable (Drive Security)?

Drive Security is a feature that prevents unauthorized access to data on secure-enabled drives when removed from the storage array. These drives can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.

What do I need to know about increasing reserved capacity?

Typically, you should increase capacity when you receive a warning that the reserved capacity is in danger of becoming full. You can increase reserved capacity only in increments of 8 GiB.

- You must have sufficient free capacity in the pool or volume group so it can be expanded if necessary.

If no free capacity exists on any pool or volume group, you can add unassigned capacity in the form of unused drives to a pool or volume group.

- The volume in the pool or volume group must have an Optimal status and must not be in any state of modification.
- Free capacity must exist in the pool or volume group that you want to use to increase capacity.
- You cannot increase reserved capacity for a snapshot volume that is read-only. Only snapshot volumes that are read-write require reserved capacity.

For snapshot operations, reserved capacity is typically 40 percent of the base volume. For asynchronous mirroring operations reserved capacity is typically 20 percent of the base volume. Use a higher percentage if you believe the base volume will undergo many changes or if the estimated life expectancy of a storage object's copy service operation will be very long.

Why can't I choose another amount to decrease by?

You can decrease reserved capacity only by the amount you used to increase it. Reserved capacity for member volumes can be removed only in the reverse order they were added.

You cannot decrease the reserved capacity for a storage object if one of these conditions exists:

- If the storage object is a mirrored pair volume.
- If the storage object contains only one volume for reserved capacity. The storage object must contain at least two volumes for reserved capacity.
- If the storage object is a disabled snapshot volume.
- If the storage object contains one or more associated snapshot images.

You can remove volumes for reserved capacity only in the reverse order that they were added.

You cannot decrease the reserved capacity for a snapshot volume that is read-only because it does not have any associated reserved capacity. Only snapshot volumes that are read-write require reserved capacity.

Why do I need reserved capacity for each member volume?

Each member volume in a snapshot consistency group must have its own reserved capacity to save any modifications made by the host application to the base volume without affecting the referenced consistency group snapshot image. Reserved capacity provides the host application with write access to a copy of the data contained in the member volume that is designated as read-write.

A consistency group snapshot image is not directly read or write accessible to hosts. Rather, the snapshot image is used to save only the data captured from the base volume.

During the creation of a consistency group snapshot volume that is designated as read-write, System Manager creates a reserved capacity for each member volume in the consistency group. This reserved capacity provides the host application with write access to a copy of the data contained in the consistency group snapshot image.

How do I view and interpret all SSD Cache statistics?

You can view nominal statistics and detailed statistics for SSD Cache. Nominal statistics are a subset of the detailed statistics. The detailed statistics can be viewed only when you export all SSD statistics to a `.csv` file. As you review and interpret the statistics, keep in mind that some interpretations are derived by looking at a combination of statistics.

Nominal statistics

To view SSD Cache statistics, select **Storage > Pools & Volume Groups**. Select the SSD Cache that you want to view statistics for, and then select **More > View Statistics**. The nominal statistics are shown on the View SSD Cache Statistics dialog.



This feature is not available on the EF600 storage system.

The following list includes nominal statistics, which are a subset of the detailed statistics.

Nominal statistic	Description
Reads/Writes	The total number of host reads from or host writes to the SSD Cache-enabled volumes. Compare the Reads relative to Writes. The Reads need to be greater than the Writes for effective SSD Cache operation. The greater the ratio of Reads to Writes, the better the operation of the cache.
Cache Hits	A count of the number of cache hits.

Nominal statistic	Description
Cache Hits (%)	<p>Derived from Cache Hits / (reads + writes). The Cache Hit percentage should be greater than 50 percent for effective SSD Cache operation. A small number could indicate several things:</p> <ul style="list-style-type: none"> • The ratio of Reads to Writes is too small • Reads are not repeated • Cache capacity is too small
Cache Allocation (%)	<p>The amount of SSD Cache storage that is allocated, expressed as a percentage of the SSD Cache storage that is available to this controller. Derived from allocated bytes / available bytes.</p> <p>Cache Allocation percentage normally shows as 100 percent. If this number is less than 100 percent, it means either the cache has not been warmed or the SSD Cache capacity is larger than all the data being accessed. In the latter case, a smaller SSD Cache capacity could provide the same level of performance. Note that this does not indicate that cached data has been placed into the SSD Cache; it is simply a preparation step before data can be placed in the SSD Cache.</p>
Cache Utilization (%)	<p>The amount of SSD Cache storage that contains data from enabled volumes, expressed as a percentage of SSD Cache storage that is allocated. This value represents the utilization or density of the SSD Cache derived from user data bytes / allocated bytes.</p> <p>Cache Utilization percentage normally is lower than 100 percent, perhaps much lower. This number shows the percent of SSD Cache capacity that is filled with cache data. This number is lower than 100 percent because each allocation unit of the SSD Cache, the SSD Cache block, is divided into smaller units called sub-blocks, which are filled somewhat independently. A higher number is generally better, but performance gains can be significant even with a smaller number.</p>

Detailed statistics

The detailed statistics consist of the nominal statistics, plus additional statistics. These additional statistics are saved along with the nominal statistics, but unlike the nominal statistics, they do not display in the **View SSD Cache Statistics** dialog. You can view the detailed statistics only after exporting the statistics to a `.csv` file.

When viewing the `.csv` file, notice that the detailed statistics are listed after the nominal statistics:

Detailed statistics	Description
Read Blocks	The number of blocks in host reads.
Write Blocks	The number of blocks in host writes.
Full Hit Blocks	<p>The number of blocks in cache hits.</p> <p>The full hit blocks indicate the number of blocks that have been read entirely from SSD Cache. The SSD Cache is only beneficial to performance for those operations that are full cache hits.</p>
Partial Hits	The number of host reads where at least one block, but not all blocks, were in the SSD Cache. A partial hit is an SSD Cache miss where the reads were satisfied from the base volume.
Partial Hits - Blocks	<p>The number of blocks in Partial Hits.</p> <p>Partial cache hits and partial cache hit blocks result from an operation that has only a portion of its data in the SSD Cache. In this case, the operation must get the data from the cached hard disk drive (HDD) volume. The SSD Cache offers no performance benefit for this type of hit. If the partial cache hit blocks count is higher than the full cache hit blocks, a different I/O characteristic type (file system, database, or web server) could improve the performance. It is expected that there will be a larger number of Partial Hits and Misses as compared to Cache Hits while the SSD Cache is warming.</p>
Misses	The number of host reads where none of the blocks were in the SSD Cache. An SSD Cache miss occurs when the reads were satisfied from the base volume. It is expected that there will be a larger number of Partial Hits and Misses as compared to Cache Hits while the SSD Cache is warming.
Misses - Blocks	The number of blocks in Misses.
Populate Actions (Host Reads)	The number of host reads where data was copied from the base volume to the SSD Cache.
Populate Actions (Host Reads) - Blocks	The number of blocks in Populate Actions (Host Reads).

Detailed statistics	Description
Populate Actions (Host Writes)	<p>The number of host writes where data was copied from the base volume to the SSD Cache.</p> <p>The Populate Actions (Host Writes) count might be zero for the cache configuration settings that do not fill the cache as a result of a Write I/O operation.</p>
Populate Actions (Host Writes) - Blocks	The number of blocks in Populate Actions (Host Writes).
Invalidate Actions	The number of times data was invalidated or removed from the SSD Cache. A cache invalidate operation is performed for each host write request, each host read request with Forced Unit Access (FUA), each verify request, and in some other circumstances.
Recycle Actions	<p>The number of times that the SSD Cache block has been re-used for another base volume and/or a different logical block addressing (LBA) range.</p> <p>For effective cache operation, the number of recycles must be small compared to the combined number of read and write operations. If the number of Recycle Actions is close to the combined number of Reads and Writes, the SSD Cache is thrashing. Either the cache capacity needs to be increased or the workload is not favorable for use with SSD Cache.</p>
Available Bytes	The number of bytes available in the SSD Cache for use by this controller.
Allocated Bytes	The number of bytes allocated from the SSD Cache by this controller. Bytes allocated from the SSD Cache might be empty or they might contain data from base volumes.
User Data Bytes	<p>The number of allocated bytes in the SSD Cache that contain data from base volumes.</p> <p>The available bytes, allocated bytes, and user data bytes are used to compute the Cache Allocation percentage and the Cache Utilization percentage.</p>

What is optimization capacity for pools?

SSD drives will have longer life and better maximum write performance when a portion of their capacity is unallocated.

For drives associated with a pool, unallocated capacity is comprised of a pool's preservation capacity, the free

capacity (capacity not used by volumes), and a portion of the usable capacity set aside as additional optimization capacity. The additional optimization capacity ensures a minimum level of optimization capacity by reducing the usable capacity, and as such, is not available for volume creation.

When a pool is created, a recommended optimization capacity is generated that provides a balance of performance, drive wear life, and available capacity. The Additional Optimization Capacity slider located in the **Pool Settings** dialog allows adjustments to the pool's optimization capacity. Adjusting the slider provides for better performance and drive wear life at the expense of available capacity, or additional available capacity at the expense of performance and drive wear life.



The Additional Optimization Capacity slider is only available for EF600 storage systems.

What is optimization capacity for volume groups?

SSD drives will have longer life and better maximum write performance when a portion of their capacity is unallocated.

For drives associated with a volume group, unallocated capacity is comprised of a volume group's free capacity (capacity not used by volumes), and a portion of the usable capacity set aside as optimization capacity. The additional optimization capacity ensures a minimum level of optimization capacity by reducing the usable capacity, and as such, is not available for volume creation.

When a volume group is created, a recommended optimization capacity is generated that provides a balance of performance, drive wear life, and available capacity. The Additional Optimization Capacity slider in the **Volume Group Settings** dialog allows adjustments to a volume group's optimization capacity. Adjusting the slider provides for better performance and drive wear life at the expense of available capacity, or additional available capacity at the expense of performance and drive wear life.



The Additional Optimization Capacity slider is only available for EF600 storage systems.

Volumes

Concepts

Volumes in the storage array

Volumes are data containers that manage and organize the storage space on your storage array. Volumes are created from the storage capacity available on your storage array and make it easy to organize and use your system's resources. This concept is similar to using folders/directories on a computer to organize files for easy and quick access.

Volumes are the only data layer visible to hosts. In a SAN environment, volumes are mapped to logical unit numbers (LUNs), which are visible to hosts. LUNs hold the user data that is accessible using one or more of the host access protocols supported by the storage array, including FC, iSCSI, and SAS.

Volume types you can create from pools and volume groups

Volumes draw their capacity from pools or volume groups. You can create the following types of volumes from the pools or volume groups that exist on your storage array.

- **From pools** — You can create volumes from a pool as either *fully-provisioned (thick) volumes* or *thinly-*

provisioned (thin) volumes.



SANtricity System Manager does not provide an option to create thin volumes. If you want to create thin volumes, use the Command Line Interface (CLI).

- **From volume groups** — You can create volumes from a volume group only as *fully-provisioned (thick) volumes*.

Thick volumes and thin volumes draw capacity from the storage array in different ways:

- The capacity for a thick volume is allocated when the volume is created.
- The capacity for a thin volume is allocated as data when written to the volume.

Thin provisioning helps to avoid wasted allocated capacity and can save businesses on up-front storage costs. However, full provisioning has the benefit of less latency because all storage is allocated at once when thick volumes are created.



The EF600 storage system does not support thin provisioning.

Characteristics of volumes

Each volume in a pool or volume group can have its own individual characteristics based on what type of data will be stored in it. Some of these characteristics include:

- **Segment size** — A segment is the amount of data in kilobytes (KiB) that is stored on a drive before the storage array moves to the next drive in the stripe (RAID group). The segment size is equal to or less than the capacity of the volume group. The segment size is fixed and cannot be changed for pools.
- **Capacity** — You create a volume from the free capacity available in either a pool or volume group. Before you create a volume, the pool or volume group must already exist, and it must have enough free capacity to create the volume.
- **Controller ownership** — All storage arrays can have either one or two controllers. On a single-controller array, a volume's workload is managed by a single controller. On a dual-controller array, a volume will have a preferred controller (A or B) that "owns" the volume. In a dual-controller configuration, volume ownership is automatically adjusted using the Automatic Load Balancing feature to correct any load balance issues when workloads shift across the controllers. Automatic load balancing provides automated I/O workload balancing and ensures that incoming I/O traffic from the hosts is dynamically managed and balanced across both controllers.
- **Volume assignment** — You can give hosts access to a volume either when you create the volume or at a later time. All host access is managed through a logical unit number (LUN). Hosts detect LUNs that are, in turn, assigned to volumes. If you are assigning a volume to multiple hosts, use clustering software to make sure that the volume is available to all of the hosts.

The host type can have specific limits on how many volumes the host can access. Keep this limitation in mind when you create volumes for use by a particular host.

- **Descriptive name** — You can name a volume whatever name you like, but we recommend making the name descriptive.

During volume creation, each volume is allocated capacity and is assigned a name, segment size (volume groups only), controller ownership, and volume-to-host assignment. Volume data is automatically load balanced across controllers, as needed.

Volume terminology

Learn how the volume terms apply to your storage array.

All volume types

Term	Description
Allocated capacity	<p>You use allocated capacity to create volumes and for copy services operations.</p> <p>Allocated capacity and reported capacity are the same for thick volumes, but are different for thin volumes. For a thick volume, the physically allocated space is equal to the space that is reported to the host. For a thin volume, reported capacity is the capacity that is reported to the hosts, whereas allocated capacity is the amount of drive space that is currently allocated for writing data.</p>
Application	<p>An application is software such as SQL Server or Exchange. You define one or more workloads to support each application. For some applications, System Manager will automatically recommend a volume configuration that optimizes storage. Characteristics such as I/O type, segment size, controller ownership, and read and write cache are included in the volume configuration.</p>
Capacity	<p>Capacity is the amount of data that you can store in a volume.</p>
Controller ownership	<p>Controller ownership defines the controller that is designated to be the owning, or primary, controller of the volume. A volume can have a preferred controller (A or B) that “owns” the volume. Volume ownership is automatically adjusted using the Automatic Load Balancing feature to correct any load balance issues when workloads shift across the controllers. Automatic Load Balancing provides automated I/O workload balancing and ensures that incoming I/O traffic from the hosts is dynamically managed and balanced across both controllers.</p>

Term	Description
Dynamic cache read prefetch	<p>Dynamic cache read prefetch allows the controller to copy additional sequential data blocks into the cache while it is reading data blocks from a drive to the cache. This caching increases the chance that future requests for data can be filled from the cache. Dynamic cache read prefetch is important for multimedia applications that use sequential I/O. The rate and amount of data that is prefetched into cache is self-adjusting based on the rate and request size of the host reads. Random access does not cause data to be prefetched into cache. This feature does not apply when read caching is disabled.</p> <p>For a thin volume, dynamic cache read prefetch is always disabled and cannot be changed.</p>
Free capacity area	<p>A free capacity area is the free capacity that can result from deleting a volume or from not using all available free capacity during volume creation. When you create a volume in a volume group that has one or more free capacity areas, the volume's capacity is limited to the largest free capacity area in that volume group. For example, if a volume group has a total of 15 GiB free capacity, and the largest free capacity area is 10 GiB, the largest volume you can create is 10 GiB.</p> <p>By consolidating free capacity, you can create additional volumes from the maximum amount of free capacity in a volume group.</p>
Host	A host is a server that sends I/O to a volume on a storage array.
Host cluster	A host cluster is a group of hosts. You create a host cluster to make it easy to assign the same volumes to multiple hosts.
Hot spare drive	Hot spare drives are supported only with volume groups. A hot spare drive contains no data and acts as a standby in case a drive fails in RAID 1, RAID 3, RAID 5, or RAID 6 volumes contained in a volume group. The hot spare drive adds another level of redundancy to your storage array.

Term	Description
LUN	<p>A logical unit number (LUN) is the number assigned to the address space that a host uses to access a volume. The volume is presented to the host as capacity in the form of a LUN.</p> <p>Each host has its own LUN address space. Therefore, the same LUN can be used by different hosts to access different volumes.</p>
Media scan	<p>A media scan provides a way of detecting drive media errors before they are found during a normal read from or write to the drives. A media scan is performed as a background operation and scans all data and redundancy information in defined user volumes.</p>
Namespace	<p>A namespace is NVM storage that is formatted for block access. It is analogous to a logical unit in SCSI, which relates to a volume in the storage array.</p>
Pool	<p>A pool is a set of drives that is logically grouped. You can use a pool to create one or more volumes accessible to a host. (You create volumes from either a pool or a volume group.)</p>
Pool or volume group capacity	<p>Pool, volume, or volume group capacity is the capacity in a storage array that has been assigned to a pool or volume group. This capacity is used to create volumes and service the various capacity needs of copy services operations and storage objects.</p>
Read cache	<p>The read cache is a buffer that stores data that has been read from the drives. The data for a read operation might already be in the cache from a previous operation, which eliminates the need to access the drives. The data stays in the read cache until it is flushed.</p>
Reported capacity	<p>Reported capacity is the capacity that is reported to the host and can be accessed by the host.</p> <p>Reported capacity and allocated capacity are the same for thick volumes, but are different for thin volumes. For a thick volume, the physically allocated space is equal to the space that is reported to the host. For a thin volume, reported capacity is the capacity that is reported to the hosts, whereas allocated capacity is the amount of drive space that is currently allocated for writing data.</p>

Term	Description
Segment size	A segment is the amount of data in kilobytes (KiB) that is stored on a drive before the storage array moves to the next drive in the stripe (RAID group). The segment size is equal to or less than the capacity of the volume group. The segment size is fixed and cannot be changed for pools.
Striping	Striping is way of storing data on the storage array. Striping splits the flow of data into blocks of a certain size (called "block size") and then writes these blocks across the drives one by one. This way of data storage is used to distribute and store data across multiple physical drives. Striping is synonymous with RAID 0 and spreads the data across all the drives in a RAID group without parity.
Volume	A volume is a container in which applications, databases, and file systems store data. It is the logical component created for the host to access storage on the storage array.
Volume assignment	Volume assignment is how host LUNs are assigned to a volume.
Volume name	A volume name is a string of characters assigned to the volume when it is created. You can either accept the default name or provide a more descriptive name indicating the type of data stored in the volume.
Volume group	A volume group is a container for volumes with shared characteristics. A volume group has a defined capacity and RAID level. You can use a volume group to create one or more volumes accessible to a host. (You create volumes from either a volume group or a pool.)
Workload	A workload is a storage object that supports an application. You can define one or more workloads, or instances, per application. For some applications, System Manager configures the workload to contain volumes with similar underlying volume characteristics. These volume characteristics are optimized based on the type of application the workload supports. For example, if you create a workload that supports a Microsoft SQL Server application and then subsequently create volumes for that workload, the underlying volume characteristics are optimized to support Microsoft SQL Server.

Term	Description
Write cache	The write cache is a buffer that stores data from the host that has not yet been written to the drives. The data stays in the write cache until it is written to the drives. Write caching can increase I/O performance.
Write caching with mirroring	Write caching with mirroring occurs when the data written to the cache memory of one controller is also written to the cache memory of the other controller. Therefore, if one controller fails, the other can complete all outstanding write operations. Write cache mirroring is available only if write caching is enabled and two controllers are present. Write caching with mirroring is the default setting at volume creation.
Write caching without batteries	The write caching without batteries setting lets write caching continue even when the batteries are missing, failed, discharged completely, or not fully charged. Choosing write caching without batteries is not typically recommended, because data might be lost if power is lost. Typically, write caching is turned off temporarily by the controller until the batteries are charged or a failed battery is replaced.

Specific to thin volumes



SANtricity System Manager does not provide an option to create thin volumes. If you want to create thin volumes, use the command line interface (CLI).




Thin volumes are not available on the EF600 storage system.

Term	Description
Allocated capacity limit	Allocated capacity limit is the cap on how large the allocated physical capacity for a thin volume can grow.
Written capacity	Written capacity is the amount of capacity that has been written from the reserved capacity allocated for thin volumes.
Warning threshold	You can set a warning threshold alert to be issued when the allocated capacity for a thin volume reaches the percent full (the warning threshold).

Workflow for creating volumes

In SANtricity System Manager, you can create volumes by following these steps.



Data integrity and data security for volumes

You can enable volumes to use the Data Assurance (DA) feature and the Drive Security feature. These features are presented at the pool and volume group level in System Manager.

Data Assurance

Data Assurance (DA) implements the T10 Protection Information (PI) standard, which increases data integrity by checking for and correcting errors that might occur as data is transferred along the I/O path. The typical use of the Data Assurance feature will check the portion of the I/O path between the controllers and drives. DA capabilities are presented at the pool and volume group level in System Manager.

When this feature is enabled, the storage array appends error-checking codes (also known as cyclic redundancy checks or CRCs) to each block of data in the volume. After a data block is moved, the storage array uses these CRC codes to determine if any errors occurred during transmission. Potentially corrupted data is neither written to disk nor returned to the host. If you want to use the DA feature, select a pool or volume group that is DA capable when you create a new volume (look for "Yes" next to "DA" in the pool and volume group candidates table).

Make sure you assign these DA-enabled volumes to a host using an I/O interface that is capable of DA. I/O interfaces that are capable of DA include Fibre Channel, SAS, iSCSI over TCP/IP, NVMe/FC, NVMe/IB, NVMe/RoCE and iSER over InfiniBand (iSCSI Extensions for RDMA/IB). DA is not supported by SRP over InfiniBand.

Drive Security

Drive Security is a feature that prevents unauthorized access to data on secure-enabled drives when removed from the storage array. These drives can be either Full Disk Encryption (FDE) drives or drives that are certified to meet Federal Information Processing Standards 140-2 level 2 (FIPS drives).

How Drive Security works at the drive level

A secure-capable drive, either FDE or FIPS, encrypts data during writes and decrypts data during reads. This encryption and decryption does not affect the performance or user workflow. Each drive has its own unique encryption key, which can never be transferred from the drive.

How Drive Security works at the volume level

When you create a pool or volume group from secure-capable drives, you can also enable Drive Security for those pools or volume groups. The Drive Security option makes the drives and associated volume groups and pools *secure-enabled*. A pool or volume group can contain both secure-capable and non-secure-capable drives, but all drives must be secure-capable to use their encryption capabilities.

How to implement Drive Security

To implement Drive Security, you perform the following steps.

1. Equip your storage array with secure-capable drives, either FDE drives or FIPS drives. (For volumes that require FIPS support, use only FIPS drives. Mixing FIPS and FDE drives in a volume group or pool will result in all drives being treated as FDE drives. Also, an FDE drive cannot be added to or used as a spare in an all-FIPS volume group or pool.)
2. Create a security key, which is a string of characters that is shared by the controller and drives for read/write access. You can create either an internal key from the controller's persistent memory or an external key from a key management server. For external key management, authentication must be established with the key management server.
3. Enable Drive Security for pools and volume groups:
 - Create a pool or volume group (look for **Yes** in the **Secure-capable** column in the Candidates table).
 - Select a pool or volume group when you create a new volume (look for **Yes** next to **Secure-capable** in

the pool and volume group Candidates table).

With the Drive Security feature, you create a security key that is shared between the secure-enabled drives and controllers in a storage array. Whenever power to the drives is turned off and on, the secure-enabled drives change to a Security Locked state until the controller applies the security key.

SSD Cache and volumes

You can add a volume to SSD Cache as a way to improve read-only performance. SSD Cache consists of a set of solid-state disk (SSD) drives that you logically group together in your storage array.



This feature is not available on the EF600 storage system.

Volumes

Simple volume I/O mechanisms are used to move data to and from the SSD Cache. After data is cached and stored on the SSDs, subsequent reads of that data are performed on the SSD Cache, thereby eliminating the need to access the HDD volume.

SSD Cache is a secondary cache for use with the primary cache in the controller's dynamic random-access memory (DRAM).

- In primary cache, the data is stored in DRAM after a host read.
- In SSD Cache, the data is copied from volumes and stored on two internal RAID volumes (one per controller) that are automatically created when you create an SSD Cache.

The internal RAID volumes are used for internal cache processing purposes. These volumes are not accessible or displayed in the user interface. However, these two volumes do count against the total number of volumes allowed in the storage array.



Any volume assigned to use a controller's SSD Cache is not eligible for an automatic load balance transfer.

Drive Security

To use SSD Cache on a volume that is also using Drive Security (is secure-enabled), the Drive Security capabilities of the volume and the SSD Cache must match. If they do not match, the volume will not be secure-enabled.

Application-specific workloads

A workload is a storage object that supports an application. You can define one or more workloads, or instances, per application. For some applications, System Manager configures the workload to contain volumes with similar underlying volume characteristics. These volume characteristics are optimized based on the type of application the workload supports. For example, if you create a workload that supports a Microsoft SQL Server application and then subsequently create volumes for that workload, the underlying volume characteristics are optimized to support Microsoft SQL Server.

During volume creation, System Manager prompts you to answer questions about a workload's use. For example, if you are creating volumes for Microsoft Exchange, you are asked how many mailboxes you need, what your average mailbox capacity requirements are, and how many copies of the database you want. System Manager uses this information to create an optimal volume configuration for you, which can be edited as needed. Optionally, you can skip this step in the volume creation sequence.

Types of workloads

You can create two types of workloads: application-specific and other.

- **Application-specific.** When you are creating volumes using an application-specific workload, the system may recommend an optimized volume configuration to minimize contention between application workload I/O and other traffic from your application instance. Volume characteristics like I/O type, segment size, controller ownership, and read and write cache are automatically recommended and optimized for workloads that are created for the following application types.

- Microsoft® SQL Server™
- Microsoft® Exchange Server™
- Video Surveillance applications
- VMware ESXi™ (for volumes to be used with Virtual Machine File System)

You can review the recommended volume configuration and edit, add, or delete the system-recommended volumes and characteristics using the **Add/Edit Volumes** dialog box.

- **Other** (or applications without specific volume creation support). Other workloads use a volume configuration that you must manually specify when you want to create a workload that is not associated with a specific application, or if System Manager does not have built-in optimization for the application you intend to use on the storage array. You must manually specify the volume configuration using the **Add/Edit Volumes** dialog box.

Application and workload views

You can view information associated with an application-specific workload in a couple of different ways:

- You can select the **Applications & Workloads** tab in the **Volumes** tile to view the storage array's volumes grouped by workload and the application type the workload is associated with.
- You can select the **Applications & Workloads** tab in the **Performance** tile to view performance metrics (latency, IOPS, and MBs) for logical objects. The objects are grouped by application and associated workload. By collecting this performance data at regular intervals, you can establish baseline measurements and analyze trends, which can help as you investigate problems related to I/O performance.

Actions you can perform on volumes

You can perform a number of different actions on a volume: increasing capacity, deleting, copying, initializing, redistributing, changing ownership, changing cache settings, and changing media scan settings.

Increase capacity

You can expand the capacity for a volume in two ways:

- Use the free capacity that is available in the pool or volume group.

You add capacity to a volume by selecting **Storage › Pools and Volume Groups › Add Capacity**.

- Add unassigned capacity (in the form of unused drives) to the pool or volume group of the volume. Use this option when no free capacity exists in the pool or volume group.

You add unassigned capacity to the pool or volume group by selecting **Storage › Pools and Volume Groups › Add Capacity**.

If free capacity is not available in the pool or volume group, you cannot increase the capacity of the volume. You must increase the size of the pool or volume group first or delete unused volumes.

After you expand the volume capacity, you must manually increase the file system size to match. How you do this depends on the file system you are using. See your host operating system documentation for details.

Delete

Typically, you delete volumes if the volumes were created with the wrong parameters or capacity, no longer meet storage configuration needs, or are snapshot images that are no longer needed for backup or application testing. Deleting a volume increases the free capacity in the pool or volume group.

Deleting volumes causes loss of all data on those volumes. Deleting a volume will also delete any associated snapshot images, schedules, and snapshot volumes and remove any mirroring relationships.

Copy

When you copy volumes, you create a point-in-time copy of two separate volumes, the source volume and the target volume, on the same storage array. You can copy volumes by selecting **Storage › Volumes › Copy Services › Copy volume**.

Initialize

Initializing a volume erases all data from the volume. A volume is automatically initialized when it is first created. However, the Recovery Guru might advise that you manually initialize a volume to recover from certain failure conditions. When you initialize a volume, the volume keeps its WWN, host assignments, allocated capacity, and reserved capacity settings. It also keeps the same Data Assurance (DA) settings and security settings.

You can initialize volumes by selecting **Storage › Volumes › More › Initialize volumes**.

Redistribute

You redistribute volumes to move volumes back to their preferred controller owners. Typically, multipath drivers move volumes from their preferred controller owner when a problem occurs along the data path between the host and storage array.

Most host multipath drivers attempt to access each volume on a path to its preferred controller owner. However, if this preferred path becomes unavailable, the multipath driver on the host fails over to an alternate path. This failover might cause the volume ownership to change to the alternate controller. After you have resolved the condition that caused the failover, some hosts might automatically move the volume ownership back to the preferred controller owner, but in some cases, you might need to manually redistribute the volumes.

You can redistribute volumes by selecting **Storage › Volumes › More › Redistribute volumes**.

Change volume ownership

Changing the ownership of a volume changes the preferred controller ownership of the volume. The preferred controller owner of a volume is listed under **Storage › Volumes › View/Edit Settings › Advanced tab**.

You can change the ownership of a volume by selecting **Storage › Volumes › More › Change ownership**.

Mirroring and volume ownership

If the primary volume of the mirrored pair is owned by controller A, then the secondary volume will also be owned by controller A of the remote storage array. Changing the primary volume's owner will automatically change the owner of the secondary volume to ensure that both volumes are owned by the same controller. Current ownership changes on the primary side automatically propagate to corresponding current ownership changes on the secondary side.

If a mirror consistency group contains a local secondary volume and the controller ownership is changed, the secondary volume is automatically transferred back to its original controller owner on the first write operation. You cannot change the controller ownership of a secondary volume by using the **Change ownership** option.

Copy volume and volume ownership

During a copy volume operation, the same controller must own both the source volume and the target volume. Sometimes both volumes do not have the same preferred controller when the copy volume operation starts. Therefore, the ownership of the target volume is automatically transferred to the preferred controller of the source volume. When the volume copy is completed or is stopped, ownership of the target volume is restored to its preferred controller.

If ownership of the source volume is changed during the copy volume operation, ownership of the target volume is also changed. Under certain operating system environments, it might be necessary to reconfigure the multipath host driver before an I/O path can be used. (Some multipath drivers require an edit to recognize the I/O path. Refer to your driver documentation for more information.)

Change cache settings

Cache memory is an area of temporary volatile storage (RAM) on the controller that has a faster access time than the drive media. If you use cache memory, you can increase overall I/O performance because of these reasons:

- Data requested from the host for a read might already be in the cache from a previous operation, thus eliminating the need for drive access.
- Write data is written initially to the cache, which frees the application to continue instead of waiting for the data to be written to the drive.

Select **Storage › Volumes › More › Change cache settings** to change the following cache settings:

- **Read and write caching** — The read cache is a buffer that stores data that has been read from the drives. The data for a read operation might already be in the cache from a previous operation, which eliminates the need to access the drives. The data stays in the read cache until it is flushed.

The write cache is a buffer that stores data from the host that has not yet been written to the drives. The data stays in the write cache until it is written to the drives. Write caching can increase I/O performance.

- **Write caching with mirroring** — Write caching with mirroring occurs when the data written to the cache memory of one controller is also written to the cache memory of the other controller. Therefore, if one

controller fails, the other can complete all outstanding write operations. Write cache mirroring is available only if write caching is enabled and two controllers are present. Write caching with mirroring is the default setting at volume creation.

- **Write caching without batteries** — The write caching without batteries setting lets write caching continue even when the batteries are missing, failed, discharged completely, or not fully charged. Choosing write caching without batteries is not typically recommended, because data might be lost if power is lost. Typically, write caching is turned off temporarily by the controller until the batteries are charged or a failed battery is replaced.

This setting is available only if you enabled write caching. This setting is not available for thin volumes.

- **Dynamic read cache prefetch** — Dynamic cache read prefetch allows the controller to copy additional sequential data blocks into the cache while it is reading data blocks from a drive to the cache. This caching increases the chance that future requests for data can be filled from the cache. Dynamic cache read prefetch is important for multimedia applications that use sequential I/O. The rate and amount of data that is prefetched into cache is self-adjusting based on the rate and request size of the host reads. Random access does not cause data to be prefetched into cache. This feature does not apply when read caching is disabled.

For a thin volume, dynamic cache read prefetch is always disabled and cannot be changed.

Change media scan settings

Media scans detect and repair media errors on disk blocks that are infrequently read by applications. This scan can prevent data loss from occurring if other drives in the pool or volume group fail as data for failed drives is reconstructed using redundancy information and data from other drives in the pool or volume group.

Media scans run continuously at a constant rate based on the capacity to be scanned and the scan duration. Background scans may be temporarily suspended by a higher priority background task (for example, reconstruction), but will resume at the same constant rate.

You can enable and set the duration over which the media scan runs by selecting **Storage › Volumes › More › Change media scan settings**.

A volume is scanned only when the media scan option is enabled for the storage array and for that volume. If redundancy check is also enabled for that volume, redundancy information in the volume will be checked for consistency with data, provided that the volume has redundancy. Media scan with redundancy check is enabled by default for each volume when it is created.

If an unrecoverable medium error is encountered during the scan, data will be repaired using redundancy information, if available. For example, redundancy information is available in optimal RAID 5 volumes, or in RAID 6 volumes that are optimal or only have one drive failed. If the unrecoverable error cannot be repaired using redundancy information, the data block will be added to the unreadable sector log. Both correctable and uncorrectable medium errors are reported to the event log.

If the redundancy check finds an inconsistency between data and the redundancy information, it is reported to the event log.

Capacity for volumes

The drives in your storage array provide the physical storage capacity for your data. Before you can begin storing data, you must configure the allocated capacity into logical components known as pools or volume groups. You use these storage objects to

configure, store, maintain, and preserve data on your storage array.

Using capacity to create and expand volumes

You can create volumes from either the unassigned capacity or free capacity in a pool or volume group.

- When you create a volume from unassigned capacity, you can create a pool or volume group and the volume at the same time.
- When you create a volume from free capacity, you are creating an additional volume on an already existing pool or volume group.

After you expand the volume capacity, you must manually increase the file system size to match. How you do this depends on the file system you are using. See your host operating system documentation for details.

Capacity types for thick volumes and thin volumes

You can create either thick volumes or thin volumes. Reported capacity and allocated capacity are the same for thick volumes, but are different for thin volumes.

- For a thick volume, the reported capacity of the volume is equal to the amount of physical storage capacity allocated. The entire amount of physical storage capacity must be present. The physically allocated space is equal to the space that is reported to the host.

You normally set the thick volume's reported capacity to be the maximum capacity to which you think the volume will grow. Thick volumes provide high and predictable performance for your applications mainly because all of the user capacity is reserved and allocated upon creation.

- For a thin volume, reported capacity is the capacity that is reported to the hosts, whereas allocated capacity is the amount of drive space that is currently allocated for writing data.

The reported capacity can be larger than the allocated capacity on the storage array. Thin volumes can be sized to accommodate growth without regard for currently available assets.



SANtricity System Manager does not provide an option to create thin volumes. If you want to create thin volumes, use the Command Line Interface (CLI).

Capacity limits for thick volumes

The minimum capacity for a thick volume is 1 MiB, and the maximum capacity is determined by the number and capacity of the drives in the pool or volume group.

When increasing reported capacity for a thick volume, keep the following guidelines in mind:

- You can specify up to three decimal places (for example, 65.375 GiB).
- Capacity needs to be less than (or equal to) the maximum available in the volume group.

When you create a volume, some additional capacity is pre-allocated for Dynamic Segment Size (DSS) migration. DSS migration is a feature of the software that allows you to change the segment size of a volume.

- Volumes larger than 2 TiB are supported by some host operating systems (maximum reported capacity is determined by the host operating system). In fact, some host operating systems support up to 128 TiB volumes. Refer to your host operating system documentation for additional details.

Capacity limits for thin volumes

You can create thin volumes with a large reported capacity and a relatively small allocated capacity, which is beneficial for storage utilization and efficiency. Thin volumes can help simplify storage administration because the allocated capacity can increase as the application needs change, without disrupting the application, allowing for better storage utilization.

In addition to reported capacity and allocated capacity, thin volumes also contain Written capacity. Written capacity is the amount of capacity that has been written from the reserved capacity allocated for thin volumes.

The following table lists the capacity limits for a thin volume.

Type of capacity	Minimum size	Maximum size
Reported	32 MiB	256 TiB
Allocated	4 MiB	64 TiB

For a thin volume, if the maximum reported capacity of 256 TiB has been reached, you cannot increase its capacity. Make sure the thin volume's reserved capacity is set to a size larger than the maximum reported capacity.

System Manager automatically expands the allocated capacity based on the allocated capacity limit. The allocated capacity limit allows you to limit the thin volume's automatic growth below the reported capacity. When the amount of data written gets close to the allocated capacity, you can change the allocated capacity limit.

To change the allocated capacity limit, select **Storage › Volumes › Thin Volume Monitoring tab › Change Limit**.

Because System Manager does not allocate the full capacity when it creates a thin volume, insufficient free capacity might exist in the pool. Insufficient space can block writes to the pool, not only for the thin volumes, but also for other operations that require capacity from the pool (for example, snapshot images or snapshot volumes). However, you can still perform read operations from the pool. If this situation occurs, you receive an alert threshold warning.

Thin volume monitoring

You can monitor thin volumes for space and generate appropriate alerts to prevent out-of-capacity conditions.

Thin-provisioned environments can allocate more logical space than they have underlying physical storage. You can select **Storage › Volumes › Thin Volume Monitoring** tab to monitor how much growth your thin volumes have before they reach the allocated capacity maximum limit.

You can use the Thin Monitoring view to perform the following actions:

- Define the limit that restricts the allocated capacity to which a thin volume can automatically expand.
- Set the percentage point at which an alert (warning threshold exceeded) is sent to the Notifications area on the Home page when a thin volume is near the maximum allocated capacity limit.

To increase capacity for a thin volume, increase its reported capacity.



SANtricity System Manager does not provide an option to create thin volumes. If you want to create thin volumes, use the Command Line Interface (CLI).



Thin volumes are not available on the EF600 storage system.

Comparison between thick volumes and thin volumes

A thick volume is always fully-provisioned, which means that all of the capacity is allocated when the volume is created. A thin volume is always thinly-provisioned, which means that the capacity is allocated as the data is being written to the volume.

You can create thick volumes from either a pool or volume group. You can create thin volumes only from a pool, not from a volume group.



SANtricity System Manager does not provide an option to create thin volumes. If you want to create thin volumes, use the Command Line Interface (CLI).



Thin volumes are not available on the EF600 storage system.

Volume type	Description
Thick volumes	<ul style="list-style-type: none">• With thick volumes, a large amount of storage space is provided in advance in anticipation of future storage needs.• Thick volumes are created with the entire size of the volume pre-allocated on physical storage at the time the volume is created. This pre-allocation means that creating a 100 GiB volume actually consumes 100 GiB of allocated capacity on your drives. However, the space might remain unused, causing under-utilization of storage capacity.• When creating thick volumes, make sure not to over-allocate capacity for a single volume. Over-allocating capacity for a single volume can quickly consume all the physical storage in your system.• Keep in mind that storage capacity is also required for copy services (snapshot images, snapshot volumes, volume copies, and asynchronous mirroring), so do not allocate all of the capacity to thick volumes. Insufficient space can block writes to the pool or volume group. You receive a free-capacity alert threshold warning if this situation occurs.

Volume type	Description
Thin volumes	<ul style="list-style-type: none"> • Unlike thick volumes, space required for the thin volume is not allocated during creation, but is supplied, on demand at a later time. • A thin volume lets you over-allocate its size. That is, you can assign a LUN size that is larger than the size of the volume. You can then expand the volume as needed (if necessary, adding drives in the process) without expanding the size of the LUN, and therefore without disconnecting users. • You can use thin provisioning block space reclamation (UNMAP) to reclaim blocks of a thin-provisioned volume on the storage array through a host-issued SCSI UNMAP command. A storage array that supports thin provisioning can re-purpose the reclaimed space to satisfy allocation requests for some other thin provisioned volume within the same storage array, which allows better reporting of disk space consumption and more efficient use of resources.

Thin volume restrictions

Thin volumes support all of the operations as thick volumes, with the following exceptions:

- You cannot change the segment size of a thin volume.
- You cannot enable the pre-read redundancy check for a thin volume.
- You cannot use a thin volume as the target volume in a Copy Volume operation.
- You cannot use a thin volume in a snapshot operation.
- You can change a thin volume's allocated capacity limit and warning threshold only on the primary side of an asynchronous mirrored pair. Any changes to these parameters on the primary side are automatically propagated to the secondary side.

Copy Volume function

The Copy Volume function enables you to create a point-in-time copy of a volume by creating two separate volumes, the source volume and the target volume, on the same storage array. This function performs a byte-by-byte copy from the source volume to the target volume, making the data on the target volume identical to the data on the source volume.

Data copying for greater access

As your storage requirements for a volume change, you can use the Copy Volume function to copy data from pools or volume groups that use smaller capacity drives to pools or volume groups that use larger capacity drives. For example, you can use the Copy Volume function to do the following:

- Move data to larger drives.
- Change to drives with a higher data transfer rate.
- Change to drives using new technologies for higher performance.
- Change a thin volume to a thick volume.

Change a thin volume to a thick volume

If you want to change a thin volume to a thick volume, use the Copy Volume operation to create a copy of the thin volume. The target of a Copy Volume operation is always a thick volume.



SANtricity System Manager does not provide an option to create thin volumes. If you want to create thin volumes, use the Command Line Interface (CLI).

Backup data

The Copy Volume function lets you back up a volume by copying data from one volume to another volume on the same storage array. You can use the target volume as a backup for the source volume, for system testing, or to back up to another device, such as a tape drive.

Restore snapshot volume data to the base volume

If you need to restore data to the base volume from its associated snapshot volume, you can use the Copy Volume function to copy data from the snapshot volume to the base volume. You can create a volume copy of the data on the snapshot volume, and then copy the data to the base volume.

Source and target volumes

The following table specifies the types of volumes that can be used for source and target volumes with the Copy Volume function.

Volume type	Offline volume copy source volume	Online volume copy source volume	Online and offline target volume
Thick volume in a pool	Yes	Yes	Yes
Thick volume in a volume group	Yes	Yes	Yes
Thin volume	Yes	Yes	No
Snapshot volume	Yes1	No	No
Snapshot base volume	Yes	No	No
Remote mirror primary volume	Yes2	No	Yes

Types of Copy Volume operations

You can perform either an *offline* Copy Volume operation or an *online* Copy Volume operation. An offline operation reads data from a source volume and copies it to a target volume. An online operation uses a snapshot volume as the source and copies its data to a target volume.

To ensure data integrity, all I/O activity to the target volume is suspended during either type of Copy Volume

operation. This suspension occurs because the state of data on the target volume is inconsistent until the procedure is complete.

The offline and online Copy Volume operations are described below.

Offline Copy Volume operation

The offline Copy Volume relationship is between a source volume and a target volume. An offline copy reads data from the source volume and copies it to a target volume, while suspending all updates to the source volume with the copy in progress. All updates to the source volume are suspended to prevent chronological inconsistencies from being created on the target volume.

What you need to know about offline copy operations	
Read and write requests	<ul style="list-style-type: none">• Source volumes that are participating in an offline copy are available for read-only I/O activity while a Copy Volume operation has a status of In Progress or Pending.• Write requests are allowed after the offline copy has completed.• To prevent write-protected error messages, do not access a source volume that is participating in a Copy Volume operation with a status of In Progress.
Journaling file system	<ul style="list-style-type: none">• If the source volume has been formatted with a journaling file system, any attempt to issue a read request to the source volume might be rejected by the storage array controllers, and an error message might appear.• The journaling file system driver issues a write request before it attempts to issue the read request. The controller rejects the write request, and the read request might not be issued due to the rejected write request. This condition might result in an error message appearing, which indicates that the source volume is write protected.• To prevent this issue from occurring, do not attempt to access a source volume that is participating in an offline copy while the Copy Volume operation has a status of In Progress.

Online Copy Volume operation

The online Copy Volume relationship is between a snapshot volume and a target volume. You can initiate a Copy Volume operation while the source volume is online and available for data writes. This function is achieved by creating a snapshot of the volume and using the snapshot as the actual source volume for the copy.

When you initiate a Copy Volume operation for a source volume, System Manager creates a snapshot image

of the base volume and a copy relationship between the snapshot image of the base volume and a target volume. Using the snapshot image as the source volume allows the storage array to continue to write to the source volume while the copy is in progress.

During an online copy operation, a performance impact is experienced due to the copy-on-write procedure. After the online copy completes, the base volume performance is restored.

What you need to know about online copy operations	
What kind of volumes can be used?	<ul style="list-style-type: none">• The volume for which the point-in-time image is created is known as the base volume and must be a standard volume or a thin volume on the storage array.• A target volume can be a standard volume in a volume group or a standard volume in a pool. A target volume cannot be a thin volume or a base volume in a snapshot group.• You can use the online Copy Volume function to copy data from a thin volume to a standard volume in a pool that resides within the same storage array. But you cannot use the Copy Volume function to copy data from a standard volume to a thin volume.
Base volume performance	<ul style="list-style-type: none">• If the snapshot volume that is used as the copy source is active, the base volume performance is degraded due to copy-on-write operations. When the copy is complete, the snapshot is disabled, and the base volume performance is restored. Although the snapshot is disabled, the reserved capacity volume and copy relationship remain intact.
Types of volumes created	<ul style="list-style-type: none">• A snapshot volume and a reserved capacity volume are created during the online copy operation.• The snapshot volume is not an actual volume containing data; rather, it is a reference to the data that was contained on a volume at a specific time.• For each snapshot that is taken, a reserved capacity volume is created to hold the data for the snapshot. The reserved capacity volume is used only to manage the snapshot image.

What you need to know about online copy operations

Reserved capacity volume

- Before a data block on the source volume is modified, the contents of the block to be modified are copied to the reserved capacity volume for safekeeping.
- Because the reserved capacity volume stores copies of the original data in those data blocks, further changes to those data blocks write only to the source volume.
- The online copy operation uses less disk space than a full physical copy because the only data blocks that are stored in the reserved capacity volume are those that have changed since the time of the snapshot.

How to

Create storage

Create workloads

You can create workloads for any type of application.

About this task

A workload is a storage object that supports an application. You can define one or more workloads, or instances, per application. For some applications, System Manager configures the workload to contain volumes with similar underlying volume characteristics. These volume characteristics are optimized based on the type of application the workload supports. For example, if you create a workload that supports a Microsoft SQL Server application and then subsequently create volumes for that workload, the underlying volume characteristics are optimized to support Microsoft SQL Server.

System Manager recommends an optimized volume configuration only for the following application types:

- Microsoft® SQL Server™
- Microsoft® Exchange Server™
- Video Surveillance
- VMware ESXi™ (for volumes to be used with Virtual Machine File System)

Keep these guidelines in mind:

- *When using an application-specific workload*, the system recommends an optimized volume configuration to minimize contention between application workload I/O and other traffic from your application instance. You can review the recommended volume configuration, and then edit, add, or delete the system-recommended volumes and characteristics using the **Add/Edit Volumes** dialog box.
- *When using other application types*, you manually specify the volume configuration using the **Add/Edit Volumes** dialog box.

Steps

1. Select **Storage > Volumes**.

2. Select **Create › Workload**.

The **Create Application Workload** dialog box appears.

3. Use the drop-down list to select the type of application that you want to create the workload for and then type a workload name.
4. Click **Create**.

After you finish

You are ready to add storage capacity to the workload you created. Use the **Create Volume** option to create one or more volumes for an application, and to allocate specific amounts of capacity to each volume.

Create volumes

You create volumes to add storage capacity to an application-specific workload, and to make the created volumes visible to a specific host or host cluster. In addition, the volume creation sequence provides options to allocate specific amounts of capacity to each volume you want to create.

About this task

Most application types default to a user-defined volume configuration. Some application types have a smart configuration applied at volume creation. For example, if you are creating volumes for Microsoft Exchange application, you are asked how many mailboxes you need, what your average mailbox capacity requirements are, and how many copies of the database you want. System Manager uses this information to create an optimal volume configuration for you, which can be edited as needed.



If you want to mirror a volume, first create the volumes that you want to mirror, and then use the **Storage › Volumes › Copy Services › Mirror a volume asynchronously** option.

The process to create a volume is a multi-step procedure.

Step 1: Select host for a volume

You create volumes to add storage capacity to an application-specific workload, and to make the created volumes visible to a specific host or host cluster. In addition, the volume creation sequence provides options to allocate specific amounts of capacity to each volume you want to create.

Before you begin

- Valid hosts or host clusters exist under the **Hosts** tile.
- Host port identifiers have been defined for the host.
- Before creating a DA-enabled volume, the host connection you are planning to use must support DA. If any of the host connections on the controllers in your storage array do not support DA, the associated hosts cannot access data on DA-enabled volumes.

About this task

Keep these guidelines in mind when you assign volumes:

- A host's operating system can have specific limits on how many volumes the host can access. Keep this limitation in mind when you create volumes for use by a particular host.
- You can define one assignment for each volume in the storage array.

- Assigned volumes are shared between controllers in the storage array.
- The same logical unit number (LUN) cannot be used twice by a host or a host cluster to access a volume. You must use a unique LUN.



Assigning a volume to a host will fail if you try to assign a volume to a host cluster that conflicts with an established assignment for a host in the host clusters.

Steps

1. Select **Storage > Volumes**.
2. Select **Create > Volume**.

The **Create Volumes** dialog box appears.

3. From the drop-down list, select a specific host or host cluster to which you want to assign volumes, or choose to assign the host or host cluster at a later time.
4. To continue the volume creation sequence for the selected host or host cluster, click **Next**, and go to [Step 2: Select a workload for a volume](#).

The **Select Workload** dialog box appears.

Step 2: Select a workload for a volume

Select a workload to customize the storage array configuration for a specific application, such as Microsoft SQL Server, Microsoft Exchange, Video Surveillance applications, or VMware. You can select "Other application" if the application you intend to use on this storage array is not listed.

About this task

This task describes how to create volumes for an existing workload.

- *When you are creating volumes using an application-specific workload*, the system may recommend an optimized volume configuration to minimize contention between application workload I/O and other traffic from your application instance. You can review the recommended volume configuration and edit, add, or delete the system-recommended volumes and characteristics using the **Add/Edit Volumes** dialog box.
- *When you are creating volumes using "Other" applications* (or applications without specific volume creation support), you manually specify the volume configuration using the **Add/Edit Volumes** dialog box.

Steps

1. Do one of the following:
 - Select the **Create volumes for an existing workload** option to create volumes for an existing workload.
 - Select the **Create a new workload** option to define a new workload for a supported application or for "Other" applications.
 - From the drop-down list, select the name of the application you want to create the new workload for.

Select one of the "Other" entries if the application you intend to use on this storage array is not listed.

- Enter a name for the workload you want to create.

2. Click **Next**.
3. If your workload is associated with a supported application type, enter the information requested; otherwise, go to [Step 3: Add or edit volumes](#).

Step 3: Add or edit volumes

System Manager may suggest a volume configuration based on the application or workload you selected. This volume configuration is optimized based on the type of application the workload supports. You can accept the recommended volume configuration or you can edit it as needed. If you selected one of the "Other" applications, you must manually specify the volumes and characteristics you want to create.

Before you begin

- The pools or volume groups must have sufficient free capacity.
- To create a Data Assurance (DA)-enabled volume, the host connection you are planning to use must support DA.

Selecting a DA capable pool or volume group

If you want to create a DA-enabled volume, select a pool or volume group that is DA capable (look for **Yes** next to "DA" in the pool and volume group candidates table).

DA capabilities are presented at the pool and volume group level in System Manager. DA protection checks for and corrects errors that might occur as data is transferred through the controllers down to the drives. Selecting a DA-capable pool or volume group for the new volume ensures that any errors are detected and corrected.

If any of the host connections on the controllers in your storage array do not support DA, the associated hosts cannot access data on DA-enabled volumes.



DA is not supported by iSCSI over TCP/IP, or by the SRP over InfiniBand.

- To create a secure-enabled volume, a security key must be created for the storage array.

Selecting a secure-capable pool or volume group

If you want to create a secure-enabled volume, select a pool or volume group that is secure capable (look for **Yes** next to "Secure-capable" in the pool and volume group candidates table).

Drive security capabilities are presented at the pool and volume group level in System Manager. Secure-capable drives prevent unauthorized access to the data on a drive that is physically removed from the storage array. A secure-enabled drive encrypts data during writes and decrypts data during reads using a unique *encryption key*.

A pool or volume group can contain both secure-capable and non-secure-capable drives, but all drives must be secure-capable to use their encryption capabilities.

About this task

You create volumes from pools or volume groups. The **Add/Edit Volumes** dialog box shows all eligible pools and volume groups on the storage array. For each eligible pool and volume group, the number of drives available and the total free capacity appears.

For some application-specific workloads, each eligible pool or volume group shows the proposed capacity based on the suggested volume configuration and shows the remaining free capacity in GiB. For other workloads, the proposed capacity appears as you add volumes to a pool or volume group and specify the reported capacity.

Steps

1. Choose one of these actions based on whether you selected Other or an application-specific workload:
 - **Other** — Click **Add new volume** in each pool or volume group that you want to use to create one or more volumes.

Field Details

Field	Description
Volume Name	A volume is assigned a default name by System Manager during the volume creation sequence. You can either accept the default name or provide a more descriptive one indicating the type of data stored in the volume.
Reported Capacity	<p>Define the capacity of the new volume and the capacity units to use (MiB, GiB, or TiB). For Thick volumes, the minimum capacity is 1 MiB, and the maximum capacity is determined by the number and capacity of the drives in the pool or volume group.</p> <p>Keep in mind that storage capacity is also required for copy services (snapshot images, snapshot volumes, volume copies, and remote mirrors); therefore, do not allocate all of the capacity to standard volumes.</p> <p>Capacity in a pool is allocated in 4-GiB increments. Any capacity that is not a multiple of 4 GiB is allocated but not usable. To make sure that the entire capacity is usable, specify the capacity in 4-GiB increments. If unusable capacity exists, the only way to regain it is to increase the capacity of the volume.</p>

Field	Description
Segment Size	<p>Shows the setting for segment sizing, which only appears for volumes in a volume group. You can change the segment size to optimize performance.</p> <p>Allowed segment size transitions — System Manager determines the segment size transitions that are allowed. Segment sizes that are inappropriate transitions from the current segment size are unavailable on the drop-down list. Allowed transitions usually are double or half of the current segment size. For example, if the current volume segment size is 32 KiB, a new volume segment size of either 16 KiB or 64 KiB is allowed.</p> <p>SSD Cache-enabled volumes — You can specify a 4-KiB segment size for SSD Cache-enabled volumes. Make sure you select the 4-KiB segment size only for SSD Cache-enabled volumes that handle small-block I/O operations (for example, 16 KiB I/O block sizes or smaller). Performance might be impacted if you select 4 KiB as the segment size for SSD Cache-enabled volumes that handle large block sequential operations.</p> <p>Amount of time to change segment size — The amount of time to change a volume's segment size depends on these variables:</p> <ul style="list-style-type: none"> • The I/O load from the host • The modification priority of the volume • The number of drives in the volume group • The number of drive channels • The processing power of the storage array controllers <p>When you change the segment size for a volume, I/O performance is affected, but your data remains available.</p>
Secure-capable	<p>Yes appears next to "Secure-capable" only if the drives in the pool or volume group are secure-capable.</p> <p>Drive Security prevents unauthorized access to the data on a drive that is physically removed from the storage array. This option is available only when the Drive Security feature has been enabled, and a security key is set up for the storage array.</p> <p>A pool or volume group can contain both secure-capable and non-secure-capable drives, but all drives must be secure-capable to use their encryption capabilities.</p>

Field	Description
DA	<p>Yes appears next to "DA" only if the drives in the pool or volume group support Data Assurance (DA).</p> <p>DA increases data integrity across the entire storage system. DA enables the storage array to check for errors that might occur as data is transferred through the controllers down to the drives. Using DA for the new volume ensures that any errors are detected.</p>

- **Application-specific workload** — Either click **Next** to accept the system-recommended volumes and characteristics for the selected workload, or click **Edit Volumes** to change, add, or delete the system-recommended volumes and characteristics for the selected workload.

Field Details

Field	Description
Volume Name	A volume is assigned a default name by System Manager during the volume creation sequence. You can either accept the default name or provide a more descriptive one indicating the type of data stored in the volume.
Reported Capacity	<p>Define the capacity of the new volume and the capacity units to use (MiB, GiB, or TiB). For Thick volumes, the minimum capacity is 1 MiB, and the maximum capacity is determined by the number and capacity of the drives in the pool or volume group.</p> <p>Keep in mind that storage capacity is also required for copy services (snapshot images, snapshot volumes, volume copies, and remote mirrors); therefore, do not allocate all of the capacity to standard volumes.</p> <p>Capacity in a pool is allocated in 4-GiB increments. Any capacity that is not a multiple of 4 GiB is allocated but not usable. To make sure that the entire capacity is usable, specify the capacity in 4-GiB increments. If unusable capacity exists, the only way to regain it is to increase the capacity of the volume.</p>
Volume Type	Volume type indicates the type of volume that was created for an application-specific workload.

Field	Description
Segment Size	<p>Shows the setting for segment sizing, which only appears for volumes in a volume group. You can change the segment size to optimize performance.</p> <p>Allowed segment size transitions — System Manager determines the segment size transitions that are allowed. Segment sizes that are inappropriate transitions from the current segment size are unavailable on the drop-down list. Allowed transitions usually are double or half of the current segment size. For example, if the current volume segment size is 32 KiB, a new volume segment size of either 16 KiB or 64 KiB is allowed.</p> <p>SSD Cache-enabled volumes — You can specify a 4-KiB segment size for SSD Cache-enabled volumes. Make sure you select the 4-KiB segment size only for SSD Cache-enabled volumes that handle small-block I/O operations (for example, 16 KiB I/O block sizes or smaller). Performance might be impacted if you select 4 KiB as the segment size for SSD Cache-enabled volumes that handle large block sequential operations.</p> <p>Amount of time to change segment size — The amount of time to change a volume's segment size depends on these variables:</p> <ul style="list-style-type: none"> • The I/O load from the host • The modification priority of the volume • The number of drives in the volume group • The number of drive channels • The processing power of the storage array controllers When you change the segment size for a volume, I/O performance is affected, but your data remains available.
Secure-capable	<p>Yes appears next to "Secure-capable" only if the drives in the pool or volume group are secure-capable.</p> <p>Drive security prevents unauthorized access to the data on a drive that is physically removed from the storage array. This option is available only when the drive security feature has been enabled, and a security key is set up for the storage array.</p> <p>A pool or volume group can contain both secure-capable and non-secure-capable drives, but all drives must be secure-capable to use their encryption capabilities.</p>

Field	Description
DA	<p>Yes appears next to "DA" only if the drives in the pool or volume group support Data Assurance (DA).</p> <p>DA increases data integrity across the entire storage system. DA enables the storage array to check for errors that might occur as data is transferred through the controllers down to the drives. Using DA for the new volume ensures that any errors are detected.</p>

2. To continue the volume creation sequence for the selected application, click **Next**, and go to [Step 4: Review volume configuration](#).

Step 4: Review volume configuration

Review a summary of the volumes you intend to create and make any necessary changes.

Steps

1. Review the volumes you want to create. Click **Back** to make any changes.
2. When you are satisfied with your volume configuration, click **Finish**.

Results

System Manager creates the new volumes in the selected pools and volume groups, and then displays the new volumes in the All Volumes table.

After you finish

- Perform any operating system modifications necessary on the application host so that the applications can use the volume.
- Run either the host-based `hot_add` utility or an operating system-specific utility (available from a third-party vendor), and then run the `SMdevices` utility to correlate volume names with host storage array names.

The `hot_add` utility and the `SMdevices` utility are included as part of the `SMutils` package. The `SMutils` package is a collection of utilities to verify what the host sees from the storage array. It is included as part of the SANtricity software installation.

Manage volumes

Increase capacity of a volume

You can increase the reported capacity (the capacity reported to hosts) of a volume by using the free capacity that is available in the pool or volume group.

Before you begin

- Enough free capacity is available in the volume's associated pool or volume group.
- The volume is Optimal and not in any state of modification.
- The maximum reported capacity of 256 TiB has not been reached for thin volumes.
- No hot spare drives are in use in the volume. (Applies only to volumes in volume groups.)

About this task

Keep in mind any future capacity requirements that you might have for other volumes in this pool or volume group. Make sure that you allow enough free capacity to create snapshot images, snapshot volumes, or remote mirrors.



Increasing the capacity of a volume is supported only on certain operating systems. If you increase the volume capacity on a host operating system that is unsupported, the expanded capacity is unusable, and you cannot restore the original volume capacity.

Steps

1. Select **Storage > Volumes**.
2. Select the volume for which you want to increase capacity, and then select **Increase Capacity**.

The **Confirm Increase Capacity** dialog box appears.

3. Select **Yes** to continue.

The **Increase Reported Capacity** dialog box appears.

This dialog box displays the volume's current reported capacity and the free capacity available in the volume's associated pool or volume group.

4. Use the **Increase reported capacity by adding...** box to add capacity to the current available reported capacity. You can change the capacity value to display in either mebibytes (MiB), gibibytes (GiB), or tebibytes (TiB).
5. Click **Increase**.

Results

- System Manager increases the volume's capacity based on your selection.
- Select **Home > View Operations in Progress** to view the progress of the increase capacity operation that is currently running for the selected volume. This operation can be lengthy and could affect system performance.

After you finish

After you expand the volume capacity, you must manually increase the file system size to match. How you do this depends on the file system you are using. See your host operating system documentation for details.

Change settings for a volume

You can change a volume's settings such as its name, host assignment, segment size, modification priority, caching, and so on.

Before you begin

The volume you want to change is in Optimal status.

Steps


1. Select **Storage > Volumes**.
2. Select the volume that you want to change, and then select **View/Edit Settings**.

The **Volume Settings** dialog box appears. The configuration settings for the volume you selected appear

in this dialog box.

3. Select the **Basic** tab to change the volume's name and host assignment.

Field Details

Setting	Description
Name	Displays the name of the volume. Change the name of a volume when the current name is no longer meaningful or applicable.
Capacities	<p>Displays the reported and allocated capacity for the selected volume.</p> <p>Reported capacity and allocated capacity are the same for thick volumes, but are different for thin volumes. For a thick volume, the physically allocated space is equal to the space that is reported to the host. For a thin volume, reported capacity is the capacity that is reported to the hosts, whereas allocated capacity is the amount of drive space that is currently allocated for writing data.</p>
Pool / Volume group	Displays the name and RAID level of the pool or volume group. Indicates whether the pool or volume group is secure-capable and secure-enabled.
Host	<p>Displays the volume assignment. You assign a volume to a host or host cluster so it can be accessed for I/O operations. This assignment grants a host or host cluster access to a particular volume or to a number of volumes in a storage array.</p> <ul style="list-style-type: none"> • Assigned to — Identifies the host or host cluster that has access to the selected volume. • LUN — A logical unit number (LUN) is the number assigned to the address space that a host uses to access a volume. The volume is presented to the host as capacity in the form of a LUN. Each host has its own LUN address space. Therefore, the same LUN can be used by different hosts to access different volumes. <div>  <p>For NVMe interfaces, this column displays Namespace ID. A namespace is NVM storage that is formatted for block access. It is analogous to a logical unit in SCSI, which relates to a volume in the storage array. The namespace ID is the NVMe controller's unique identifier for the namespace, and can be set to a value between 1 and 255. It is analogous to a logical unit number (LUN) in SCSI.</p> </div>
Identifiers	<p>Displays the identifiers for the selected volume.</p> <ul style="list-style-type: none"> • World-wide identifier (WWID) — A unique hexadecimal identifier for the volume. • Extended unique identifier (EUI) — An EUI-64 identifier for the volume. • Subsystem identifier (SSID) — The storage array subsystem identifier of a volume.

4. Select the **Advanced** tab to change additional configuration settings for a volume in a pool or in a volume group.

Field Details

Setting	Description
Application & workload information	<p>During volume creation, you can create application-specific workloads or other workloads. If applicable, the workload name, application type, and volume type appears for the selected volume.</p> <p>You can change the workload name if desired.</p>
Quality of Service settings	<p>Permanently disable data assurance — This setting appears only if the volume is Data Assurance (DA)-enabled. DA checks for and corrects errors that might occur as data is transferred through the controllers down to the drives. Use this option to permanently disable DA on the selected volume. When disabled, DA cannot be re-enabled on this volume.</p> <p>Enable pre-read redundancy check — This setting appears only if the volume is a thick volume. Pre-read redundancy checks determine whether the data on a volume is consistent any time a read is performed. A volume that has this feature enabled returns read errors if the data is determined to be inconsistent by the controller firmware.</p>
Controller ownership	<p>Defines the controller that is designated to be the owning, or primary, controller of the volume.</p> <p>Controller ownership is very important and should be planned carefully. Controllers should be balanced as closely as possible for total I/Os.</p>

Setting	Description
Segment sizing	<p>Shows the setting for segment sizing, which appears only for volumes in a volume group. You can change the segment size to optimize performance.</p> <p>Allowed segment size transitions — System Manager determines the segment size transitions that are allowed. Segment sizes that are inappropriate transitions from the current segment size are unavailable on the drop-down list. Allowed transitions usually are double or half of the current segment size. For example, if the current volume segment size is 32 KiB, a new volume segment size of either 16 KiB or 64 KiB is allowed.</p> <p>SSD Cache-enabled volumes — You can specify a 4-KiB segment size for SSD Cache-enabled volumes. Make sure you select the 4-KiB segment size only for SSD Cache-enabled volumes that handle small-block I/O operations (for example, 16 KiB I/O block sizes or smaller). Performance might be impacted if you select 4 KiB as the segment size for SSD Cache-enabled volumes that handle large block sequential operations.</p> <p>Amount of time to change segment size — The amount of time to change a volume's segment size depends on these variables:</p> <ul style="list-style-type: none"> • The I/O load from the host • The modification priority of the volume • The number of drives in the volume group • The number of drive channels • The processing power of the storage array controllers When you change the segment size for a volume, I/O performance is affected, but your data remains available.
Modification priority	<p>Shows the setting for modification priority, which only appears for volumes in a volume group.</p> <p>The modification priority defines how much processing time is allocated for volume modification operations relative to system performance. You can increase the volume modification priority, although this might affect system performance.</p> <p>Move the slider bars to select a priority level.</p> <p>Modification priority rates — The lowest priority rate benefits system performance, but the modification operation takes longer. The highest priority rate benefits the modification operation, but system performance might be compromised.</p>
Caching	Shows the caching setting, which you can change to impact the overall I/O performance of a volume.

Setting	Description
SSD Cache	<p>Shows the SSD Cache setting, which you can enable on compatible volumes as a way to improve read-only performance. Volumes are compatible if they share the same Drive Security and Data Assurance capabilities.</p> <p>The SSD Cache feature uses a single or multiple Solid State Disks (SSDs) to implement a read cache. Application performance is improved because of the faster read times for SSDs. Because the read cache is in the storage array, caching is shared across all applications using the storage array. Simply select the volume that you want to cache, and then caching is automatic and dynamic.</p>

5. Click **Save**.

Results

System Manager changes the volume's settings based on your selections.

After you finish

Select **Home > View Operations in Progress** to view the progress of the change operations that are currently running for the selected volume.

Initialize volumes

A volume is automatically initialized when it is first created. However, the Recovery Guru might advise that you manually initialize a volume to recover from certain failure conditions. Use this option only under the guidance of technical support. You can select one or more volumes to initialize.

Before you begin

- All I/O operations have been stopped.
- Any devices or file systems on the volumes you want to initialize must be unmounted.
- The volume is in Optimal status and no modification operations are in progress on the volume.



You cannot cancel the operation after it starts. All volume data is erased. Do not try this operation unless the Recovery Guru advises you to do so. Contact technical support before you begin this procedure.

About this task

When you initialize a volume, the volume keeps its WWN, host assignments, allocated capacity, and reserved capacity settings. It also keeps the same Data Assurance (DA) settings and security settings.

The following types of volumes *cannot* be initialized:

- Base volume of a snapshot volume
- Primary volume in a mirror relationship
- Secondary volume in a mirror relationship

- Source volume in a volume copy
- Target volume in a volume copy
- Volume that already has an initialization in progress

This topic applies only to standard volumes created from pools or volume groups.

Steps

1. Select **Storage › Volumes**.
2. Select any volume, and then select **More › Initialize volumes**.

The **Initialize Volumes** dialog box appears. All volumes on the storage array appear in this dialog box.

3. Select one or more volumes that you want to initialize, and confirm that you want to perform the operation.

Results

System Manager performs the following actions:

- Erases all data from the volumes that were initialized.
- Clears the block indices, which causes unwritten blocks to be read as if they are zero-filled (the volume appears to be completely empty).

Select **Home › View Operations in Progress** to view the progress of the initialize operation that is currently running for the selected volume. This operation can be lengthy and could affect system performance.

Redistribute volumes

You redistribute volumes to move volumes back to their preferred controller owners. Typically, multipath drivers move volumes from their preferred controller owner when a problem occurs along the data path between the host and storage array.

Before you begin

- The volumes you want to redistribute are not in use, or I/O errors will occur.
- A multipath driver is installed on all hosts using the volumes you want to redistribute, or I/O errors will occur.

If you want to redistribute volumes without a multipath driver on the hosts, all I/O activity to the volumes *while the redistribution operation is in progress* must be stopped to prevent application errors.

About this task

Most host multipath drivers attempt to access each volume on a path to its preferred controller owner. However, if this preferred path becomes unavailable, the multipath driver on the host fails over to an alternate path. This failover might cause the volume ownership to change to the alternate controller. After you have resolved the condition that caused the failover, some hosts might automatically move the volume ownership back to the preferred controller owner, but in some cases, you might need to manually redistribute the volumes.

Steps

1. Select **Storage › Volumes**.
2. Select **More › Redistribute volumes**.

The **Redistribute Volumes** dialog box appears. All volumes on the storage array whose preferred controller owner does not match its current owner appear in this dialog box.

3. Select one or more volumes that you want to redistribute, and confirm that you want to perform the operation.

Results

System Manager moves the selected volumes to their preferred controller owners or you might see a **Redistribute Volumes Unnecessary** dialog box.

Change controller ownership of a volume

You can change the preferred controller ownership of a volume, so that I/O for host applications is directed through the new path.

Before you begin

If you do not use a multipath driver, any host applications that are currently using the volume must be shut down. This action prevents application errors when the I/O path changes.

About this task

You can change controller ownership for one or more volumes in a pool or volume group.

Steps

1. Select **Storage > Volumes**.
2. Select any volume, and then select **More > Change ownership**.

The **Change Volume Ownership** dialog box appears. All volumes on the storage array appear in this dialog box.

3. Use the **Preferred Owner** drop-down list to change the preferred controller for each volume that you want to change, and confirm that you want to perform the operation.

Results

- System Manager changes the controller ownership of the volume. I/O to the volume is now directed through this I/O path.
- The volume might not use the new I/O path until the multipath driver reconfigures to recognize the new path. This action usually takes less than five minutes.

Change cache settings for a volume

You can change read cache and write cache settings to impact the overall I/O performance of a volume.

About this task

Keep these guidelines in mind when you change cache settings for a volume:

- After opening the **Change Cache Settings** dialog box, you might see an icon shown next to the selected cache properties. This icon indicates that the controller has temporarily suspended caching operations.

This action might occur when a new battery is charging, when a controller has been removed, or if a mismatch in cache sizes has been detected by the controller. After the condition has cleared, the cache

properties selected in the dialog box become active. If the selected cache properties do not become active, contact technical support.

- You can change the cache settings for a single volume or for multiple volumes on a storage array. You can change the cache settings for all standard volumes or all thin volumes at the same time.


Steps

1. Select **Storage › Volumes**.
2. Select any volume, and then select **More › Change cache settings**.

The **Change Cache Settings** dialog box appears. All volumes on the storage array appear in this dialog box.


3. Select the **Basic** tab to change the settings for read caching and write caching.

Field Details

Cache setting	Description
Read Caching	The read cache is a buffer that stores data that has been read from the drives. The data for a read operation might already be in the cache from a previous operation, which eliminates the need to access the drives. The data stays in the read cache until it is flushed.
Write Caching	<div><div>The write cache is a buffer that stores data from the host that has not yet been written to the drives. The data stays in the write cache until it is written to the drives. Write caching can increase I/O performance.</div><div><div></div><div>Cache is automatically flushed after the Write caching is disabled for a volume.</div></div></div>

4. Select the **Advanced** tab to change the advanced settings for thick volumes. The advanced cache settings are available only for thick volumes.

Field Details

Cache setting	Description
Dynamic Read Cache Prefetch	<p>Dynamic cache read prefetch allows the controller to copy additional sequential data blocks into the cache while it is reading data blocks from a drive to the cache. This caching increases the chance that future requests for data can be filled from the cache. Dynamic cache read prefetch is important for multimedia applications that use sequential I/O. The rate and amount of data that is prefetched into cache is self-adjusting based on the rate and request size of the host reads. Random access does not cause data to be prefetched into cache. This feature does not apply when read caching is disabled.</p> <p>For a thin volume, dynamic cache read prefetch is always disabled and cannot be changed.</p>
Write Caching without Batteries	<p>The write caching without batteries setting lets write caching continue even when the batteries are missing, failed, discharged completely, or not fully charged. Choosing write caching without batteries is not typically recommended, because data might be lost if power is lost. Typically, write caching is turned off temporarily by the controller until the batteries are charged or a failed battery is replaced.</p> <div><p>Possible loss of data — If you select this option and do not have a universal power supply for protection, you could lose data. In addition, you could lose data if you do not have controller batteries and you enable the Write caching without batteries option.</p></div> <p>This setting is available only if you enabled write caching. This setting is not available for thin volumes.</p>
Write Caching with Mirroring	<p>Write caching with mirroring occurs when the data written to the cache memory of one controller is also written to the cache memory of the other controller. Therefore, if one controller fails, the other can complete all outstanding write operations. Write cache mirroring is available only if write caching is enabled and two controllers are present. Write caching with mirroring is the default setting at volume creation.</p> <p>This setting is available only if you enabled write caching. This setting is not available for thin volumes.</p>

5. Click **Save** to change the cache settings.

Change media scan settings for a volume

A media scan is a background operation that scans all data and redundancy information in the volume. Use this option to enable or disable the media scan settings for one or more volumes, or to change the scan duration.

Before you begin

Understand the following:

- Media scans run continuously at a constant rate based on the capacity to be scanned and the scan duration. Background scans may be temporarily suspended by a higher priority background task (e.g. reconstruction), but will resume at the same constant rate.
- A volume is scanned only when the media scan option is enabled for the storage array and for that volume. If redundancy check is also enabled for that volume, redundancy information in the volume will be checked for consistency with data, provided that the volume has redundancy. Media scan with redundancy check is enabled by default for each volume when it is created.
- If an unrecoverable medium error is encountered during the scan, data will be repaired using redundancy information, if available.

For example, redundancy information is available in optimal RAID 5 volumes, or in RAID 6 volumes that are optimal or only have one drive failed. If the unrecoverable error cannot be repaired using redundancy information, the data block will be added to the unreadable sector log. Both correctable and uncorrectable medium errors are reported to the event log.

If the redundancy check finds an inconsistency between data and the redundancy information, it is reported to the event log.

About this task

Media scans detect and repair media errors on disk blocks that are infrequently read by applications. This can prevent data loss in the event of a drive failure, as data for failed drives is reconstructed using redundancy information and data from other drives in the volume group or pool.

You can perform the following actions:

- Enable or disable background media scans for the entire storage array
- Change the scan duration for the entire storage array
- Enable or disable media scan for one or more volumes
- Enable or disable the redundancy check for one or more volumes

Steps

1. Select **Storage > Volumes**.
2. Select any volume, and then select **More > Change media scan settings**.

The **Change Drive Media Scan Settings** dialog box appears. All volumes on the storage array appear in this dialog box.

3. To enable the media scan, select the **Scan media over the course of...** check box.

Disabling the media scan check box suspends all media scan settings.

4. Specify the number of days over which you want the media scan to run.
5. Select the **Media Scan** check box for each volume you want to perform a media scan on.

System Manager enables the Redundancy Check option for each volume on which you choose to run a media scan. If there are individual volumes for which you do not want to perform a redundancy check, deselect the **Redundancy Check** check box.

6. Click **Save**.

Results

System Manager applies changes to background media scans based on your selection.

Delete volume

Typically, you delete volumes if the volumes were created with the wrong parameters or capacity, no longer meet storage configuration needs, or are snapshot images that are no longer needed for backup or application testing. Deleting a volume increases the free capacity in the pool or volume group. You can select one or more volumes to delete.

Before you begin

On the volumes that you plan to delete, make sure of the following:

- All data is backed up.
- All Input/Output (I/O) is stopped.
- Any devices and file systems are unmounted.

About this task

You cannot delete a volume that has one of these conditions:

- The volume is initializing.
- The volume is reconstructing.
- The volume is part of a volume group that contains a drive that is undergoing a copyback operation.
- The volume is undergoing a modification operation, such as a change of segment size, unless the volume is now in Failed status.
- The volume is holding any type of persistent reservation.
- The volume is a source volume or a target volume in a Copy Volume that has a status of Pending, In Progress, or Failed.



Deleting a volume causes loss of all data on those volumes.



When a volume exceeds a given size (currently 128 TB) the delete is being performed in background and the freed space may not be immediately available.

Steps

1. Select **Storage > Volumes**.
2. Click **Delete**.

The **Delete Volumes** dialog box appears.

3. Select one or more volumes that you want to delete, and confirm that you want to perform the operation.
4. Click **Delete**.

Results

System Manager performs the following actions:

- Deletes any associated snapshot images, schedules, and snapshot volumes.
- Removes any mirroring relationships.
- Increases the free capacity in the pool or volume group.

Manage applications and workloads

Add to workload

You can add one or more volumes to an existing or new workload for volumes that are not currently associated with a workload.

About this task

Volumes are not associated with a workload if they have been created using the command line interface (CLI) or if they have been migrated (imported/exported) from a different storage array.

Steps

1. Select **Storage › Volumes**.

2. Select the **Applications & Workloads** tab.

The Applications & Workloads view appears.

3. Select **Add to Workload**.

The **Select Workload** dialog box appears.

4. Do one of the following actions:

- **Add volumes to an existing workload** — Select this option to add volumes to an existing workload.

Use the drop-down list to select a workload. The workload's associated application type is assigned to the volumes you add to this workload.

- **Add volumes to a new workload** — Select this option to define a new workload for an application type and add volumes to the new workload.

5. Select **Next** to continue with the add to workload sequence.

The **Select Volumes** dialog box appears.

6. Select the volumes you want to add to the workload.
7. Review the volumes that you want to add to the selected workload.
8. When you are satisfied with your workload configuration, click **Finish**.

Change workload settings

You can change the name for a workload and view its associated application type. Change the name of a workload when the current name is no longer meaningful or applicable.

Steps

1. Select **Storage › Volumes**.

2. Select the **Applications & Workloads** tab.

The Applications & Workloads view appears.

3. Select the workload that you want to change, and then select **View/Edit Settings**.

The **Applications & Workloads Settings** dialog box appears.

4. **Optional:** Change the user-supplied name of the workload.

5. Click **Save**.

Work with copy services

Copy volume

You can copy data from one volume to another volume in the same storage array, and create a physical, point-in-time duplicate (clone) of a source volume.

Before you begin

- All I/O activity to the source volume and the target volume must be stopped.
- Any file systems on the source volume and the target volume must be unmounted.
- If you have used the target volume in a Copy Volume operation before, you no longer need that data or that you have backed up the data.

About this task

The source volume is the volume that accepts host I/O and stores application data. When a Copy Volume is started, data from the source volume is copied in its entirety to the target volume.

The target volume is a standard volume that maintains a copy of the data from the source volume. The target volume is identical to the source volume after the Copy Volume operation completes. The target volume must have the same or greater capacity as the source volume; however, it can have a different RAID level.

More about online and offline copies

Online copy

An online copy creates a point-in-time copy of any volume within a storage array, while it is still possible to write to the volume with the copy in progress. This function is achieved by creating a snapshot of the volume and using the snapshot as the actual source volume for the copy. The volume for which the point-in-time image is created is known as the base volume and it can be a standard volume or a thin volume in the storage array.

Offline copy

An offline copy reads data from the source volume and copies it to a target volume, while suspending all updates to the source volume with the copy in progress. All updates to the source volume are suspended to prevent chronological inconsistencies from being created on the target volume. The offline volume copy relationship is between a source volume and a target volume.



A Copy Volume operation overwrites data on the target volume and fails all snapshot volumes associated with the target volume, if any exist.

Steps

1. Select **Storage › Volumes**.
2. Select the volume that you want to use as the source for the Copy Volume operation, and then select **Copy Services › Copy volume**.

The **Copy Volume-Select Target** dialog box appears.

3. Select the target volume to which you want to copy the data.

The table shown in this dialog box lists all the eligible target volumes.

4. Use the slider bar to set the copy priority for the Copy Volume operation.

The copy priority determines how much of the system resources are used to complete the Copy Volume operation as compared to service I/O requests.

More about copy priority rates

There are five copy priority rates:

- Lowest
- Low
- Medium
- High
- Highest

If the copy priority is set to the lowest rate, I/O activity is prioritized, and the Copy Volume operation takes longer. If the copy priority is set to the highest rate, the Copy Volume operation is prioritized, but I/O activity for the storage array might be affected.

5. Select whether you want to create an online copy or an offline copy. To create an online copy, select the **Keep source volume online during copy operation** check box.
6. Do one of the following:
 - To perform an *online* copy operation, click **Next** to continue to the **Reserve Capacity** dialog box.
 - To perform an *offline* copy operation, click **Finish** to start the offline copy.
7. If you chose to create an online copy, set the reserved capacity needed to store data and other information for the online copy, and then click **Finish** to start the online copy.

The volume candidate table displays only the candidates that support the reserved capacity specified. Reserved capacity is the physical allocated capacity that is used for any copy service operation and storage object. It is not directly readable by the host.

Allocate the reserved capacity using the following guidelines:

- The default setting for reserved capacity is 40% of the capacity of the base volume, and usually this capacity is sufficient.
- Reserved capacity, however, varies depending on the number of changes to the original data. The longer a storage object is active, the larger the reserved capacity should be.

Results

System Manager copies all data from the source volume to the target volume. After the Copy Volume operation is complete, the target volume automatically becomes read-only to the hosts.

After you finish

Select **Home** > **View Operations in Progress** to view the progress of the Copy Volume operation. This operation can be lengthy and could affect system performance.

Take action on a Copy Volume operation

You can view a Copy Volume operation in progress and stop, change priority, re-copy, or clear a Copy Volume operation.

Steps


1. Select **Home** > **View Operations in Progress**.

The **Operations in Progress** dialog box appears.

2. Find the Copy Volume operation that you want to take action on, and then click the link in the **Actions** column to take one of the following actions.

Read all cautionary text provided in dialogs, particularly when stopping an operation.

Action	Description
Stop	<p>You can stop a Copy Volume operation while the operation has a status of In Progress, Pending, or Failed.</p> <p>When the Copy Volume is stopped, all of the mapped hosts have write access to the source volume. If data is written to the source volume, the data on the target volume no longer matches the data on the source volume.</p>
Change priority	<p>You can change the priority of a Copy Volume operation while the operation has a status of In Progress to select the rate at which a Copy Volume operation completes.</p>
Re-copy	<p>You can re-copy a volume when you have stopped a Copy Volume operation and want to start it again or when a Copy Volume operation has failed or halted. The Copy Volume operation starts over from the beginning.</p> <p>The re-copy action overwrites existing data on the target volume and fails all snapshot volumes associated with the target volume, if any exist.</p>

Action	Description
Clear	<p>You can remove the Copy Volume operation while the operation has a status of In Progress, Pending, or Failed.</p> <div>  <p>Be sure that you want to do this operation before selecting Clear. There is no confirmation dialog.</p> </div>

Change allocated capacity limit for a thin volume

For thin volumes capable of allocating space on demand, you can change the limit that restricts the allocated capacity to which a thin volume can automatically expand. You also can change the percentage point at which an alert (warning threshold exceeded) is sent to the Notifications area on the Home page when a thin volume is near the allocated capacity limit. You can choose to enable or disable this alert notification.



SANtricity System Manager does not provide an option to create thin volumes. If you want to create thin volumes, use the Command Line Interface (CLI).



This feature is not available on the EF600 storage system.

About this task

System Manager automatically expands the allocated capacity based on the allocated capacity limit. The allocated capacity limit allows you to limit the thin volume's automatic growth below the reported capacity. When the amount of data written gets close to the allocated capacity, you can change the allocated capacity limit.

When changing a thin volume's allocated capacity limit and warning threshold, you must take into account the space to be consumed by both the volume's user data and copy services data.

Steps

1. Select **Storage > Volumes**.

2. Select the **Thin Volume Monitoring** tab.

The Thin Volume Monitoring view appears.

3. Select the thin volume that you want to change, and then select **Change Limit**.

The **Change Limit** dialog box appears. The allocated capacity limit and warning threshold setting for the thin volume you selected appear in this dialog box.

4. Change the allocated capacity limit and warning threshold as needed.

Field Details

Setting	Description
Change allocated capacity limit to...	The threshold at which writes fail, preventing the thin volume from consuming additional resources. This threshold is a percentage of the volume's reported capacity size.
Alert me when... (warning threshold)	Select the check box if you want the system to generate an alert when a thin volume is near the allocated capacity limit. The alert is sent to the Notifications area on the Home page. This threshold is a percentage of the volume's reported capacity size. Clear the check box to disable the warning threshold alert notification.

5. Click **Save**.

FAQs

What is a volume?

A volume is a container in which applications, databases, and file systems store data. It is the logical component created for the host to access storage on the storage array.

A volume is created from the capacity available in a pool or a volume group. A volume has a defined capacity. Although a volume might consist of more than one drive, a volume appears as one logical component to the host.

Why am I seeing a capacity over-allocation error when I have enough free capacity in a volume group to create volumes?

The selected volume group might have one or more free capacity areas. A free capacity area is the free capacity that can result from deleting a volume or from not using all available free capacity during volume creation.

When you create a volume in a volume group that has one or more free capacity areas, the volume's capacity is limited to the largest free capacity area in that volume group. For example, if a volume group has a total of 15 GiB free capacity, and the largest free capacity area is 10 GiB, the largest volume you can create is 10 GiB.

If a volume group has free capacity areas, the volume group graph contains a link indicating the number of existing free capacity areas. Select the link to display a pop-over that indicates the capacity of each area.

By consolidating free capacity, you can create additional volumes from the maximum amount of free capacity in a volume group. You can consolidate the existing free capacity on a selected volume group using one of the following methods:

- When at least one free capacity area is detected for a volume group, the **Consolidate free capacity** recommendation appears on the **Home** page in the Notification area. Click the **Consolidate free capacity** link to launch the dialog box.
- You can also select **Pools & Volume Groups > Uncommon Tasks > Consolidate volume group free capacity** to launch the dialog box.

If you want to use a specific free capacity area rather than the largest free capacity area, use the Command Line Interface (CLI).

How does my selected workload impact volume creation?

A workload is a storage object that supports an application. You can define one or more workloads, or instances, per application. For some applications, System Manager configures the workload to contain volumes with similar underlying volume characteristics. These volume characteristics are optimized based on the type of application the workload supports. For example, if you create a workload that supports a Microsoft SQL Server application and then subsequently create volumes for that workload, the underlying volume characteristics are optimized to support Microsoft SQL Server.

- **Application-specific.** When you are creating volumes using an application-specific workload, the system may recommend an optimized volume configuration to minimize contention between application workload I/O and other traffic from your application instance. Volume characteristics like I/O type, segment size, controller ownership, and read and write cache are automatically recommended and optimized for workloads that are created for the following application types.
 - Microsoft® SQL Server™
 - Microsoft® Exchange Server™
 - Video surveillance applications
 - VMware ESXi™ (for volumes to be used with Virtual Machine File System) You can review the recommended volume configuration and edit, add, or delete the system-recommended volumes and characteristics using the **Add/Edit Volumes** dialog box.
- **Other** (or applications without specific volume creation support). Other workloads use a volume configuration that you must manually specify when you want to create a workload that is not associated with a specific application, or if System Manager does not have built-in optimization for the application you intend to use on the storage array. You must manually specify the volume configuration using the **Add/Edit Volumes** dialog box.

Why aren't these volumes associated with a workload?

Volumes are not associated with a workload if they have been created using the command line interface (CLI) or if they have been migrated (imported/exported) from a different storage array.

Why can't I delete the selected workload?

This workload consists of a group of volumes that were created using the command line interface (CLI) or migrated (imported/exported) from a different storage array. As a result, the volumes in this workload are not associated with an application-specific workload, so the workload cannot be deleted.

How do application-specific workloads help me manage my storage array?

An application is software such as SQL Server or Exchange. You define one or more workloads to support each application. For some applications, System Manager will

automatically recommend a volume configuration that optimizes storage. Characteristics such as I/O type, segment size, controller ownership, and read and write cache are included in the volume configuration.

The volume characteristics of your application-specific workload dictate how the workload interacts with the components of your storage array and helps determine the performance of your environment under a given configuration.

How does providing this information help create storage?

The workload information is used to optimize the volume characteristics such as I/O type, segment size, and read/write cache for the workload selected. These optimized characteristics dictate how your workload interacts with the storage array components.

Based on the workload information you provide, System Manager creates the appropriate volumes and places them on the available pools or volume groups that currently exist on the system. The system creates the volumes and optimizes their characteristics based on the current best practices for the workload you selected.

Before you finish creating volumes for a given workload, you can review the recommended volume configuration and edit, add, or delete the system-recommended volumes and characteristics using the **Add/Edit Volumes** dialog box.

Refer to your application-specific documentation for best practice information.

What do I need to do to recognize the expanded capacity?

If you increase the capacity for a volume, the host might not immediately recognize the increase in volume capacity.

Most operating systems recognize the expanded volume capacity and automatically expand after the volume expansion is initiated. However, some might not. If your OS does not automatically recognize the expanded volume capacity, you might need to perform a disk rescan or reboot.

After you have expanded the volume capacity, you must manually increase the file system size to match. How you do this depends on the file system you are using.

Refer to your host operating system documentation for additional details.

Why don't I see all my pools and/or volume groups?

Any pool or volume group to which you cannot move the volume does not display in the list.

Pools or volume groups are not eligible for any of the following reasons:

- The Data Assurance (DA) capabilities of a pool or volume group pool do not match.
- A pool or volume group is in a non-optimal state.
- The capacity of a pool or volume group is too small.

What is segment size?

A segment is the amount of data in kilobytes (KiB) that is stored on a drive before the

storage array moves to the next drive in the stripe (RAID group). Segment size applies only to volume groups, not pools.

Segment size is defined by the number of data blocks it contains. When determining segment size, you must know what type of data you will store in a volume. If an application typically uses small, random reads and writes (IOPS), a smaller segment size typically works better. Alternatively, if the application has large, sequential reads and writes (throughput), a large segment size is generally better.

Whether an application uses small random reads and writes, or large sequential reads and writes, the storage array performs better if the segment size is larger than the typical data block chunk size. This normally makes it easier and faster for the drives to access the data, which is important for better storage array performance.

In an environment where IOPS performance is important

In an I/O operations per second (IOPS) environment, the storage array performs better if you use a segment size that is larger than the typical data block size ("chunk") that is read/written to a drive. This ensures that each chunk is written to a single drive.

In an environment where throughput is important

In a throughput environment, the segment size should be an even fraction of the total drives for data and the typical data chunk size (I/O size). This spreads the data as a single stripe across the drives in the volume group leading to faster reads and writes.

What is preferred controller ownership?

Preferred controller ownership defines the controller that is designated to be the owning, or primary, controller of the volume.

Controller ownership is very important and should be planned carefully. Controllers should be balanced as closely as possible for total I/Os.

For example, if one controller reads primarily large, sequential data blocks and the other controller has small data blocks with frequent reads and writes, the loads are very different. Knowing which volumes contain what type of data allows you to balance I/O transfers equally over both controllers.

What is Automatic Load Balancing?

The Automatic Load Balancing feature provides automated I/O balancing and ensures that incoming I/O traffic from the hosts is dynamically managed and balanced across both controllers.

The Automatic Load Balancing feature provides improved I/O resource management by reacting dynamically to load changes over time and automatically adjusting volume controller ownership to correct any load imbalance issues when workloads shift across the controllers.

The workload of each controller is continually monitored and, with cooperation from the multipath drivers installed on the hosts, can be automatically brought into balance whenever necessary. When workload is automatically re-balanced across the controllers, the storage administrator is relieved of the burden of manually adjusting volume controller ownership to accommodate load changes on the storage array.

When Automatic Load Balancing is enabled, it performs the following functions:

- Automatically monitors and balances controller resource utilization.
- Automatically adjusts volume controller ownership when needed, thereby optimizing I/O bandwidth between the hosts and the storage array.



Any volume assigned to use a controller's SSD Cache is not eligible for an automatic load balance transfer.

Hosts

Concepts

Host terminology


Learn how the host terms apply to your storage array.

Component	Definition
Host	A host is a server that sends I/O to a volume on a storage array.
Host name	The host name should equate to the system name of the host.
Host cluster	A host cluster is a group of hosts. You create a host cluster to make it easy to assign the same volumes to multiple hosts.
Host interface protocol	A host interface protocol is the connection (such as Fibre Channel, iSCSI, etc.) between the controllers and the hosts.
HBA or Network Interface Card (NIC)	A host bus adapter (HBA) is a board that resides in a host and contains one or more host ports.
Host port	A host port is a port on a host bus adapter (HBA) that provides the physical connection to a controller and is used for I/O operations.

Component	Definition
Host port identifier	<p>A host port identifier is a unique world-wide name associated with each host port on a host bus adapter (HBA).</p> <ul style="list-style-type: none"> • Internet Small Computer System Interface (iSCSI) host port identifiers must have between 1 and 233 characters. iSCSI host port identifiers display in standard IQN format (e.g., <code>iqn.xxx.com.xxx:8b3ad</code>). • Non-iSCSI host port identifiers such as Fibre Channel and Serial Attached SCSI (SAS) display as colon-delimited after every two characters (e.g., <code>xx:yy:zz</code>). Fibre Channel host port identifiers must have 16 characters.
Host operating system type	<p>The host operating system type is a configuration setting that defines how the controllers in the storage array react to I/O depending on the operating system (or variant) of the host. This is also sometimes called <i>host type</i> for short.</p>
Controller host port	<p>A controller host port is a port on the controller that provides the physical connection to a host and is used for I/O operations.</p>
LUN	<p>A logical unit number (LUN) is the number assigned to the address space that a host uses to access a volume. The volume is presented to the host as capacity in the form of a LUN.</p> <p>Each host has its own LUN address space. Therefore, the same LUN can be used by different hosts to access different volumes.</p>

Workflow for host creation and volume assignment

The following figure illustrates how to configure host access.



Automatic versus manual host creation

Creating a host is one of the steps required to let the storage array know which hosts are attached to it and to allow I/O access to the volumes. You can create a host automatically or manually.

Automatic creation

Automatic host creation for SCSI-based (not NVMe-oF) hosts is initiated by the Host Context Agent (HCA). The HCA is a utility that you can install on each host attached to the storage array. Each host that has the HCA installed pushes its configuration information to the storage array controllers through the I/O path. Based on the host information, the controllers automatically create the host and the associated host ports and set the host type. If needed, you can make any additional changes to the host configuration using System Manager.

After the HCA performs its automatic detection, the host automatically appears in the Hosts page with the

following attributes:

- The host name derived from the system name of the host.
- The host identifier ports that are associated with the host.
- The Host Operating System Type of the host.

Hosts are created as stand-alone hosts; the HCA does not automatically create or add to host clusters.

Manual creation

You might want to manually create a host for one of the following reasons:

1. You chose not to install the HCA utility on your hosts.
2. You want to ensure that the host port identifiers that were detected by the storage array controllers are associated correctly with the hosts.

During manual host creation, you associate host port identifiers by selecting them from a list or manually entering them. After you create a host, you can assign volumes to it or add it to a host cluster if you plan to share access to volumes.

How volumes are assigned to hosts and host clusters

For a host or host cluster to send I/O to a volume, you must assign the volume to the host or host cluster.

You can select a host or host cluster when you create a volume or you can assign a volume to a host or host cluster later. A host cluster is a group of hosts. You create a host cluster to make it easy to assign the same volumes to multiple hosts.

Assigning volumes to hosts is flexible, allowing you to meet your particular storage needs.

- **Stand-alone host, not part of a host cluster** — You can assign a volume to an individual host. The volume can be accessed only by the one host.
- **Host cluster** — You can assign a volume to a host cluster. The volume can be accessed by all the hosts in the host cluster.
- **Host within a host cluster** — You can assign a volume to an individual host that is part of a host cluster. Even though the host is part of a host cluster, the volume can be accessed only by the individual host and not by any other hosts in the host cluster.

When volumes are created, logical unit numbers (LUNs) are assigned automatically. The LUN serves as the "address" between the host and the controller during I/O operations. You can change LUNs after the volume is created.

Access volumes

An access volume is a factory-configured volume on the storage array that is used for communication with the storage array and the host through the host I/O connection. The access volume requires a Logical Unit Number (LUN).

The access volume is used in two instances:

- **Automatic host creation** — The access volume is used by the Host Context Agent (HCA) utility to push

host information (name, ports, host type) to System Manager for automatic host creation.

- **In-band management** — The access volume is used for an in-band connection to manage the storage array. This can only be done if you are managing the storage array with the command line interface (CLI).



In-band management is not available for EF600 storage systems.

An access volume is automatically created the first time you assign a volume to a host. For example, if you assign Volume_1 and Volume_2 to a host, when you view results of that assignment, you see three volumes (Volume_1, Volume_2, and Access).

If you are not automatically creating hosts or managing a storage array in-band with the CLI, you do not need the access volume, and you can free up the LUN by deleting the access volume. This action removes the volume-to-LUN assignment as well as any in-band management connections to the host.

Maximum number of LUNs

The storage array has a maximum number of logical unit numbers (LUNs) that can be used for each host.

The maximum number depends on the operating system of the host. The storage array tracks the number of LUNs used. If you try to assign a volume to a host that exceeds the maximum number of LUNs, the host cannot access the volume.

How to

Configure host access

Create host automatically

You can allow the Host Context Agent (HCA) to automatically detect the hosts, and then verify that the information is correct. Creating a host is one of the steps required to let the storage array know which hosts are attached to it and to allow I/O access to the volumes.

Before you begin

The Host Context Agent (HCA) is installed and running on every host connected to the storage array. Hosts with the HCA installed and connected to the storage array are created automatically. To install the HCA, install SANtricity Storage Manager on the host and select the Host option. The HCA is not available on all supported operating systems. If it is not available, you must create the host manually.

Steps

1. Select **Storage > Hosts**.

The table lists the automatically-created hosts.

2. Verify that the information provided by the HCA is correct (name, host type, host port identifiers).

If you need to change any of the information, select the host, and then click **View/Edit Settings**.

3. **Optional:** If you want the automatically-created host to be in a cluster, create a host cluster and add the host or hosts.

Results

After a host is created automatically, the system displays the following items in the Hosts tile table:

- The host name derived from the system name of the host.
- The host identifier ports that are associated with the host.
- The Host Operating System Type of the host.

Create host manually

For hosts that cannot be automatically discovered, you can manually create a host. Creating a host is one of the steps required to let the storage array know which hosts are attached to it and to allow I/O access to the volumes.

About this task

Keep these guidelines in mind when you create a host:

- You must define the host identifier ports that are associated with the host.
- Make sure that you provide the same name as the host's assigned system name.
- This operation does not succeed if the name you choose is already in use.
- The length of the name cannot exceed 30 characters.

Steps

1. Select **Storage › Hosts**.
2. Click **Create › Host**.

The **Create Host** dialog box appears.

3. Select the settings for the host as appropriate.

Field details

Setting	Description
Name	Type a name for the new host.
Host operating system type	Select the operating system that is running on the new host from the drop-down list.
Host interface type	Optional: If you have more than one type of host interface supported on your storage array, select the host interface type that you want to use.
Host ports	<p>Do one of the following:</p> <ul style="list-style-type: none">• Select I/O Interface <p>Generally, the host ports should have logged in and be available from the drop-down list. You can select the host port identifiers from the list.</p> <ul style="list-style-type: none">• Manual add <p>If a host port identifier is not displayed in the list, it means that the host port has not logged in. An HBA utility or the iSCSI initiator utility may be used to find the host port identifiers and associate them with the host.</p> <p>You can manually enter the host port identifiers or copy/paste them from the utility (one at a time) into the Host ports field.</p> <p>You must select one host port identifier at a time to associate it with the host, but you can continue to select as many identifiers that are associated with the host. Each identifier is displayed in the Host ports field. If necessary, you also can remove an identifier by selecting the X next to it.</p>

Setting	Description
CHAP initiator	<p>Optional: If you selected or manually entered a host port with an iSCSI IQN, and if you want to require a host that tries to access the storage array to authenticate using Challenge Handshake Authentication Protocol (CHAP), select the CHAP initiator checkbox. For each iSCSI host port you selected or manually entered, do the following:</p> <ul style="list-style-type: none"> • Enter the same CHAP secret that was set on each iSCSI host initiator for CHAP authentication. If you are using mutual CHAP authentication (two-way authentication that enables a host to validate itself to the storage array and for a storage array to validate itself to the host), you also must set the CHAP secret for the storage array at initial setup or by changing settings. • Leave the field blank if you do not require host authentication. Currently, the only iSCSI authentication method used by System Manager is CHAP.

4. Click **Create**.

Results

After the host is successfully created, the system creates a default name for each host port configured for the host (user label).

The default alias is <Hostname_Port Number>. For example, the default alias for the first port created for host IPT is IPT_1.

Create host cluster

You create a host cluster when two or more hosts require I/O access to the same volumes.

About this task

Keep these guidelines in mind when you create a host cluster:

- This operation does not start unless there are two or more hosts available to create the cluster.
- Hosts in host clusters can have different operating systems (heterogeneous).
- To create a Data Assurance (DA)-enabled volume, the host connection you are planning to use must support DA.

If any of the host connections on the controllers in your storage array do not support DA, the associated hosts cannot access data on DA-enabled volumes.

DA is **not** supported by iSCSI over TCP/IP, or by the SRP over InfiniBand.

- This operation does not succeed if the name you choose is already in use.
- The length of the name cannot exceed 30 characters.

Steps

1. Select **Storage › Hosts**.
2. Select **Create › Host Cluster**.

The **Create Host Cluster** dialog box appears.

3. Select the settings for the host cluster as appropriate.

Field details

Setting	Description
Name	Type the name for the new host cluster.
Hosts	Select two or more hosts from the drop-down list. Only those hosts that are not already part of a host cluster appear in the list.

4. Click **Create**.

If the selected hosts are attached to interface types that have different Data Assurance (DA) capabilities, a dialog box appears with the message that DA will be unavailable on the host cluster. This unavailability prevents DA-enabled volumes from being added to the host cluster. Select **Yes** to continue or **No** to cancel.

DA increases data integrity across the entire storage system. DA enables the storage array to check for errors that might occur when data is moved between the hosts and the drives. Using DA for the new volume ensures that any errors are detected.

Results

The new host cluster appears in the table with the assigned hosts in the rows beneath.

Create volumes

You create volumes to add storage capacity to an application-specific workload, and to make the created volumes visible to a specific host or host cluster. In addition, the volume creation sequence provides options to allocate specific amounts of capacity to each volume you want to create.

About this task

Most application types default to a user-defined volume configuration. Some application types have a smart configuration applied at volume creation. For example, if you are creating volumes for Microsoft Exchange application, you are asked how many mailboxes you need, what your average mailbox capacity requirements are, and how many copies of the database you want. System Manager uses this information to create an optimal volume configuration for you, which can be edited as needed.



If you want to mirror a volume, first create the volumes that you want to mirror, and then use the **Storage › Volumes › Copy Services › Mirror a volume asynchronously** option.

The process to create a volume is a multi-step procedure.

Step 1: Select host for a volume

You create volumes to add storage capacity to an application-specific workload, and to make the created volumes visible to a specific host or host cluster. In addition, the volume creation sequence provides options to allocate specific amounts of capacity to each volume you want to create.

Before you begin

- Valid hosts or host clusters exist under the **Hosts** tile.
- Host port identifiers have been defined for the host.
- Before creating a DA-enabled volume, the host connection you are planning to use must support DA. If any of the host connections on the controllers in your storage array do not support DA, the associated hosts cannot access data on DA-enabled volumes.

About this task

Keep these guidelines in mind when you assign volumes:

- A host's operating system can have specific limits on how many volumes the host can access. Keep this limitation in mind when you create volumes for use by a particular host.
- You can define one assignment for each volume in the storage array.
- Assigned volumes are shared between controllers in the storage array.
- The same logical unit number (LUN) cannot be used twice by a host or a host cluster to access a volume. You must use a unique LUN.



Assigning a volume to a host will fail if you try to assign a volume to a host cluster that conflicts with an established assignment for a host in the host clusters.

Steps

1. Select **Storage › Volumes**.
2. Select **Create › Volume**.

The **Create Volumes** dialog box appears.

3. From the drop-down list, select a specific host or host cluster to which you want to assign volumes, or choose to assign the host or host cluster at a later time.
4. To continue the volume creation sequence for the selected host or host cluster, click **Next**, and go to [Step 2: Select a workload for a volume](#).

The **Select Workload** dialog box appears.

Step 2: Select a workload for a volume

Select a workload to customize the storage array configuration for a specific application, such as Microsoft SQL Server, Microsoft Exchange, Video Surveillance applications, or VMware. You can select "Other application" if the application you intend to use on this storage array is not listed.

About this task

This task describes how to create volumes for an existing workload.

- *When you are creating volumes using an application-specific workload*, the system may recommend an optimized volume configuration to minimize contention between application workload I/O and other traffic from your application instance. You can review the recommended volume configuration and edit, add, or delete the system-recommended volumes and characteristics using the **Add/Edit Volumes** dialog box.
- *When you are creating volumes using "Other" applications* (or applications without specific volume creation support), you manually specify the volume configuration using the **Add/Edit Volumes** dialog box.

Steps

1. Do one of the following:
 - Select the **Create volumes for an existing workload** option to create volumes for an existing workload.
 - Select the **Create a new workload** option to define a new workload for a supported application or for "Other" applications.
 - From the drop-down list, select the name of the application you want to create the new workload for.

Select one of the "Other" entries if the application you intend to use on this storage array is not listed.
 - Enter a name for the workload you want to create.
2. Click **Next**.
3. If your workload is associated with a supported application type, enter the information requested; otherwise, go to [Step 3: Add or edit volumes](#).

Step 3: Add or edit volumes

Before you begin

- The pools or volume groups must have sufficient free capacity.
- To create a Data Assurance (DA)-enabled volume, the host connection you are planning to use must support DA.

Selecting a DA capable pool or volume group

If you want to create a DA-enabled volume, select a pool or volume group that is DA capable (look for **Yes** next to "DA" in the pool and volume group candidates table).

DA capabilities are presented at the pool and volume group level in System Manager. DA protection checks for and corrects errors that might occur as data is transferred through the controllers down to the drives. Selecting a DA-capable pool or volume group for the new volume ensures that any errors are detected and corrected.

If any of the host connections on the controllers in your storage array do not support DA, the associated hosts cannot access data on DA-enabled volumes.



DA is not supported by iSCSI over TCP/IP, or by the SRP over InfiniBand.

- To create a secure-enabled volume, a security key must be created for the storage array.

Selecting a secure-capable pool or volume group

If you want to create a secure-enabled volume, select a pool or volume group that is secure capable (look for **Yes** next to "Secure-capable" in the pool and volume group candidates table).

Drive security capabilities are presented at the pool and volume group level in System Manager. Secure-capable drives prevent unauthorized access to the data on a drive that is physically removed from the storage array. A secure-enabled drive encrypts data during writes and decrypts data during reads using a unique *encryption key*.

A pool or volume group can contain both secure-capable and non-secure-capable drives, but all drives must be secure-capable to use their encryption capabilities.

About this task

You create volumes from pools or volume groups. The **Add/Edit Volumes** dialog box shows all eligible pools and volume groups on the storage array. For each eligible pool and volume group, the number of drives available and the total free capacity appears.

For some application-specific workloads, each eligible pool or volume group shows the proposed capacity based on the suggested volume configuration and shows the remaining free capacity in GiB. For other workloads, the proposed capacity appears as you add volumes to a pool or volume group and specify the reported capacity.

Steps

1. Choose one of these actions based on whether you selected Other or an application-specific workload:
 - **Other** — Click **Add new volume** in each pool or volume group that you want to use to create one or more volumes.

Field Details

Field	Description
Volume Name	<p>A volume is assigned a default name by System Manager during the volume creation sequence. You can either accept the default name or provide a more descriptive one indicating the type of data stored in the volume.</p>
Reported Capacity	<p>Define the capacity of the new volume and the capacity units to use (MiB, GiB, or TiB). For Thick volumes, the minimum capacity is 1 MiB, and the maximum capacity is determined by the number and capacity of the drives in the pool or volume group.</p> <p>Keep in mind that storage capacity is also required for copy services (snapshot images, snapshot volumes, volume copies, and remote mirrors); therefore, do not allocate all of the capacity to standard volumes.</p> <p>Capacity in a pool is allocated in 4-GiB increments. Any capacity that is not a multiple of 4 GiB is allocated but not usable. To make sure that the entire capacity is usable, specify the capacity in 4-GiB increments. If unusable capacity exists, the only way to regain it is to increase the capacity of the volume.</p>

Field	Description
Segment Size	<p data-bbox="867 155 1456 289">Shows the setting for segment sizing, which only appears for volumes in a volume group. You can change the segment size to optimize performance.</p> <p data-bbox="867 323 1456 667">Allowed segment size transitions — System Manager determines the segment size transitions that are allowed. Segment sizes that are inappropriate transitions from the current segment size are unavailable on the drop-down list. Allowed transitions usually are double or half of the current segment size. For example, if the current volume segment size is 32 KiB, a new volume segment size of either 16 KiB or 64 KiB is allowed.</p> <p data-bbox="867 701 1456 1045">SSD Cache-enabled volumes — You can specify a 4-KiB segment size for SSD Cache-enabled volumes. Make sure you select the 4-KiB segment size only for SSD Cache-enabled volumes that handle small-block I/O operations (for example, 16 KiB I/O block sizes or smaller). Performance might be impacted if you select 4 KiB as the segment size for SSD Cache-enabled volumes that handle large block sequential operations.</p> <p data-bbox="867 1079 1456 1213">Amount of time to change segment size — The amount of time to change a volume's segment size depends on these variables:</p> <ul data-bbox="889 1247 1456 1516" style="list-style-type: none"> • The I/O load from the host • The modification priority of the volume • The number of drives in the volume group • The number of drive channels • The processing power of the storage array controllers <p data-bbox="867 1549 1456 1650">When you change the segment size for a volume, I/O performance is affected, but your data remains available.</p>

Field	Description
Secure-capable	<p>Yes appears next to "Secure-capable" only if the drives in the pool or volume group are secure-capable.</p> <p>Drive Security prevents unauthorized access to the data on a drive that is physically removed from the storage array. This option is available only when the Drive Security feature has been enabled, and a security key is set up for the storage array.</p> <p>A pool or volume group can contain both secure-capable and non-secure-capable drives, but all drives must be secure-capable to use their encryption capabilities.</p>
DA	<p>Yes appears next to "DA" only if the drives in the pool or volume group support Data Assurance (DA).</p> <p>DA increases data integrity across the entire storage system. DA enables the storage array to check for errors that might occur as data is transferred through the controllers down to the drives. Using DA for the new volume ensures that any errors are detected.</p>

- **Application-specific workload** — Either click **Next** to accept the system-recommended volumes and characteristics for the selected workload, or click **Edit Volumes** to change, add, or delete the system-recommended volumes and characteristics for the selected workload.

Field Details

Field	Description
Volume Name	A volume is assigned a default name by System Manager during the volume creation sequence. You can either accept the default name or provide a more descriptive one indicating the type of data stored in the volume.
Reported Capacity	<p>Define the capacity of the new volume and the capacity units to use (MiB, GiB, or TiB). For Thick volumes, the minimum capacity is 1 MiB, and the maximum capacity is determined by the number and capacity of the drives in the pool or volume group.</p> <p>Keep in mind that storage capacity is also required for copy services (snapshot images, snapshot volumes, volume copies, and remote mirrors); therefore, do not allocate all of the capacity to standard volumes.</p> <p>Capacity in a pool is allocated in 4-GiB increments. Any capacity that is not a multiple of 4 GiB is allocated but not usable. To make sure that the entire capacity is usable, specify the capacity in 4-GiB increments. If unusable capacity exists, the only way to regain it is to increase the capacity of the volume.</p>
Volume Type	Volume type indicates the type of volume that was created for an application-specific workload.

Field	Description
Segment Size	<p>Shows the setting for segment sizing, which only appears for volumes in a volume group. You can change the segment size to optimize performance.</p> <p>Allowed segment size transitions — System Manager determines the segment size transitions that are allowed. Segment sizes that are inappropriate transitions from the current segment size are unavailable on the drop-down list. Allowed transitions usually are double or half of the current segment size. For example, if the current volume segment size is 32 KiB, a new volume segment size of either 16 KiB or 64 KiB is allowed.</p> <p>SSD Cache-enabled volumes — You can specify a 4-KiB segment size for SSD Cache-enabled volumes. Make sure you select the 4-KiB segment size only for SSD Cache-enabled volumes that handle small-block I/O operations (for example, 16 KiB I/O block sizes or smaller). Performance might be impacted if you select 4 KiB as the segment size for SSD Cache-enabled volumes that handle large block sequential operations.</p> <p>Amount of time to change segment size — The amount of time to change a volume's segment size depends on these variables:</p> <ul style="list-style-type: none"> • The I/O load from the host • The modification priority of the volume • The number of drives in the volume group • The number of drive channels • The processing power of the storage array controllers <p>When you change the segment size for a volume, I/O performance is affected, but your data remains available.</p>
Secure-capable	<p>Yes appears next to "Secure-capable" only if the drives in the pool or volume group are secure-capable.</p> <p>Drive security prevents unauthorized access to the data on a drive that is physically removed from the storage array. This option is available only when the drive security feature has been enabled, and a security key is set up for the storage array.</p> <p>A pool or volume group can contain both secure-capable and non-secure-capable drives, but all drives must be secure-capable to use their encryption capabilities.</p>

Field	Description
DA	<p>Yes appears next to "DA" only if the drives in the pool or volume group support Data Assurance (DA).</p> <p>DA increases data integrity across the entire storage system. DA enables the storage array to check for errors that might occur as data is transferred through the controllers down to the drives. Using DA for the new volume ensures that any errors are detected.</p>

2. To continue the volume creation sequence for the selected application, click **Next**, and go to [Step 4: Review volume configuration](#).

Step 4: Review volume configuration

Review a summary of the volumes you intend to create and make any necessary changes.

Steps

1. Review the volumes you want to create. Click **Back** to make any changes.
2. When you are satisfied with your volume configuration, click **Finish**.

Results

System Manager creates the new volumes in the selected pools and volume groups, and then displays the new volumes in the All Volumes table.

After you finish

- Perform any operating system modifications necessary on the application host so that the applications can use the volume.
- Run either the host-based `hot_add` utility or an operating system-specific utility (available from a third-party vendor), and then run the `SMdevices` utility to correlate volume names with host storage array names.

The `hot_add` utility and the `SMdevices` utility are included as part of the `SMutils` package. The `SMutils` package is a collection of utilities to verify what the host sees from the storage array. It is included as part of the SANtricity software installation.

Assign volumes

You must assign a volume to a host or a host cluster so it can be used for I/O operations. This assignment grants a host or host cluster access to one or more volumes in a storage array.

Before you begin

Keep these guidelines in mind when you assign volumes:

- You can assign a volume to only one host or host cluster at a time.
- Assigned volumes are shared between controllers in the storage array.
- The same logical unit number (LUN) cannot be used twice by a host or a host cluster to access a volume. You must use a unique LUN.

Assigning a volume fails under these conditions:

- All volumes are assigned.
- The volume is already assigned to another host or host cluster.

The ability to assign a volume is unavailable under these conditions:

- No valid hosts or host clusters exist.
- No host port identifiers have been defined for the host.
- All volume assignments have been defined.

About this task

All unassigned volumes are displayed, but functions for hosts with or without Data Assurance (DA) apply as follows:

- For a DA-capable host, you can select volumes that are either DA-enabled or not DA-enabled.
- For a host that is not DA-capable, if you select a volume that is DA-enabled, a warning states that the system must automatically turn off DA on the volume before assigning the volume to the host.

Steps

1. Select **Storage > Hosts**.
2. Select the host or host cluster to which you want to assign volumes, and then click **Assign Volumes**.

A dialog box appears that lists all the volumes that can be assigned. You can sort any of the columns or type something in the **Filter** box to make it easier to find particular volumes.

3. Select the check box next to each volume that you want to assign or select the check box in the table header to select all volumes.
4. Click **Assign** to complete the operation.

Results

After successfully assigning a volume or volumes to a host or a host cluster, the system performs the following actions:

- The assigned volume receives the next available LUN number. The host uses the LUN number to access the volume.
- The user-supplied volume name appears in volume listings associated to the host. If applicable, the factory-configured access volume also appears in volume listings associated to the host.

Manage hosts and host clusters

Change the settings for a host

You can change the name, host operating system type, and associated host clusters for a host.

Steps

1. Select **Storage > Hosts**.
2. Select the host that you want to edit, and then click **View/Edit Settings**.

A dialog box appears that shows the current host settings.

3. If it is not already selected, click the **Properties** tab.
4. Change the settings as appropriate.

Field Details

Setting	Description
Name	You can change the user-supplied name of the host. Specifying a name for the host is required.
Associated host cluster	<p>You can choose one of the following options:</p> <ul style="list-style-type: none">• None — The host remains a standalone host. If the host was associated to a host cluster, the system removes the host from the cluster.• <Host Cluster> — The system associates the host to the selected cluster.
Host operating system type	You can change the type of operating system running on the host you defined.

5. Click **Save**.

Change the settings for a host cluster

You can change the host cluster name, or add or remove hosts in a host cluster.

Steps

1. Select **Storage > Hosts**.
2. Select the host cluster you want to edit, and then click **View/Edit Settings**.

A dialog box appears that shows the current host cluster settings.

3. Change the settings for the host cluster as appropriate.

Field Details

Setting	Description
Name	You can specify the user-supplied name of the host cluster. Specifying a name for a cluster is required.
Associated Hosts	<p>To add a host, click the Associated Hosts box, and then select a host name from the drop-down list. You cannot manually enter a host name.</p> <p>To delete a host, click the X next to the host name.</p>

4. Click **Save**.

Unassign volumes

Unassign volumes from hosts or host clusters if you no longer need I/O access to that volume from the host or host cluster.

About this task

Keep these guidelines in mind when you unassign a volume:

- If you are removing the last assigned volume from a host cluster, and the host cluster also has hosts with specific assigned volumes, make sure that you remove or move those assignments before removing the last assignment for the host cluster.
- If a host cluster, a host, or a host port is assigned to a volume that is registered to the operating system, you must clear this registration before you can remove these nodes.

Steps

1. Select **Storage > Hosts**.
2. Select the host or host cluster that you want to edit, and then click **Unassign Volumes**.

A dialog box appears that shows all the volumes that are currently assigned.

3. Select the check box next to each volume that you want to unassign or select the check box in the table header to select all volumes.
4. Click **Unassign**.

Results

- The volumes that were unassigned are available for a new assignment.
- Until the changes are configured on the host, the volume is still recognized by the host operating system.

Change host port identifiers for a host

Change the host port identifiers when you want to change the user label on a host port identifier, add a new host port identifier to the host, or delete a host port identifier from the host.

About this task

When changing host port identifiers, keep the following guidelines in mind:

- **Add** — When you add a host port, you are associating the host port identifier to the host you created to connect to your storage array. You can manually enter port information using a host bus adapter (HBA) utility.
- **Edit** — You can edit the host ports to move (associate) a host port to a different host. You might have moved the host bus adapter or iSCSI initiator to a different host, so you must move (associate) the host port to the new host.
- **Delete** — You can delete host ports to remove (unassociate) host ports from a host.

Steps

1. Select **Storage > Hosts**.

2. Select the host to which the ports will be associated, and then click **View/Edit Settings**.


If you want to add ports to a host in a host cluster, expand the host cluster and select the desired host. You cannot add ports at the host cluster level.

A dialog box appears that shows the current host settings.

3. Click the **Host Ports** tab.

The dialog box shows the current host port identifiers.

4. Change the host port identifier settings as appropriate.

Setting	Description
Host Port	<p>You can choose one of the following options:</p> <ul style="list-style-type: none">• Add — Use Add to associate a new host port identifier to the host. The length of the host port identifier name is determined by the host interface technology.<ul style="list-style-type: none">◦ Fibre Channel host port identifier names must have 16 characters.◦ Infiniband host port identifier names must have 16 characters.◦ iSCSI host port identifier names have a maximum of 223 characters.◦ The port must be unique.◦ A port number that has already been configured is not allowed.• Delete — Use Delete to remove (unassociate) a host port identifier. The Delete option does not physically remove the host port. This option removes the association between the host port and the host. Unless you remove the host bus adapter or the iSCSI initiator, the host port is still recognized by the controller. <div> If you delete a host port identifier, it is no longer associated with this host. Also, the host loses access to any of its assigned volumes through this host port identifier.</div>
Label	To change the port label name, click the Edit icon (pencil). The port label name must be unique. A label name that has already been configured is not allowed.
CHAP Secret	<p>Appears only for iSCSI hosts. You can set or change the CHAP secret for the initiators (iSCSI hosts).</p> <p>System Manager uses the Challenge Handshake Authentication Protocol (CHAP) method, which validates the identity of targets and initiators during the initial link. Authentication is based on a shared security key called a CHAP secret.</p>

5. Click **Save**.

Delete host or host cluster

You can delete a host or host cluster.

About this task

Keep these guidelines in mind when you delete a host or a host cluster:

- Any specific volume assignments are deleted, and the associated volumes are available for a new assignment.
- If the host is part of a host cluster that has its own specific assignments, the host cluster is unaffected. However, if the host is part of a host cluster that does not have any other assignments, the host cluster and any other associated hosts or host port identifiers inherit any default assignments.
- Any host port identifiers that were associated with the host become undefined.

Steps

1. Select **Storage › Hosts**.
2. Select the host or host cluster that you want to delete, and then click **Delete**.

The **confirmation** dialog box appears.

3. Confirm that you want to perform the operation, and then click **Delete**.

Results

If you deleted a host, the system performs the following actions:

- Deletes the host and, if applicable, removes it from the host cluster.
- Removes access to any assigned volumes.
- Returns the associated volumes to an unassigned state.
- Returns any host port identifiers associated with the host to an unassociated state.

If you deleted a host cluster, the system performs the following actions:

- Deletes the host cluster and its associated hosts (if any).
- Removes access to any assigned volumes.
- Returns the associated volumes to an unassigned state.
- Returns any host port identifiers associated with the hosts to an unassociated state.

FAQs

What are hosts and host clusters?

A host is a server that sends I/O to a volume on a storage array. A host cluster is a group of hosts. You create a host cluster to make it easy to assign the same volumes to multiple hosts.

You define a host separately. It can either be an independent entity or be added to a host cluster. You can assign volumes to an individual host, or a host can be part of a host cluster that shares access to one or more volumes with other hosts in the host cluster.

The host cluster is a logical entity that you create in SANtricity System Manager. You must add hosts to the host cluster before you can assign volumes.

Why would I need to create a host cluster?

You need to create a host cluster if you want to have two or more hosts share access to the same set of volumes. Normally, the individual hosts have clustering software installed on them to coordinate volume access.

How do I know which host operating system type is correct?

The Host Operating System Type field contains the operating system of the host. You can select the recommended host type from the drop-down list or allow the Host Context Agent (HCA) to configure the host and appropriate host operating system type.

The host types that appear in the drop-down list depend on the storage array model and the firmware version. The most recent versions display the most common options first, which are the most likely to be appropriate. Appearance on this list does not imply the option is fully supported.



For more information about host support, refer to the [NetApp Interoperability Matrix](#) tool.

Some of the following host types might appear in the list:

Host Operating System type	Operating System (OS) and multipath driver
Linux DM-MP (Kernel 3.10 or later)	Supports Linux operating systems using a Device Mapper multipath failover solution with a 3.10 or later Kernel.
VMware ESXi	Supports VMware ESXi operating systems running the Native Multipathing Plug-in (NMP) architecture using the VMware built-in Storage Array Type Policy module SATP_ALUA.
Windows (clustered or non-clustered)	Supports Windows clustered or non-clustered configurations that are not running the ATTO multipathing driver.
ATTO Cluster (all operating systems)	Supports all cluster configurations using the ATTO Technology, Inc., multipathing driver.
Linux (Veritas DMP)	Supports Linux operating systems using a Veritas DMP multipathing solution.
Linux (ATTO)	Supports Linux operating systems using an ATTO Technology, Inc., multipathing driver.
Mac OS (ATTO)	Supports Mac OS versions using an ATTO Technology, Inc., multipathing driver.

Host Operating System type	Operating System (OS) and multipath driver
Windows (ATTO)	Supports Windows operating systems using an ATTO Technology, Inc., multipathing driver.
FlexArray (ALUA)	Supports a NetApp FlexArray system using ALUA for multipathing.
IBM SVC	Supports an IBM SAN Volume Controller configuration.
Factory Default	Reserved for the initial start-up of the storage array. If your host operating system type is set to Factory Default, change it to match the host operating system and multipath driver running on the connected host.
Linux DM-MP (Kernel 3.9 or earlier)	Supports Linux operating systems using a Device Mapper multipath failover solution with a 3.9 or earlier Kernel.
Window Clustered (deprecated)	If your host operating system type is set to this value, use the Windows (clustered or non-clustered) setting instead.

After the HCA is installed and the storage is attached to the host, the HCA sends the host topology to the storage controllers through the I/O path. Based on the host topology, the storage controllers automatically define the host and the associated host ports, and then set the host type.



If the HCA does not select the recommended host type, you must manually set the host type in System Manager.

What are HBAs and adapter ports?

A host bus adapter (HBA) is a board that resides in a host and contains one or more host ports. A host port is a port on a host bus adapter (HBA) that provides the physical connection to a controller and is used for I/O operations.

The adapter ports on the HBA are called host ports. Most HBAs have either one or two host ports. The HBA has a unique World Wide Identifier (WWID), and each HBA host port has a unique WWID. The host port identifiers are used to associate the appropriate HBA with the physical host when you are either manually creating the host through SANtricity System Manager or automatically creating the host using the host context agent.

How do I match the host ports to a host?

If you are manually creating a host, you first must use the appropriate host bus adapter (HBA) utility available on the host to determine the host port identifiers associated with each HBA installed in the host.

When you have this information, select the host port identifiers that have logged into the storage array from the

list provided in the **Create Host** dialog of System Manager.



Make sure you select the appropriate host port identifiers for the host you are creating. If you associate the wrong host port identifiers, you might cause unintended access from another host to this data.

If you are automatically creating hosts using the host context agent (HCA) installed on each host, the HCA should automatically associate the host port identifiers with each host and configure them appropriately.

How do I create CHAP secrets?

If you set up Challenge Handshake Authentication Protocol (CHAP) authentication on any iSCSI host connected to the storage array, you must re-enter that initiator CHAP secret for each iSCSI host. To do this, you can use System Manager either as part of the Create Host operation or through the View/Edit Settings option.

If you are using CHAP mutual authentication, you also must define a target CHAP secret for the storage array in the Settings page and re-enter that target CHAP secret on each iSCSI host.

What is the default cluster?

The default cluster is a system-defined entity that allows any unassociated host bus adapter (HBA) host port identifier that has logged into the storage array to gain access to any volumes assigned to the default cluster. An unassociated host port identifier is a host port that while physically installed in a host and logged into the storage array is not logically associated with a particular host.



If you want your hosts to have specific access to certain volumes in the storage array, you must *not* use the default cluster. Instead, you must associate the host port identifiers with their corresponding hosts. This can be done either manually using System Manager during the Create Host operation or automatically using the host context agent (HCA) installed on each host. Then, you assign volumes either to an individual host or to a host cluster.

You should *only* use the default cluster in special situations where your external storage environment is conducive to allowing all the hosts and all the logged-in host port identifiers connected to the storage array have access to all of the volumes (all-access mode) without specifically making the hosts known to the storage array or System Manager.

Initially, you can assign volumes only to the default cluster through the command line interface (CLI). However, after you assign at least one volume to the default cluster, this entity (called Default Cluster) is displayed in System Manager, and you can then use System Manager to manage this entity.

Performance

Concepts

Performance overview

The Performance page provides graphs and tables of data that enable you to assess the storage array's performance in several key areas.

Performance functions allow you to accomplish these tasks:

- View performance data in near real-time to help you determine whether a storage array is experiencing problems.
- Export performance data to construct a historical view of a storage array and identify when a problem started or what caused a problem.
- Select the objects, performance metrics, and time frame you want to view.
- Compare metrics.

You can view performance data in three formats:

- **Real-time graphical** — Plots performance data on a graph in near real-time.
- **Near real-time tabular** — Lists performance data in a table in near real-time.
- **Exported CSV file** — Allows you to save tabular performance data in a file of comma-separated values for further viewing and analysis.

Characteristics of performance data formats

Type of performance monitoring	Sampling interval	Length of time displayed	Maximum number of objects displayed	Ability to save data
Real-time graphical, live Real-time graphical, historical	10 sec (live) 5 min (historical) Data points shown depend on selected time frame	Default time frame is 1 hour. Choices: <ul style="list-style-type: none">• 5 minutes• 1 hour• 8 hours• 1 day• 7 days• 30 days	5	No
Near real-time tabular (table view)	10 sec -1 hr	Most current value	Unlimited	Yes
Comma-separated values (CSV) file	Depends on selected time frame	Depends on selected time frame	Unlimited	Yes

Guidelines for viewing performance data

- Performance data collection is always on. There is no option to turn it off.
- Each time the sampling interval elapses, the storage array is queried and the data is updated.
- For graphical data, the 5-minute time frame supports 10-second updating averaged over 5 minutes. All other time frames are updated every 5 minutes, averaged over the selected time frame.
- Performance data in the graphical views is updated in real time. Performance data in the table view is

updated in near real time.

- If a monitored object changes during the time data is collected, the object might not have a complete set of data points spanning the selected time frame. For example, volume sets can change as volumes are created, deleted, assigned, or unassigned; or drives can be added, removed, or failed.

Performance terminology

Learn how the performance terms apply to your storage array.

Term	Description
Application	An application is a software program, such as SQL or Exchange.
CPU	CPU is short for "central processing unit." CPU indicates the percentage of the storage array's processing capacity being used.
Host	A host is a server that sends I/O to a volume on a storage array.
IOPS	IOPS stands for input/output operations per second.
Latency	Latency is the time interval between a request, such as for a read or write command, and the response from the host or the storage array.
LUN	<p>A logical unit number (LUN) is the number assigned to the address space that a host uses to access a volume. The volume is presented to the host as capacity in the form of a LUN.</p> <p>Each host has its own LUN address space. Therefore, the same LUN can be used by different hosts to access different volumes.</p>
MiB	MiB is an abbreviation for mebibyte (mega binary byte). One MiB is 220, or 1,048,576 bytes. Compare with MB, which signifies a base 10 value. One MB equals 1,024 bytes.
Object	<p>An object is any logical or physical storage component.</p> <p>Logical objects include volume groups, pools, and volumes. Physical objects include the storage array, array controllers, hosts, and drives.</p>

Term	Description
Pool	A pool is a set of drives that is logically grouped. You can use a pool to create one or more volumes accessible to a host. (You create volumes from either a pool or a volume group.)
Read	Read is short for "read operation," which occurs when the host requests data from the storage array.
Volume	<p>A volume is a container in which applications, databases, and file systems store data. It is the logical component created for the host to access storage on the storage array.</p> <p>A volume is created from the capacity available in a pool or a volume group. A volume has a defined capacity. Although a volume might consist of more than one drive, a volume appears as one logical component to the host.</p>
Volume name	A volume name is a string of characters assigned to the volume when it is created. You can either accept the default name or provide a more descriptive name indicating the type of data stored in the volume.
Volume group	A volume group is a container for volumes with shared characteristics. A volume group has a defined capacity and RAID level. You can use a volume group to create one or more volumes accessible to a host. (You create volumes from either a volume group or a pool.)
Workload	A workload is a storage object that supports an application. You can define one or more workloads, or instances, per application. For some applications, System Manager configures the workload to contain volumes with similar underlying volume characteristics. These volume characteristics are optimized based on the type of application the workload supports. For example, if you create a workload that supports a Microsoft SQL Server application and then subsequently create volumes for that workload, the underlying volume characteristics are optimized to support Microsoft SQL Server.
Write	Write is short for "write operation," when data is sent from the host to the array for storage.

How tos

View graphical performance data

You can view graphical performance data for logical objects, physical objects, applications, and workloads.

About this task

The performance graphs show historical data as well as live data currently being captured. A vertical line on the graph, labeled Live updating, distinguishes historical data from live data.

Home page view

The Home page contains a graph showing storage array level performance. You can select limited metrics from this view, or you can click **View Performance Details** to select all the available metrics.

Detailed view

The graphs available from the detailed performance view are arranged under three tabs:

- **Logical View** — Displays performance data for logical objects grouped by volume groups and pools. Logical objects include volume groups, pools, and volumes.
- **Physical View** — Displays performance data for the controller, host channels, drive channels, and drives.
- **Applications & Workloads View** — Displays a list of logical objects (volumes) grouped by the application types and workloads you have defined.

Steps

1. Select **Home**.
2. To select an array-level view, click the IOPS, MiB/s, or CPU button.
3. To see more details, click **View Performance Details**.
4. Select **Logical View** tab, **Physical View** tab, or **Applications & Workloads View** tab.

Depending on the object type, different graphs appear in each tab.

View tabs	Performance data displayed for each object type
Logical View	<ul style="list-style-type: none">• Storage array: IOPS, MiB/s• Pools: Latency, IOPS, MiB/s• Volume groups: Latency, IOPS, MiB/s• Volumes: Latency, IOPS, MiB/s
Physical View	<ul style="list-style-type: none">• Controllers: IOPS, MiB/s, CPU, Headroom• Host channels: Latency, IOPS, MiB/s, Headroom• Drive channels: Latency, IOPS, MiB/s• Drives: Latency, IOPS, MiB/s

View tabs	Performance data displayed for each object type
Applications & Workloads View	<ul style="list-style-type: none"> • Storage array: IOPS, MiB/s • Applications: Latency, IOPS, MiB/s • Workloads: Latency, IOPS, MiB/s • Volumes: Latency, IOPS, MiB/s


5. Use the options to view the objects and information you need.

Options

Options for viewing objects	Description
Expand a drawer to see the list of objects.	<p><i>Navigation drawers</i> contain storage objects, such as pools, volume groups, and drives.</p> <p>Click the drawer to view the list of objects in the drawer.</p>
Select objects to view.	Select the check box to the left of each object to choose the performance data you want to view.
Use Filter to find object names or partial names.	In the Filter box, enter the name or a partial name of objects to list just those objects in the drawer.
Click Refresh Graphs after selecting objects.	After selecting objects from the drawers, select Refresh Graphs to view graphical data for the items you have selected.
Hide or show graph	Select the graph title to hide or show the graph.

6. As needed, use the additional options for viewing performance data.

Additional options

Option	Description
Time frame	<p>Select the length of time you want to view (5 minutes, 1 hour, 8 hours, 1 day, 7 days, or 30 days). The default is 1 hour.</p> <div><p>Loading performance data for a 30-day time frame can take several minutes. Do not navigate away from the web page, refresh the web page, or close the browser while data is loading.</p></div>
Data point details	<p>Hover the cursor over the graph to see metrics for a particular data point.</p>
Scroll bar	<p>Use the scroll bar below the graph to view an earlier or later time span.</p>
Zoom bar	<p>Below the graph, drag the zoom bar handles to zoom out on a time span. The wider the zoom bar, the less granular the details of the graph.</p> <p>To reset the graph, select one of the time frame options.</p>
Drag and drop	<p>On the graph, drag the cursor from one point in time to another to zoom in on a time span.</p> <p>To reset the graph, select one of the time frame options.</p>

= View and save tabular performance data :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

You can view and save performance graphs data in tabular format. This allows you to filter the data you want displayed.

Steps

1. From any performance data graph, click **Launch table view**.

A table appears that lists all the performance data for the selected objects.

2. Use the object selection pull-down and the filter as needed.
3. Click the Show/Hide Columns button to select the columns you want to include in the table.

You can click each check box to select or deselect an item.

4. Select **Export** at the bottom of the screen to save the tabular view to a file of comma-separated values (CSV).

The **Export Table** dialog box appears, indicating the number of rows to be exported and the file format of the export (comma-separated values, or CSV format).

5. Click **Export** to proceed with the download, or click **Cancel**.

Depending on your browser settings, the file is either saved, or you are prompted to choose a name and location for the file.

The default file name format is `performanceStatistics-yyyy-mm-dd_hh-mm-ss.csv`, which includes the date and time when the file was exported.

= Interpret performance data :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

Performance data can guide you in tuning the performance of your storage array.

When interpreting Performance data, keep in mind that several factors affect the performance of your storage array. The following table describes the main areas to consider.

Performance data	Implications for performance tuning
Latency (milliseconds, or ms)	<p data-bbox="841 155 1365 189">Monitor the I/O activity of a specific object.</p> <p data-bbox="841 222 1422 256">Potentially identify objects that are bottlenecks:</p> <ul data-bbox="867 289 1453 1306" style="list-style-type: none"> <li data-bbox="867 289 1453 457">• If a volume group is shared among several volumes, the individual volumes might need their own volume groups to improve the sequential performance of the drives and decrease latency. <li data-bbox="867 478 1453 613">• With pools, larger latencies are introduced and uneven workloads might exist between drives, making the latency values less meaningful and, in general, higher. <li data-bbox="867 634 1453 768">• Drive type and speed influence latency. With random I/O, faster spinning drives spend less time moving to and from different locations on the disk. <li data-bbox="867 789 1453 924">• Too few drives result in more queued commands and a greater period of time for the drive to process the command, increasing the general latency of the system. <li data-bbox="867 945 1453 1037">• Larger I/Os have greater latency due to the additional time involved with transferring data. <li data-bbox="867 1058 1453 1192">• Higher latency might indicate that the I/O pattern is random in nature. Drives with random I/O will have greater latency than those with sequential streams. <li data-bbox="867 1213 1453 1306">• A disparity in latency among drives or volumes of a common volume group could indicate a slow drive.

Performance data	Implications for performance tuning
IOPS	<p>Factors that affect input/output operations per second (IOPS or IOs/sec) include these items:</p> <ul style="list-style-type: none"> • Access pattern (random or sequential) • I/O size • RAID level • Cache block size • Whether read caching is enabled • Whether write caching is enabled • Dynamic cache read prefetch • Segment size • The number of drives in the volume groups or storage array <p>The higher the cache hit rate, the higher I/O rates will be. Higher write I/O rates are experienced with write caching enabled compared to disabled. In deciding whether to enable write caching for an individual volume, look at the current IOPS and the maximum IOPS. You should see higher rates for sequential I/O patterns than for random I/O patterns. Regardless of your I/O pattern, enable write caching to maximize the I/O rate and to shorten the application response time.</p> <p>You can see performance improvements caused by changing the segment size in the IOPS statistics for a volume. Experiment to determine the optimal segment size, or use the file system size or database block size.</p>
MiB/s	<p>Transfer or throughput rates are determined by the application I/O size and the I/O rate. Generally, small application I/O requests result in a lower transfer rate but provide a faster I/O rate and shorter response time. With larger application I/O requests, higher throughput rates are possible.</p> <p>Understanding your typical application I/O patterns can help you determine the maximum I/O transfer rates for a specific storage array.</p>

Performance data	Implications for performance tuning
CPU	<p>This value is a percentage of processing capacity that is being used.</p> <p>You might notice a disparity in the CPU usage of the same types of objects. For example, the CPU usage of one controller is heavy or is increasing over time while that of the other controller is lighter or more stable. In this case, you might want to change the controller ownership of one or more volumes to the controller with the lower CPU percentage.</p> <p>You might want to monitor CPU across the storage array. If CPU continues to increase over time while application performance decreases, you might need to add storage arrays. By adding storage arrays to your enterprise, you can continue to meet application needs at an acceptable performance level.</p>
Headroom	<p>Headroom refers to the remaining performance capability of the controllers, the controller host channels, and the controller drive channels. This value is expressed as a percentage and represents the gap between the maximum possible performance these objects are able to deliver and the current performance levels.</p> <ul style="list-style-type: none"> • For the controllers, headroom is a percentage of maximum possible IOPS. • For the channels, headroom is a percentage of maximum throughput, or MiB/s. Read throughput, write throughput, and bidirectional throughput are included in the calculation.

= FAQs

= How do performance statistics for individual volumes relate to the total? :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

The statistics for pools and volume groups are calculated by aggregating all volumes, including reserved capacity volumes.

Reserved capacity is used internally by the storage system to support thin volumes, snapshots, and asynchronous mirroring, and are not visible to I/O hosts. As a result, the pool, controller, and storage array statistics may not add up to be the sum of the viewable volumes.

However, for application and workload statistics, only the visible volumes are aggregated.

= Why does data display as zero in the graphs and table? :icons: font :relative_path: ./sm-storage/

:imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

When a zero is displayed for a data point in the graphs and table, it means there is no I/O activity for the object for that point in time. This situation could occur because the host is not initiating I/O to that object, or it could be a problem with the object itself.

The historical data for the object is still available for viewing. The graphs and table will show non-zero data once I/O activity begins occurring for the object.

The following table lists the most common reasons why a data point value might be zero for any given object.

Array-level object type	Reason data displays as zero
Volume	<ul style="list-style-type: none">• Volume had no host assignment.
Volume group	<ul style="list-style-type: none">• Volume group is being imported.• Volume group does not contain a volume that is assigned to a host, and volume group does not contain any reserved capacity.
Drive	<ul style="list-style-type: none">• Drive has failed.• Drive has been removed.• Drive is in an unknown state.
Controller	<ul style="list-style-type: none">• Controller is offline.• Controller has failed.• Controller has been removed.• Controller is in an unknown state.
Storage array	<ul style="list-style-type: none">• Storage array does not contain volumes.

= What does the Latency graph show? :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

The Latency graph provides latency statistics, in milliseconds (ms), for volumes, volume groups, pools, applications, and workloads. This graph appears in the Logical View, Physical View, and Applications & Workloads View tabs.

Latency refers to any delay that occurs as data is read or written. Hover your cursor over a point on the graph to view the following values, in milliseconds (ms), for that point in time:

- Read time.
- Write time.
- Average I/O size.

= What does the IOPS graph show? :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

The IOPS graph displays statistics for input/output operations per second. On the Home page, this graph displays statistics for the storage array. In the Logical View, Physical View, and Applications & Workloads View tabs of the Performance tile, this graph displays statistics for the storage array, volumes, volume groups, pools, applications, and workloads.

IOPS is an abbreviation for *input/output (I/O) operations per second*. Hover your cursor over a point on the graph to view the following values for that point in time:

- Number of read operations.
- Number of write operations.
- Total read and write operations combined.

= What does the MiB/s graph show? :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

The MiB/s graph displays transfer speed statistics in mebibytes per second. On the Home page, this graph displays statistics for the storage array. In the Logical View, Physical View, and Applications & Workloads View tabs of the Performance tile, this graph displays statistics for the storage array, volumes, volume groups, pools, applications, and workloads.

MiB/s is an abbreviation for *mebibytes per second*, or 1,048,576 bytes per second. Hover your cursor over a point on the graph to view the following values for that point in time:

- The amount of data read.
- The amount of data written.
- The combined total amount of data read and written.

= What does the CPU graph show? :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

The CPU graph displays processing capacity statistics for each controller (controller A and controller B). CPU is an abbreviation for *central processing unit*. On the Home page, this graph displays statistics for the storage array. On the Physical View tab of the Performance tile, this graph displays statistics for the storage array and drives.

The CPU graph shows the percentage of CPU processing capacity being used against operations on the array. Even when no external I/O is occurring, the CPU utilization percentage can be non-zero because the storage operating system might be doing background operations and monitoring. Hover your cursor over a point on the graph to view a percentage of processing capability being used at that point in time.

= What does the Headroom graph show? :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

The Headroom graph is related to remaining performance capability for the storage array controllers. This graph is visible on the Home page and on the Physical View tab of the Performance tile.

The Headroom graph shows the remaining performance capability of the physical objects in the storage system. Hover your cursor over a point on the graph to view the percentages of IOPS and MiB/s capability remaining for controller A and for controller B.

= Snapshots

= Concepts


= Overview of snapshot storage :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

A snapshot image is a logical, read-only copy of volume content, captured at a particular point in time. You can use snapshots to protect against data loss.

Snapshot images also are useful for test environments. By creating a virtual copy of data, you can test data using the snapshot without altering the actual volume itself. In addition, hosts do not have write access to snapshot images, so your snapshots are always a secure backup resource.

The Snapshots feature uses copy-on-write technology to store snapshot images and use allocated reserved capacity. As snapshots are created, the Snapshots feature stores image data as follows:

- When a snapshot image is created, it exactly matches the base volume. The Snapshots feature uses copy-on-write technology. After the snapshot is taken, the first write to any block or set of blocks on the base volume causes the original data to be copied to the reserved capacity before writing the new data to the base volume.
- Subsequent snapshots include only changed data blocks. Before data is overwritten on the base volume, the Snapshots feature uses its copy-on-write technology to save the required images of the affected sectors to the snapshot reserved capacity.



¹ Base volume (physical disk capacity); ² Snapshots (logical disk capacity); ³ Reserved capacity (physical disk capacity)

- The reserved capacity stores original data blocks for portions of the base volume that have changed after the snapshot was taken and includes an index for tracking changes. Generally, the size of the reserved capacity defaults to 40 percent of the base volume. (If you need more

reserved capacity, you can increase reserved capacity.)

- Snapshot images are stored in a specific order, based on their timestamp. Only the oldest snapshot image of a base volume is available for manual deletion.

To restore data to a base volume, you can use either a snapshot volume or snapshot image:

- **Snapshot volume** — If you need to retrieve deleted files, create a snapshot volume from a known good snapshot image, and then assign it to the host.
- **Snapshot image** — If you need to restore a base volume to a specific point-in-time, use a previous snapshot image to roll back data to the base volume.

= Requirements and guidelines for snapshots :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

When creating and using snapshots, review the following requirements and guidelines.

== Snapshot images and snapshot groups

- Each snapshot image is associated with exactly one snapshot group.
- A snapshot group is created the first time you create a scheduled or instant snapshot image for an associated object. This creates reserved capacity.

You can view snapshot groups from the **Pools & Volume Groups** page.

- Scheduled snapshot images do not occur when the storage array is offline or powered off.
- If you delete a snapshot group that has a snapshot schedule, the snapshot schedule is also deleted.
- If you have a snapshot volume that you no longer need, you can reuse it, along with any associated reserved capacity, instead of deleting it. This creates a different snapshot volume of the same base volume. You can re-associate the snapshot volume or snapshot consistency group snapshot volume with the same snapshot image or a different snapshot image, as long as the snapshot image is in the same base volume.

== Snapshot consistency group

- A snapshot consistency group contains one snapshot group for each volume that is a member of the snapshot consistency group.
- You can associate a snapshot consistency group with only one schedule.
- If you delete a snapshot consistency group that has a snapshot schedule, the snapshot schedule is also deleted.
- You cannot individually manage a snapshot group that is associated with a snapshot consistency group. Instead, you must perform the manage operations (create snapshot image, delete snapshot image or snapshot group, and rollback snapshot image) at the snapshot consistency group level.

== Base volume

- A snapshot volume must have the same Data Assurance (DA) and security settings as the associated base volume.

- You cannot create a snapshot volume of a failed base volume.
- If the base volume resides on a volume group, the member volumes for any associated snapshot consistency group can reside on either a pool or volume group.
- If a base volume resides on a pool, all member volumes for any associated snapshot consistency group must reside on the same pool as the base volume.

== Reserved capacity

- Reserved capacity is associated with only one base volume.
- Using a schedule can result in a large number of snapshot images. Make sure you have sufficient reserved capacity for scheduled snapshots.
- The reserved capacity volume for a snapshot consistency group must have the same Data Assurance (DA) and security settings as its associated base volume for the snapshot consistency group member volume.

== Pending snapshot images

Snapshot image creation might remain in a Pending state in the following conditions:

- The base volume that contains this snapshot image is a member of an asynchronous mirror group.
- The base volume is currently in a synchronization operation. The snapshot image creation completes as soon as the synchronization operation is complete.

== Maximum number of snapshot images

- If a volume is a member of a snapshot consistency group, System Manager creates a snapshot group for that member volume. This snapshot group counts towards the maximum allowable number of snapshot groups per base volume.
- If you attempt to create a snapshot image on a snapshot group or snapshot consistency group, but the associated group has reached its maximum number of snapshot images, you have two options:
 - Enable automatic deletion for the snapshot group or snapshot consistency group.
 - Manually delete one or more snapshot images from the snapshot group or snapshot consistency group and retry the operation.

== Auto-deletion

If the snapshot group or snapshot consistency group is enabled for automatic deletion, System Manager deletes the oldest snapshot image when the system creates a new one for the group.

== Rollback operation

- You cannot perform the following actions when a rollback operation is in progress:
 - Delete the snapshot image that is being used for the rollback.
 - Create a new snapshot image for a base volume that is participating in a rollback operation.
 - Change the associated snapshot group's Repository-Full Policy.
- You cannot start a rollback operation when any of these operations are in progress:
 - Capacity expansion (adding capacity to a pool or volume group)


- Volume expansion (increasing the capacity of a volume)
- RAID level change for a volume group
- Segment size change for a volume
- You cannot start a rollback operation if the base volume is participating in a volume copy.
- You cannot start a rollback operation if the base volume is a secondary volume in a remote mirror.
- A rollback operation fails if any of the used capacity in the associated snapshot repository volume has unreadable sectors.

= Base volumes, reserved capacity, and snapshot groups :icons: font :relative_path: ./sm-storage/
:imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

A *base volume* is the volume used as the source for a snapshot image. A base volume can be either a thick volume or a thin volume and can reside in either a pool or volume group.

To take snapshots of the base volume, you can create an instant image at any time, or you can automate the process by defining a regular schedule for snapshots.

The following figure shows the relationship between snapshot objects and the base volume.



¹ Base volume; ² Snapshot objects in the group (images and reserved capacity); ³ Reserved capacity for the snapshot group

== Reserved capacity and snapshot groups

System Manager organizes snapshot images into *snapshot groups*. When System Manager establishes the snapshot group, it automatically creates associated *reserved capacity* to hold the snapshot images for the group and to keep track of subsequent changes to additional snapshots.

If the base volume resides in a volume group, the reserved capacity can be located in either a pool or volume group. If the base volume resides in a pool, the reserved capacity must be located in the same pool as the base volume.

Snapshot groups require no user action, but you can adjust reserved capacity on a snapshot group at any time. Additionally, you might be prompted to create reserved capacity when the following conditions are met:

- Any time you take a snapshot of a base volume that does not yet have a snapshot group, System

Manager automatically creates a snapshot group. This action also creates reserved capacity for the base volume that is used to store subsequent snapshot images.

- Any time you create a snapshot schedule for a base volume, System Manager automatically creates a snapshot group.

== Auto-deletion

When working with snapshots, use the default option to have auto-deletion turned on. Auto-deletion automatically deletes the oldest snapshot image when the snapshot group reaches the snapshot group limit of 32 images. If you turn off auto-deletion, then snapshot group limits are eventually exceeded, and you must take manual actions to configure snapshot group settings and manage reserved capacity.

= Snapshot schedules and snapshot consistency groups :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

Use schedules for collection of snapshot images, and use snapshot consistency groups to manage multiple base volumes.

To easily manage snapshot operations for base volumes, you can use the following features:

- **Snapshot schedule** — Automate snapshots for a single base volume.
- **Snapshot consistency group** — Manage multiple base volumes as one entity.


== Snapshot schedule

If you want to automatically take snapshots for a base volume, you can create a schedule. For example, you can define a schedule that takes snapshot images every Saturday at midnight, on the first of every month, or on any dates and times you decide. After the maximum of 32 snapshots is reached for a single schedule, you can suspend scheduled snapshots, create more reserved capacity, or you can delete snapshots. Snapshots can be deleted manually or by automating the deletion process. After a snapshot image is deleted, additional reserved capacity is available for reuse.

== Snapshot consistency group

You create a snapshot consistency group when you want to make sure snapshot images are taken on multiple volumes at the same time. Snapshot image actions are performed on the snapshot consistency group as a whole. For example, you can schedule synchronized snapshots of all volumes with the same timestamp. Snapshot consistency groups are ideal for applications that span multiple volumes, such as database applications that store logs on one volume and the database files on another volume.

The volumes included in a snapshot consistency group are called member volumes. When you add a volume to a consistency group, System Manager automatically creates new reserved capacity that corresponds to that member volume. You can define a schedule to automatically create a snapshot image of each member volume.




¹ Reserved capacity; ² Member volume; ³ Consistency group snapshot images

= Snapshot volumes :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

You can create a snapshot volume and assign it to a host if you want to read or write snapshot data. The snapshot volume shares the same characteristics as the base volume (RAID level, I/O characteristics, and so on).

When you create a snapshot volume, you can designate it as *read-only* or *read-write accessible*.

When you create read-only snapshot volumes, you do not need to add reserved capacity. When you create read-write snapshot volumes, you must add reserved capacity to provide write-access.



¹ Base volume; ² Host; ³ Read-only snapshot volume; ⁴ Read-write snapshot volume; ⁵ Reserved capacity

= Snapshot rollback :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

A rollback operation returns a base volume to a previous state, determined by the selected snapshot.


For the rollback, you can select a snapshot image from either of the following sources:

- **Snapshot image rollback**, for a full restore of a base volume.
- **Snapshot consistency group rollback**, which can be used to roll back one or more volumes.

During the rollback, the Snapshots feature preserves all snapshot images in the group. It also allows the host to access the base volume during this process, if needed for I/O operations.

When a rollback is launched, a background process sweeps through the logical block addresses (LBAs) for the base volume, and then finds copy-on-write data in the rollback snapshot image to be restored. Because the base volume is host-accessible for reads and writes, and all previously written data is available immediately, the reserved capacity volume must be large enough to contain all

changes while the rollback is processing. The data transfer continues as a background operation until the rollback completes.



¹ Base volume; ² Snapshot objects in a group; ³ Snapshot group reserved capacity

= Snapshot terminology :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

Learn how the snapshot terms apply to your storage array.

Term	Description
Snapshots feature	The Snapshots feature is used to create and manage images of volumes.


Term	Description
Snapshot image	A snapshot image is a logical copy of volume data, captured at a particular point-in-time. Like a restore point, snapshot images allow you to roll back to a known good data set. Although the host can access the snapshot image, it cannot directly read or write to it.
Base volume	A base volume is the source from which a snapshot image is created. It can be a thick or thin volume and is typically assigned to a host. The base volume can reside in either a volume group or disk pool.
Snapshot volume	A snapshot volume allows the host to access data in the snapshot image. The snapshot volume contains its own reserved capacity, which saves any modifications to the base volume without affecting the original snapshot image.
Snapshot group	A snapshot group is a collection of snapshot images from a single base volume.
Reserved capacity volume	A reserved capacity volume tracks which data blocks of the base volume are overwritten and the preserved content of those blocks.
Snapshot schedule	A snapshot schedule is a timetable for creating automated snapshot images. Through the schedule, you can control the frequency of image creations.
Snapshot consistency group	A snapshot consistency group is a collection of volumes that are treated as a single entity when a snapshot image is created. Each of these volumes has its own snapshot image, but all the images are created at the same point in time.
Snapshot consistency group member volume	Each volume that belongs to a snapshot consistency group is referred to as a member volume. When you add a volume to a snapshot consistency group, System Manager automatically creates a new snapshot group that corresponds to this member volume.
Rollback	A rollback is the process of returning data in a base volume to a previous point in time.

Term	Description
Reserved capacity	Reserved capacity is the physical allocated capacity that is used for any copy service operation and storage object. It is not directly readable by the host.


= Workflow for creating snapshot images and snapshot volumes :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

In System Manager, you can create snapshot images and snapshot volumes by following these steps.

== Workflow for creating snapshot images



== Workflow for creating snapshot volumes



= How tos

= Create snapshots and snapshot objects

= Create snapshot image :experimental: :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

You can manually create a snapshot image from a base volume or snapshot consistency group. This is also called an *instant snapshot* or *instant image*.

Before you begin

- The base volume must be optimal.
- The drive must be optimal.
- The snapshot group cannot be designated as “reserved.”
- The reserved capacity volume must have the same Data Assurance (DA) settings as the associated base volume for the snapshot group.

Steps

1. Do one of the following actions to create a snapshot image:
 - Select **Storage > Volumes**. Select the object (base volume or snapshot consistency group),

and then select **Copy Services › Create instant snapshot**.

- Select **Storage › Snapshots**. Select the **Snapshot Images** tab, and then select **Create › Instant snapshot image**. The **Create Snapshot Image** dialog box appears. Select the object (base volume or snapshot consistency group), and then click **Next**. If a previous snapshot image was created for the volume or snapshot consistency group, then the system creates the instant snapshot immediately. Otherwise, if this is the first time a snapshot image is created for the volume or snapshot consistency group, the **Confirm Create Snapshot Image** dialog box appears.

2. Click **Create** to accept the notification that reserved capacity is needed and to proceed to the **Reserve Capacity** step.

The **Reserve Capacity** dialog box appears.

3. Use the spinner box to adjust the capacity percentage, and then click **Next** to accept the candidate volume highlighted in the table.

The **Edit Settings** dialog box appears.

4. Select the settings for the snapshot image as appropriate, and confirm that you want to perform the operation.

Setting	Description
Snapshot image settings	
Snapshot image limit	Keep the check box selected if you want snapshot images automatically deleted after the specified limit; use the spinner box to change the limit. If you clear this check box, snapshot image creation stops after 32 images.
Reserved capacity settings	
Alert me when...	<p>Use the spinner box to adjust the percentage point at which the system sends an alert notification when the reserved capacity for a snapshot group is nearing full.</p> <p>When the reserved capacity for the snapshot group exceeds the specified threshold, use the advance notice to increase reserved capacity or to delete unnecessary objects before the remaining space runs out.</p>

Policy for full reserved capacity	<p>Choose one of the following policies:</p> <ul style="list-style-type: none"> • Purge oldest snapshot image: The system automatically purges the oldest snapshot image in the snapshot group, which releases the snapshot image reserved capacity for reuse within the group. • Reject writes to base volume: When the reserved capacity reaches its maximum defined percentage, the system rejects any I/O write request to the base volume that triggered the reserved capacity access.
-----------------------------------	---

Results

- System Manager displays the new snapshot image in the Snapshot Images table. The table lists the new image by timestamp and associated base volume or snapshot consistency group.
- Snapshot creation might remain in a Pending state because of the following conditions:
 - The base volume that contains this snapshot image is a member of an asynchronous mirror group.
 - The base volume is currently in a synchronization operation. The snapshot image creation completes as soon as the synchronization operation is complete.

= Schedule snapshot images :experimental: :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

You create a snapshot schedule to enable recovery in case of a problem with the base volume and to perform scheduled backups. Snapshots of base volumes or snapshot consistency groups can be created on a daily, weekly, or monthly schedule, at any time of day.

Before you begin

The base volume must be Optimal.

About this task

This task describes how to create a snapshot schedule for an existing snapshot consistency group or base volume.

You also can create a snapshot schedule at the same time you create a snapshot image of a base volume or snapshot consistency group.

Steps

1. Do one of the following actions to create a snapshot schedule:

- Select **Storage › Volumes**.

Select the object (volume or snapshot consistency group) for this snapshot schedule, and then select **Copy Services › Create snapshot schedule**.

- Select **Storage › Snapshots**.

Select the **Schedules** tab, and then click **Create**.

2. Select the object (volume or snapshot consistency group) for this snapshot schedule, and then click **Next**.

The **Create Snapshot Schedule** dialog box appears.

3. Do one of the following actions:

- **Use a previously defined schedule from another snapshot object.**

Make sure advanced options are displayed. Click **Show more options**. Click **Import Schedule**, select the object with the schedule you want to import, and then click **Import**.

- **Modify the basic or advanced options.**

In the upper right of the dialog box, click **Show more options** to display all options, and then refer to the following table.

Field	Description
Basic settings	
Select days	Select individual days of the week for snapshot images.
Start time	From the drop-down list, select a new start time for the daily snapshots (selections are provided in half-hour increments). The start time defaults to one half-hour ahead of the current time.
Time zone	From the drop-down list, select your array's time zone.
Advanced settings	

Field	Description
Day / month	<p>Choose one of the following options:</p> <ul style="list-style-type: none"> • Daily / Weekly — Select individual days for synchronization snapshots. You also can select the Select all days check box in the upper right if you want a daily schedule. • Monthly / Yearly — Select individual months for synchronization snapshots. In the On day(s) field, enter the days of the month for synchronizations to occur. Valid entries are 1 through 31 and Last. You can separate multiple days with a comma or semi-colon. Use a hyphen for inclusive dates. For example: 1,3,4,10-15,Last. You also can select the Select all months check box in the upper right if you want a monthly schedule.
Start time	From the drop-down list, select a new start time for the daily snapshots (selections are provided in half-hour increments). The start time defaults to one half-hour ahead of the current time.
Time zone	From the drop-down list, select your array's time zone.
Snapshots per day / Time between snapshots	Select the number of snapshot images to create per day. If you select more than one, also select the time between snapshot images. For multiple snapshot images, be sure that you have adequate reserved capacity.
Create snapshot image right now?	Select this check box to create an instant image in addition to the automatic images you are scheduling.
Start/End date or No end date	Enter the start date for synchronizations to begin. Also enter an end date or select No end date .

1. Do one of the following actions:

- If the object is a snapshot consistency group, click **Create** to accept the settings and create the schedule.
- If the object is a volume, click **Next** to allocate reserved capacity for the snapshot images.

The volume candidate table displays only the candidates that support the reserved capacity specified. Reserved capacity is the physical allocated capacity that is used for any copy service operation and storage object. It is not directly readable by the host.

2. Use the spinner box to allocate the reserved capacity for the snapshot images. Do one of the following actions:

- **Accept the default settings.**

Use this recommended option to allocate the reserved capacity for the snapshot images with the default settings.

- **Allocate your own reserved capacity settings to meet your data storage needs.**

If you change the default reserved capacity setting, click **Refresh Candidates** to refresh the candidate list for the reserved capacity you specified.

Allocate the reserved capacity using the following guidelines:

- The default setting for reserved capacity is 40% of the capacity of the base volume. Usually this capacity is sufficient.
- The capacity needed varies, depending on the frequency and size of I/O writes to the volumes and the quantity and duration of snapshot image collection.

3. Click **Next**.

The **Edit Settings** dialog box appears.

4. Edit the settings for the snapshot schedule as needed, and then click **Finish**.

Setting	Description
Snapshot image limit	
Enable automatic deletion of snapshot images when...	Keep the check box selected if you want snapshot images automatically deleted after the specified limit; use the spinner box to change the limit. If you clear this check box, snapshot image creation stops after 32 images.
Reserved capacity settings	
Alert me when...	<p>Use the spinner box to adjust the percentage point at which the system sends an alert notification when the reserved capacity for a schedule is nearing full.</p> <p>When the reserved capacity for the schedule exceeds the specified threshold, use the advance notice to increase reserved capacity or to delete unnecessary objects before the remaining space runs out.</p>

Setting	Description
Policy for full reserved capacity	<p>Choose one of the following policies:</p> <ul style="list-style-type: none"> • Purge oldest snapshot image — The system automatically purges the oldest snapshot image, which releases the snapshot image reserved capacity for reuse within the snapshot group. • Reject writes to base volume — When the reserved capacity reaches its maximum defined percentage, the system rejects any I/O write request to the base volume that triggered the reserved capacity access.

= Create snapshot consistency group :experimental: :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

To ensure that you have consistent copies, you can create a set of multiple volumes called a *snapshot consistency group*. This group allows you to make snapshot images of all the volumes at the same time for consistency. Each volume that belongs to a snapshot consistency group is referred to as a *member volume*. When you add a volume to a snapshot consistency group, the system automatically creates a new snapshot group that corresponds to this member volume.

About this task

The snapshot consistency group creation sequence lets you select member volumes for the group and allocate capacity to the member volumes.

The process to create a snapshot consistency group is a multi-step procedure.

== Step 1: Add members to snapshot consistency group

Select members to specify a collection of volumes that comprise the snapshot consistency group. Any actions you perform on the snapshot consistency group extend uniformly to selected member volumes.

Before you begin

The member volumes must be Optimal.

Steps

1. Select **Storage > Snapshots**.
2. Click the **Snapshot Consistency Groups** tab.
3. Select **Create > Snapshot consistency group**.

The **Create Snapshot Consistency Group** dialog box appears.

4. Select the volume(s) to be added as member volumes to the snapshot consistency group.
5. Click **Next**, and go to [\[Step 2: Reserve capacity for snapshot consistency group\]](#).

== Step 2: Reserve capacity for snapshot consistency group

Associate reserved capacity to the snapshot consistency group. System Manager suggests the volumes and capacity based on the properties of the snapshot consistency group. You can accept the recommended reserved capacity configuration or customize the allocated storage.

About this task

On the **Reserve Capacity** dialog box, the volume candidate table displays only the candidates that support the reserved capacity specified. Reserved capacity is the physical allocated capacity that is used for any copy service operation and storage object. It is not directly readable by the host.

Steps

1. Use the spinner box to allocate the reserved capacity for the snapshot consistency group. Do one of the following actions:
 - **Accept the default settings.**

Use this recommended option to allocate the reserved capacity for each member volume with the default settings.

- **Allocate your own reserved capacity settings to meet your data storage needs.**

Allocate the reserved capacity using the following guidelines.

- The default setting for reserved capacity is 40% of the capacity of the base volume. Usually this capacity is sufficient.
- The capacity needed varies, depending on the frequency and size of I/O writes to the volumes and the quantity and duration of snapshot image collection.

2. **Optional:** If you change the default reserved capacity setting, click **Refresh Candidates** to refresh the candidate list for the reserved capacity you specified.

3. Click **Next**, and go to [\[Step 3: Edit settings for snapshot consistency group\]](#).

== Step 3: Edit settings for snapshot consistency group

Accept or choose automatic deletion settings and reserved capacity alert thresholds for the snapshot consistency group.

About this task

The snapshot consistency group creation sequence lets you select member volumes for the group and allocate capacity to the member volumes.

Steps

1. Accept or change the default settings for the snapshot consistency group as appropriate.

Setting	Description
Snapshot consistency group settings	
Name	Specify the name for the snapshot consistency group.
Enable automatic deletion of snapshot images when...	Keep the check box selected if you want snapshot images automatically deleted after the specified limit; use the spinner box to change the limit. If you clear this check box, snapshot image creation stops after 32 images.
Reserved capacity settings	
Alert me when...	<p>Use the spinner box to adjust the percentage point at which the system sends an alert notification when the reserved capacity for a snapshot consistency group is nearing full.</p> <p>When the reserved capacity for the snapshot consistency group exceeds the specified threshold, use the advance notice to increase reserved capacity or to delete unnecessary objects before the remaining space runs out.</p>

Setting	Description
Policy for full reserved capacity	<p>Choose one of the following policies:</p> <ul style="list-style-type: none"> • Purge oldest snapshot image — The system automatically purges the oldest snapshot image in the snapshot consistency group, which releases the snapshot image reserved capacity for reuse within the group. • Reject writes to base volume — When the reserved capacity reaches its maximum defined percentage, the system rejects any I/O write request to the base volume that triggered the reserved capacity access.

1. After you are satisfied with your snapshot consistency group configuration, click **Finish**.

= Create snapshot volume :experimental: :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

You create a snapshot volume to provide host access to a snapshot image of a volume or snapshot consistency group. You can designate the snapshot volume as either read-only or read-write.

About this task

The snapshot volume creation sequence lets you create a snapshot volume from a snapshot image and provides options to allocate reserved capacity if the volume is read/write. A snapshot volume can be designated as one of the following:

- A read-only snapshot volume provides a host application with read access to a copy of the data contained in the snapshot image, but without the ability to modify the snapshot image. A read-only snapshot volume does not have associated reserved capacity.
- A read-write snapshot volume provides the host application with write access to a copy of the data contained in the snapshot image. It has its own reserved capacity that is used to save any subsequent modifications made by the host application to the base volume without affecting the referenced snapshot image.

The process to create a snapshot volume is a multi-step procedure.

== Step 1: Review members for a snapshot volume

Select either a snapshot image of a base volume or a snapshot consistency group. If you select a snapshot consistency group snapshot image, the member volumes of the snapshot consistency group appear for review.

Steps

1. Select **Storage > Snapshots**.
2. Select the **Snapshot Volumes** tab.
3. Select **Create**.

The **Create Snapshot Volume** dialog box appears.

4. Select the snapshot image (volume or snapshot consistency group) you want to convert into a snapshot volume, and then click **Next**. Use a text entry in the **Filter** field to narrow down the list.

If the selection was for a snapshot consistency group snapshot image, the **Review Members** dialog box appears.

On the **Review Members** dialog box, review the list of volumes that are selected for conversion to snapshot volumes, and then click **Next**.

5. Go to [\[Step 2: Assign snapshot volume to host\]](#).

== Step 2: Assign snapshot volume to host Select a specific host or host cluster to assign it to the snapshot volume. This assignment grants a host or host cluster access to the snapshot volume. You can choose to assign a host later, if needed.

Before you begin

- Valid hosts or host clusters exist under the **Hosts** page.
- Host port identifiers must have been defined for the host.
- Before creating a DA-enabled volume, verify that your planned host connection supports the Data Assurance (DA) feature. If any of the host connections on the controllers in your storage array do not support DA, the associated hosts cannot access data on DA-enabled volumes.

About this task

When you assign volumes, keep these guidelines in mind:

- A host's operating system can have specific limits on how many volumes the host can access.
- You can define one host assignment for each snapshot volume in the storage array.
- Assigned volumes are shared between controllers in the storage array.
- The same logical unit number (LUN) cannot be used twice by a host or a host cluster to access a snapshot volume. You must use a unique LUN.

Assigning a volume to a host fails if you try to assign a volume to a host cluster that conflicts with an established assignment for a host in the host cluster.

Steps

1. On the **Assign to Host** dialog box, select the host or host cluster that you want to assign to the new volume. If you want to create the volume without assigning a host, select **Assign later** from the drop-down list.
2. Select the access mode. Choose one of the following:
 - **Read/write** — This option provides the host with read/write access to the snapshot volume and requires reserved capacity.
 - **Read only** — This option provides the host with read-only access to the snapshot volume and does not require reserved capacity.
3. Click **Next**, and do one of the following:
 - If your snapshot volume is read/write, the **Review Capacity** dialog box appears. Go to [\[Step 3: Reserve capacity for a snapshot volume\]](#).
 - If your snapshot volume is read only, the **Edit Priority** dialog box appears. Go to [\[Step 4: Edit settings for a snapshot volume\]](#).

== Step 3: Reserve capacity for a snapshot volume

Associate reserved capacity to a read/write snapshot volume. System Manager suggests the volumes and capacity based on the properties of the base volume or snapshot consistency group. You can accept the recommended reserved capacity configuration or customize the allocated storage.

About this task

You can increase or decrease the reserved capacity for the snapshot volume as needed. If you find that the snapshot reserved capacity is larger than you need, you can reduce its size to free up space that is needed by other logical volumes.

Steps

1. Use the spinner box to allocate the reserved capacity for the snapshot volume.

The Volume Candidate table displays only the candidates that support the reserved capacity specified.

Do one of the following actions:

- **Accept the default settings.**

Use this recommended option to allocate the reserved capacity for the snapshot volume with the default settings.

- **Allocate your own reserved capacity settings to meet your data storage needs.**

If you change the default reserved capacity setting, click **Refresh Candidates** to refresh the candidate list for the reserved capacity you specified.

Allocate the reserved capacity using the following guidelines.

- The default setting for reserved capacity is 40% of the capacity of the base volume, and usually this capacity is sufficient.
- The capacity needed varies, depending on the frequency and size of I/O writes to the volumes and the quantity and duration of snapshot image collection.

2. **Optional:** If you are creating the snapshot volume for a snapshot consistency group, the option to **Change candidate** appears in the Reserved Capacity Candidates table. Click **Change candidate** to select an alternate reserved capacity candidate.
3. Click **Next**, and go to [\[Step 4: Edit settings for a snapshot volume\]](#).

== Step 4: Edit settings for a snapshot volume


Change the settings for a snapshot volume such as its name, caching, reserved capacity alert thresholds, and so on.

About this task

You can add the volume to solid-state disk (SSD) cache as a way to improve read-only performance. SSD cache consists of a set of SSD drives that you logically group together in your storage array.

Steps

1. Accept or change the settings for the snapshot volume as appropriate.

Setting	Description
Snapshot volume settings	
Name	Specify the name for the snapshot volume.
Enable SSD Cache	Choose this option to enable read-only caching on SSDs.  This feature is not available on the EF600 storage system.
Reserved capacity settings	
Alert me when...	Appears only for a read/write snapshot volume. Use the spinner box to adjust the percentage point at which the system sends an alert notification when the reserved capacity for a snapshot group is nearing full. When the reserved capacity for the snapshot group exceeds the specified threshold, use the advance notice to increase reserved capacity or to delete unnecessary objects before the remaining space runs out.

1. Review the snapshot volume configuration. Click **Back** to make any changes.
2. When you are satisfied with your snapshot volume configuration, click **Finish**.

= Manage snapshot schedules

= Change the settings for a snapshot schedule :experimental: :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

For a snapshot schedule, you can change automatic collection times or the frequency of collection.

About this task

You can import settings from an existing snapshot schedule, or you can modify settings as needed.

Because a snapshot schedule is associated to a snapshot group or snapshot consistency group, reserved capacity might be affected by changes to schedule settings.

Steps

1. Select **Storage > Snapshots**.
2. Click the **Schedules** tab.
3. Select the snapshot schedule that you want to change, and then click **Edit**.

The **Edit Snapshot Schedule** dialog box appears.

4. Do one of the following:
 - **Use a previously defined schedule from another snapshot object** — Click **Import Schedule**, select the object with the schedule you want to import, and then click **Import**.
 - **Edit the schedule settings** — Refer to Field Details below.

Setting	Description
Day / month	<p>Choose one of the following options:</p> <ul style="list-style-type: none"> • Daily / Weekly — Select individual days for synchronization snapshots. You also can select the Select all days check box in the upper right if you want a daily schedule. • Monthly / Yearly — Select individual months for synchronization snapshots. In the On day(s) field, enter the days of the month for synchronizations to occur. Valid entries are 1 through 31 and Last. You can separate multiple days with a comma or semi-colon. Use a hyphen for inclusive dates. For example: 1,3,4,10-15,Last. You also can select the Select all months check box in the upper right if you want a monthly schedule.
Start time	<p>From the drop-down list, select a new start time for the daily snapshots. Selections are provided in half-hour increments. The start time defaults to one half-hour ahead of the current time.</p>

Setting	Description
Time zone	From the drop-down list, select your storage array's time zone.
Snapshots per day	Select the number of snapshot images to create per day.
Time between snapshots	If you select more than one, also select the time between restore points. For multiple restore points, be sure that you have adequate reserved capacity.
Start date	Enter the start date for synchronizations to begin. Also enter an end date or select No end date .
End date	
No end date	

1. Click **Save**.

= Activate and suspend snapshot schedule :experimental: :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

You can temporarily suspend scheduled collection of snapshot images when you need to conserve storage space. This method is more efficient than deleting and later re-creating the snapshot schedule.

About this task

The state of the snapshot schedule stays suspended until you use the **Activate** option to resume scheduled snapshot activity.

Steps

1. Select **Storage > Snapshots**.
2. If it is not already displayed, click the **Schedules** tab.

The schedules are listed on the page.

3. Select an active snapshot schedule that you want to suspend, and then click **Activate / Suspend**.

The **State** column status changes to **Suspended**, and the snapshot schedule stops collection of all snapshot images.

4. To resume collecting snapshot images, select the suspended snapshot schedule that you want to resume, and then click **Activate / Suspend**.

The **State** column status changes to **Active**.

= Delete snapshot schedule :experimental: :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

If you no longer want to collect snapshot images, you can delete an existing snapshot schedule.

About this task

When you delete a snapshot schedule, the associated snapshot images are not deleted along with it. If you think the collection of snapshot images might be resumed at some point, you should suspend the snapshot schedule rather than delete it.

Steps

1. Select **Storage > Snapshots**.
2. Click the **Schedules** tab.
3. Select the snapshot schedule that you want to delete, and confirm the operation.

Results

The system removes all schedule attributes from the base volume or snapshot consistency group.

= Manage snapshot images

= View snapshot image settings :experimental: :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

You can view the properties, status, reserved capacity, and associated objects assigned to each snapshot image.

About this task

Associated objects for a snapshot image include the base volume or snapshot consistency group for which this snapshot image is a restore point, the associated snapshot group, and any snapshot volumes created from the snapshot image. Use the snapshot settings to determine whether you want to copy or convert the snapshot image.

Steps

1. Select **Storage > Snapshots**.
2. Click the **Snapshot Images** tab.
3. Select the snapshot image that you want to view, and then click **View Settings**.

The **Snapshot Image Settings** dialog box appears.

4. View the settings for the snapshot image.

= Start snapshot image rollback for a base volume :experimental: :experimental: :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

You can perform a rollback operation to change the content of a base volume to match the content that is saved in a snapshot image. The rollback operation does not change the content of the snapshot images that are associated with the base volume.

Before you begin

- Enough reserved capacity is available to start a rollback operation.
- The selected snapshot image is Optimal and the selected volume is Optimal.
- The selected volume does not have a rollback operation already in progress.

About this task

The rollback start sequence lets you start rollback on a snapshot image of a base volume while providing options to add storage capacity. You cannot start more than one rollback operation for a base volume at a time.

The host can immediately access the new rolled-back base volume, but the existing base volume does not allow the host read-write access after the rollback begins. You can create a snapshot of the base volume just before starting the rollback to preserve the pre-rollback base volume for recovery.

Steps

1. Select **Storage > Snapshots**.
2. Select the **Snapshot Images** tab.
3. Select the snapshot image, and then select **Rollback > Start**.

The **Confirm Start Rollback** dialog box appears.

4. **Optional:** Select the option to **Increase Capacity** if necessary.

The **Increase Reserved Capacity** dialog box appears.

- a. Use the spinner box to adjust the capacity percentage.

If free capacity does not exist on the pool or volume group that contains the storage object you selected and the storage array has Unassigned Capacity, you can add capacity. You can create a new pool or volume group and then retry this operation using the new free capacity on that pool or volume group.

- b. Click **Increase**.

5. Confirm that you want to perform this operation, and then click **Rollback**.

Results

System Manager performs the following actions:

- Restores the volume with the content saved on the selected snapshot image.
- Makes the rolled-back volumes immediately available for host access. You do not need to wait for the rollback operation to complete.

Select **Home > View Operations in Progress** to view the progress of the rollback operation.

If the rollback operation is not successful, the operation pauses. You can resume the paused operation and, if still unsuccessful, follow the Recovery Guru procedure to correct the problem or contact technical support.

= Start snapshot image rollback for snapshot consistency group member volumes :experimental: :icons:
font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

You can perform a rollback operation to change the content of snapshot consistency group member volumes to match the content that is saved in a snapshot image. The rollback operation does not change the content of the snapshot images that are associated with the snapshot consistency group.

Before you begin

- Enough reserved capacity is available to start a rollback operation.
- The selected snapshot image is Optimal and the selected volume is Optimal.
- The selected volume does not have a rollback operation already in progress.

About this task

The rollback start sequence lets you start rollback on a snapshot image of a snapshot consistency group

while providing options to add storage capacity. You cannot start more than one rollback operation for a snapshot consistency group at a time.

The host has immediate access to the new rolled-back volumes, but the existing member volumes no longer allow host read-write access after the rollback starts. You can create a snapshot image of the member volumes just before starting the rollback to preserve the pre-rollback base volumes for recovery purposes.

The process to start rollback of a snapshot image of a snapshot consistency group is a multi-step procedure:

- [\[Step 1: Select members\]](#)
- [\[Step 2: Review capacity\]](#)
- [\[Step 3: Edit priority\]](#)

== Step 1: Select members

You must select the member volumes to be rolled back.

Steps

1. Select **Storage** > **Snapshots**.
2. Select the **Snapshot Images** tab.
3. Select the snapshot consistency group snapshot image, and then select **Rollback** > **Start**.

The **Start Rollback** dialog box appears.

4. Select the member volume or volumes.
5. Click **Next**, and do one of the following:
 - If any of the selected member volumes are associated with more than one reserved capacity object that stores snapshot images, the **Review Capacity** dialog box appears. Go to [Step 2: Review capacity](#).
 - If none of the selected member volumes are associated with more than one reserved capacity object that stores snapshot images, the **Edit Priority** dialog box appears. Go to [Step 3: Edit priority](#).

== Step 2: Review capacity

If you selected member volumes associated to more than one reserved capacity object, such as a snapshot group and reserved capacity volume, you can review and increase reserved capacity for the rolled-back volume(s).

Steps

1. Next to any member volumes with very low (or zero) reserved capacity, click the **Increase capacity** link in the **Edit** column.

The **Increase Reserved Capacity** dialog box appears.

2. Use the spinner box to adjust the capacity percentage, and then click **Increase**.

If free capacity does not exist on the pool or volume group that contains the storage object you selected and the storage array has Unassigned Capacity, you can add capacity. You can create a new pool or volume group and retry this operation using the new free capacity on that pool or volume group.

3. Click **Next**, and go to [Step 3: Edit priority](#).

The **Edit Priority** dialog box appears.

== Step 3: Edit priority

You can edit the priority of the rollback operation if needed.

About this task

The rollback priority determines how many system resources are dedicated to the rollback operation at the expense of system performance.

Steps

1. Use the slider to adjust the rollback priority as needed.
2. Confirm that you want to perform this operation, and then click **Finish**.

Results

System Manager performs the following actions:

- Restores the snapshot consistency group member volumes with the content saved on the selected snapshot image.
- Makes the rolled-back volumes immediately available for host access. You do not need to wait for the rollback operation to complete.

Select **Home** > **View Operations in Progress** to view the progress of the rollback operation.

If the rollback operation is not successful, the operation pauses. You can resume the paused operation and, if still unsuccessful, follow the Recovery Guru procedure to correct the problem or contact technical support.

= Resume snapshot image rollback :experimental: :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

If an error occurs during a snapshot image rollback operation, the operation is automatically paused. You can resume a rollback operation that is in a paused state.

Steps

1. Select **Storage** > **Snapshots**.
2. Click the **Snapshot Images** tab.
3. Highlight the paused rollback, and then select **Rollback** > **Resume**.

The operation resumes.

Results

System Manager performs the following actions:

- If the rollback operation resumes successfully, you can view the progress of the rollback operation in the Operations in Progress window.
- If the rollback operation is not successful, the operation pauses again. You can follow the Recovery Guru procedure to correct the problem or contact technical support.

= Cancel snapshot image rollback :experimental: :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

You can cancel an active rollback that is in progress (actively copying data), a pending rollback (in a pending queue awaiting resources to start), or a rollback that has been paused due to an error.

About this task

When you cancel a rollback operation that is in progress, the base volume reverts to an unusable state and appears as failed. Therefore, consider canceling a rollback operation only when recovery options exist for restoring the content of the base volume.

If the snapshot group on which the snapshot image resides has one or more snapshot images that have been automatically purged, the snapshot image used for the rollback operation might not be available for future rollbacks.

Steps

1. Select **Storage > Snapshots**.
2. Click the **Snapshot Images** tab.
3. Select the active or paused rollback, and then select **Rollback > Cancel**.

The **Confirm Cancel Rollback** dialog box appears.

4. Click **Yes** to confirm.

Results

System Manager stops the rollback operation. The base volume is usable but might have data that is inconsistent or not intact.

After you finish

After you cancel a rollback operation, you must take one of the following actions:

- Reinitialize the content of the base volume.
- Perform a new rollback operation to restore the base volume using either the same snapshot image that was used in the Cancel Rollback operation or a different snapshot image to perform the new rollback operation.

= Delete snapshot image :experimental: :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

You delete snapshot images to clean up the oldest snapshot image from a snapshot group or snapshot consistency group.

About this task

You can delete a single snapshot image, or you can delete snapshot images from snapshot consistency groups that have the same creation timestamp. You also can delete snapshot images from a snapshot group.

You cannot delete a snapshot image if it is not the oldest snapshot image for the associated base volume or snapshot consistency group.

Steps

1. Select **Storage > Snapshots**.
2. Click the **Snapshot Images** tab.
3. Select the snapshot image that you want to delete, and confirm that you want to perform the operation.

If you selected a snapshot image of a snapshot consistency group, select each member volume that you want to delete, and confirm that you want to perform the operation.

4. Click **Delete**.

Results

System Manager performs the following actions:

- Deletes the snapshot image from the storage array.
- Releases the reserved capacity for reuse within the snapshot group or snapshot consistency group.
- Disables all the associated snapshot volumes that exist for the deleted snapshot image.
- From a snapshot consistency group deletion, moves any member volume associated with the deleted snapshot image to a Stopped state.

= Manage snapshot consistency groups

= Add member volume to a snapshot consistency group :experimental: :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

You can add a new member volume to an existing snapshot consistency group. When you add a new member volume, you also must reserve capacity for the member volume.

Before you begin

- The member volume must be Optimal.
- The snapshot consistency group must have less than the maximum number of allowable volumes (as defined by your configuration).
- Each reserved capacity volume must have the same Data Assurance (DA) and security settings as the associated member volume.

About this task

You can add standard volumes or thin volumes to the snapshot consistency group. The base volume can reside in either a pool or volume group.

Steps

1. Select **Storage > Snapshots**.
2. Select the **Snapshot Consistency Groups** tab.

The table appears and displays all the snapshot consistency groups associated with the storage array.

3. Select the snapshot consistency group you want to modify, and then click **Add Members**.

The **Add Members** dialog box appears.

4. Select the member volume(s) you want to add, and then click **Next**.

The Reserve Capacity step appears. The Volume Candidate table displays only the candidates that support the reserved capacity specified.

5. Use the spinner box to allocate the reserved capacity for the member volume. Do one of the following actions:

- **Accept the default settings.**

Use this recommended option to allocate the reserved capacity for the member volume with the default settings.

- **Allocate your own reserved capacity settings to meet your data storage needs.**

If you change the default reserved capacity setting, click **Refresh Candidates** to refresh the candidate list for the reserved capacity you specified.

Allocate the reserved capacity using the following guidelines.

- The default setting for reserved capacity is 40% of the capacity of the base volume, and usually this capacity is sufficient.
- The capacity needed varies, depending on the frequency and size of I/O writes to the volumes and the quantity and duration of snapshot image collection.

6. Click **Finish** to add the member volumes.

= Remove a member volume from a snapshot consistency group :experimental: :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

You can remove a member volume from an existing snapshot consistency group.

About this task

When you remove a member volume from a snapshot consistency group, System Manager automatically deletes the snapshot objects associated with that member volume.

Steps

1. Select **Storage > Snapshots**.
2. Click the **Snapshot Consistency Groups** tab.
3. Expand the snapshot consistency group you want to modify by selecting the plus (+) sign next to it.
4. Select the member volume that you want to remove, and then click **Remove**.
5. Confirm that you want to perform the operation, and then click **Remove**.

Results

System Manager performs the following actions:

- Deletes all snapshot images and snapshot volumes associated with the member volume.
- Deletes the snapshot group associated with the member volume.
- The member volume is not otherwise changed or deleted.

= Change the settings for a snapshot consistency group :experimental: :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

Change the settings for a snapshot consistency group when you want to change its name, automatic deletion settings, or the maximum number of allowed snapshot images.

Steps

1. Select **Storage > Snapshots**.
2. Click the **Snapshot Consistency Groups** tab.
3. Select the snapshot consistency group that you want to edit, and then click **View/Edit Settings**.

The **Snapshot Consistency Group Settings** dialog box appears.

4. Change the settings for the snapshot consistency group as appropriate.

Setting	Description
Snapshot consistency group settings	
Name	You can change the name for the snapshot consistency group.
Auto-deletion	Keep the check box selected if you want snapshot images automatically deleted after the specified limit; use the spinner box to change the limit. If you clear this check box, snapshot image creation stops after 32 images.
Snapshot image limit	You can change the maximum number of snapshot images allowed for a snapshot group.
Snapshot schedule	This field indicates whether a schedule is associated with the snapshot consistency group.
Associated objects	
Member volumes	You can view the quantity of member volumes associated with the snapshot consistency group.

1. Click **Save**.

= Delete snapshot consistency group :experimental: :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

You can delete snapshot consistency groups that are no longer needed.

Before you begin

Confirm that the images for all member volumes are no longer needed for backup or testing purposes.

About this task

This operation deletes all the snapshot images or schedules associated with the snapshot consistency group.

Steps

1. Select **Storage > Snapshots**.
2. Select the **Snapshot Consistency Groups** tab.
3. Select the snapshot consistency group that you want to delete, and then select **Uncommon Tasks > Delete**.

The **Confirm Delete Snapshot Consistency Group** dialog box appears.

4. Confirm that you want to perform this operation, and then click **Delete**.

Results

System Manager performs the following actions:

- Deletes all existing snapshot images and snapshot volumes from the snapshot consistency group.
- Deletes all the associated snapshot images that exist for each member volume in the snapshot consistency group.
- Deletes all the associated snapshot volumes that exist for each member volume in the snapshot consistency group.
- Deletes all associated reserved capacity for each member volume in the snapshot consistency group (if selected).

= Manage snapshot volumes

= Convert snapshot volume to read-write mode :experimental: :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

You can convert a read-only snapshot volume or a snapshot consistency group snapshot volume to read-write mode if needed. A snapshot volume that is converted to read-write accessible contains its own reserved capacity. This capacity is used to save any subsequent modifications made by the host application to the base volume without affecting the referenced snapshot image.

Steps

1. Select **Storage > Snapshots**.

2. Select the **Snapshot Volumes** tab.

The Snapshot Volumes table appears and displays all the snapshot volumes associated with the storage array.

3. Select the read-only snapshot volume you want to convert, and then click **Convert to Read/Write**.

The **Convert to Read/Write** dialog box appears with the **Reserve Capacity** step activated. The Volume Candidate table displays only the candidates that support the reserved capacity specified.

4. To allocate the reserved capacity for the read-write snapshot volume, do one of the following actions:

- **Accept the default settings** — Use this recommended option to allocate the reserved capacity for the snapshot volume with the default settings.
- **Allocate your own reserved capacity settings to meet your data storage needs** — Allocate the reserved capacity using the following guidelines.
 - The default setting for reserved capacity is 40% of the capacity of the base volume, and usually this capacity is sufficient.
 - The capacity needed varies, depending on the frequency and size of I/O writes to the volume.

5. Select **Next** to review or edit settings.

The **Edit Settings** dialog box appears.

6. Accept or specify the settings for the snapshot volume as appropriate, and then select **Finish** to convert the snapshot volume.

Setting	Description
Reserved capacity settings	
Alert me when...	<p>Use the spinner box to adjust the percentage point at which the system sends an alert notification when the reserved capacity for a snapshot group is nearing full.</p> <p>When the reserved capacity for the snapshot volume exceeds the specified threshold, the system sends an alert, allowing you time to increase reserved capacity or to delete unnecessary objects.</p>

= Change the volume settings for a snapshot volume :experimental: :icons: font :relative_path: ./sm-storage/ :imagesdir: /tmp/d20240212-7976-wmsuyj/source/./sm-storage/./media/

You can change the settings for a snapshot volume or snapshot consistency group snapshot volume to rename it, enable or disable SSD caching, or change the host, host cluster, or logical unit number (LUN) assignment.

Steps

1. Select **Storage > Snapshots**.
2. Click the **Snapshot Volumes** tab.
3. Select the snapshot volume that you want to change, and then click **View/Edit Settings**.

The **Snapshot Volume Settings** dialog box appears.

4. View or edit the settings for the snapshot volume as appropriate.

Setting	Description
Snapshot volume	
Name	You can change the name for the snapshot volume.
Assigned to	You can change the host or host cluster assignment for the snapshot volume.
LUN	You can change the LUN assignment for the snapshot volume.
SSD Cache	You can enable/disable read-only caching on solid state disks (SSDs).
Associated objects	
Snapshot image	You can view the snapshot images associated with the snapshot volume. A snapshot image is a logical copy of volume data, captured at a particular point-in-time. Like a restore point, snapshot images allow you to roll back to a known good data set. Although the host can access the snapshot image, it cannot directly read or write to it.
Base volume	You can view the base volume associated with the snapshot volume. A base volume is the source from which a snapshot image is created. It can be a thick or thin volume and is typically assigned to a host. The base volume can reside in either a volume group or disk pool.
Snapshot group	You can view the snapshot group associated with the snapshot volume. A snapshot group is a collection of snapshot images from a single base volume.

Copy snapshot volume

You can perform a Copy Volume process on a snapshot volume or a snapshot consistency group snapshot volume.

About this task

You can copy a snapshot volume to the target volume as performed in a normal Copy Volume operation. However, snapshot volumes cannot remain online during the copy volume process.

Steps

1. Select **Storage > Snapshots**.
2. Select the **Snapshot Volumes** tab.

The Snapshot Volumes table appears and displays all the snapshot volumes associated with the storage array.

3. Select the snapshot volume that you want to copy, and then select **Copy Volume**.

The **Copy Volume** dialog box appears, prompting you to select a target.

4. Select the target volume to be used as the copy destination, and then click **Finish**.

Re-create snapshot volume

You can re-create a snapshot volume or a snapshot consistency group snapshot volume that you previously disabled. Re-creating a snapshot volume takes less time than creating a new one.

Before you begin

- The snapshot volume must be in either an Optimal or Disabled state.
- All member snapshot volumes must be in a Disabled state before you can re-create the snapshot consistency group snapshot volume.

About this task

You cannot re-create an individual member snapshot volume; you can re-create only the overall snapshot consistency group snapshot volume.



If the snapshot volume or snapshot consistency group snapshot volume is part of an online copy relationship, you cannot perform the re-create option on the volume.

Steps

1. Select **Storage > Snapshots**.
2. Select the **Snapshot Volumes** tab.

The Snapshot Volumes table appears and displays all the snapshot volumes associated with the storage array.

3. Select the snapshot volume that you want to re-create, and then select **Uncommon Tasks > Recreate**.

The **Recreate Snapshot Volume** dialog box appears.

4. Select one of the following options:

- **An existing snapshot image created from volume <name>**

Select this option to indicate an existing snapshot image from which to re-create the snapshot volume.

- **A new (instant) snapshot image of volume <name>**

Select this option to create a new snapshot image from which to re-create the snapshot volume.

5. Click **Recreate**.

Results

System Manager performs the following actions:

- Deletes all `write` data on any associated snapshot repository volume.
- Snapshot volume or snapshot consistency group snapshot volume parameters remain the same as the previously disabled volume parameters.
- Retains the original names for the snapshot volume or snapshot consistency group snapshot volume.

Disable snapshot volume

You can disable a snapshot volume or a snapshot volume in a snapshot consistency group when you no longer need it or want to temporarily stop using it.

About this task

Use the **Disable** option if one of these conditions applies:

- You are finished with the snapshot volume or snapshot consistency group snapshot volume for the time being.
- You intend to re-create the snapshot volume or snapshot consistency group snapshot volume (that is designated as read-write) at a later time and want to retain the associated reserved capacity so you do not need to create it again.
- You want to increase the storage array performance by stopping write activity to a read-write snapshot volume.

If the snapshot volume or snapshot consistency group snapshot volume is designated as read-write, this option also lets you stop any further write activity to its associated reserved capacity volume. If you decide to re-create the snapshot volume or snapshot consistency group snapshot volume, you must choose a snapshot image from the same base volume.



If the snapshot volume or snapshot consistency group snapshot volume is part of an online copy relationship, you cannot perform the **Disable** option on the volume.

Steps

1. Select **Storage > Snapshots**.
2. Select the **Snapshot Volumes** tab.

System Manager displays all the snapshot volumes associated with the storage array.

3. Select the snapshot volume that you want to disable, and then select **Uncommon Tasks > Disable**.

4. Confirm that you want to perform the operation, and then click **Disable**.

Results

- The snapshot volume remains associated with its base volume.
- The snapshot volume retains its World Wide Name (WWN).
- If read-write, the snapshot volume retains its associated reserved capacity.
- The snapshot volume retains any host assignments and access. However, read-write requests fail.
- The snapshot volume loses its association with its snapshot image.

Delete snapshot volume

You can delete a snapshot volume or a snapshot consistency group snapshot volume that is no longer needed for backup or software application testing purposes. You can also specify whether you want to delete the snapshot reserved capacity volume associated with a `read-write` snapshot volume or retain the snapshot reserved capacity volume as an unassigned volume.

About this task

Deleting a base volume automatically deletes any associated snapshot volume or consistency group snapshot volume. You cannot delete a snapshot volume that is in a volume copy with a status of **In Progress**.

Steps

1. Select **Storage > Snapshots**.
2. Select the **Snapshot Volumes** tab.

System Manager displays all the snapshot volumes associated with the storage array.

3. Select the snapshot volume that you want to delete, and then select **Uncommon Tasks > Delete**.
4. Confirm that you want to perform the operation, and then click **Delete**.

Results

System Manager performs the following actions:

- Deletes all member snapshot volumes (for a snapshot consistency group snapshot volume).
- Removes all associated host assignments.

FAQs

Why don't I see all my volumes, hosts, or host clusters?

Snapshot volumes with a DA-enabled base volume are not eligible to be assigned to a host that is not Data Assurance (DA) capable. You must disable DA on the base volume before a snapshot volume can be assigned to a host that is not DA capable.

Consider the following guidelines for the host to which you are assigning the snapshot volume:

- A host is not DA capable if it is connected to the storage array through an I/O interface that is not DA capable.

- A host cluster is not DA capable if it has at least one host member that is not DA capable.



You cannot disable DA on a volume that is associated with snapshots (consistency groups, snapshot groups, snapshot images, and snapshot volumes), volume copies, and mirrors. All associated reserved capacity and snapshot objects must be deleted before DA can be disabled on the base volume.

What is a snapshot image?

A snapshot image is a logical copy of volume content, captured at a particular point in time. Snapshot images use minimal storage space.

Snapshot image data is stored as follows:

- When a snapshot image is created, it exactly matches the base volume. After the snapshot is taken, when the first write request occurs for any block or set of blocks on the base volume, the original data is copied to the snapshot reserved capacity before the new data is written to the base volume.
- Subsequent snapshots include only data blocks that have changed since the first snapshot image was created. Each subsequent copy-on-write operation saves original data that is about to be overwritten on the base volume to the snapshot reserved capacity before the new data is written to the base volume.

Why use snapshot images?

You can use snapshots to protect against and allow recovery from accidental or malicious loss or corruption of data.

Select a base volume or a group of base volumes, called a snapshot consistency group, and then capture snapshot images in one or more of the following ways:

- You can create a snapshot image of a single base volume or a snapshot consistency group consisting of multiple base volumes.
- You can take snapshots manually or create a schedule for a base volume or snapshot consistency group to automatically capture periodic snapshot images.
- You can create a host-accessible snapshot volume of a snapshot image.
- You can perform a rollback operation to restore a snapshot image.

The system retains multiple snapshot images as restore points you can use to roll back to known good data sets at specific points in time. The ability to roll back provides protection against accidental deletion of data and data corruption.

What kinds of volumes can be used for snapshots?

Standard volumes and thin volumes are the only types of volumes that can be used to store snapshot images. Non-standard volumes cannot be used. The base volume can reside on either a pool or volume group.

Why would I create a snapshot consistency group?

You create a snapshot consistency group when you want to make sure snapshot images are taken on multiple volumes at the same time. For example, a database made up of

multiple volumes that need to stay consistent for recovery purposes would require a snapshot consistency group to collect coordinated snapshots of all volumes and use them to restore the entire database.

The volumes included in a snapshot consistency group are called member volumes.

You can perform the following snapshot operations on a snapshot consistency group:

- Create a snapshot image of a snapshot consistency group to get simultaneous images of the member volumes.
- Create a schedule for the snapshot consistency group to automatically capture periodic simultaneous images of the member volumes.
- Create a host-accessible snapshot volume of a snapshot consistency group image.
- Perform a rollback operation for a snapshot consistency group.

What is a snapshot volume and when does it need reserved capacity?

A snapshot volume allows the host to access data in the snapshot image. The snapshot volume contains its own reserved capacity, which saves any modifications to the base volume without affecting the original snapshot image. Snapshot images are not read- or write-accessible to hosts. If you want to read or write to snapshot data, create a snapshot volume and assign it to a host.

You can create two types of snapshot volumes. The type of snapshot volume determines if it uses reserved capacity.

- **Read-only** — A snapshot volume that is created as read-only provides a host application with read access to a copy of the data contained in the snapshot image. A read-only snapshot volume does not use reserved capacity.
- **Read-write** — A snapshot volume that is created as read-write allows you to make changes to the snapshot volume without affecting the referenced snapshot image. A read-write snapshot volume uses reserved capacity to store these changes. You can convert a read-only snapshot volume to read-write at any time.

What is a snapshot group?

A snapshot group is a collection of point-in-time snapshot images of a single associated base volume.

System Manager organizes snapshot images into *snapshot groups*. Snapshot groups require no user action, but you can adjust reserved capacity on a snapshot group at any time. Additionally, you might be prompted to create reserved capacity when the following conditions are met:

- Any time you take a snapshot of a base volume that does not yet have a snapshot group, System Manager automatically creates a snapshot group. This creates reserved capacity for the base volume that is used to store subsequent snapshot images.
- Any time you create a snapshot schedule for a base volume, System Manager automatically creates a snapshot group.

Why would I disable a snapshot volume?

You disable a snapshot volume when you want to assign a different snapshot volume to the snapshot image. You can reserve the disabled snapshot volume for later use.

If you no longer need the snapshot volume or the consistency group snapshot volume and do not intend to re-create it at a later time, you should delete the volume instead of disabling it.

What is the Disabled state?

A snapshot volume in Disabled status is not currently assigned to a snapshot image. To enable the snapshot volume, you must use the re-create operation to assign a new snapshot image to the disabled snapshot volume.

The snapshot volume characteristics are defined by the snapshot image assigned to it. Read-write activity is suspended on a snapshot volume in Disabled status.

Why would I suspend a snapshot schedule?

When a schedule is suspended, the scheduled snapshot image creations do not occur. You can pause a snapshot schedule to conserve storage space, and then resume the scheduled snapshots at a later time.

If you do not need the snapshot schedule, you should delete the schedule instead of suspending it.

Mirroring

Asynchronous mirroring

Concepts

How asynchronous mirroring works

Asynchronous mirroring copies data volumes on demand or on a schedule, which minimizes or avoids downtime that might result from data corruption or loss.



Mirroring is not available on the EF600 storage array.

Asynchronous mirroring captures the state of the primary volume at a particular point in time and copies just the data that has changed since the last image capture. The primary site can be updated immediately and the secondary site can be updated as bandwidth allows. The information is cached and sent later, as network resources become available.


This type of mirroring is ideal for satisfying the demand for non-stop operations and, in general, is far more network efficient for periodic processes, such as backup and archive. The reasons for using asynchronous mirroring include the following:

- Remote backup consolidation.
- Protect against local or wide-area disasters.
- Application development and testing on a point-in-time image of live data.

Asynchronous mirroring session

Asynchronous mirroring captures the state of the primary volume at a particular point in time and copies just the data that has changed since the last image capture. Asynchronous mirroring allows the primary site to be updated immediately and the secondary site to be updated as bandwidth allows. The information is cached and sent later, as network resources become available.

There are four primary steps in an active asynchronous mirroring session.



1. A write operation first occurs on the primary volume's storage array.
2. The status of the operation is returned to the host.
3. All changes on the primary volume are logged and tracked.
4. All changes are sent to the secondary volume's storage array as a background process.

These steps are repeated according to the defined synchronization intervals or the steps can be repeated manually if no intervals are defined.

Asynchronous mirroring transfers data to the remote site only at set intervals, so local I/O is not affected nearly as much by slow network connections. Because this transfer is not tied to the local I/O, it does not affect application performance. Therefore, asynchronous mirroring can use slower connections, such as iSCSI, and run across longer distances between the local and remote storage systems.

The storage arrays must have a minimum firmware version of 7.84. (They can each run different OS versions.)

Mirror consistency groups and mirrored pairs

You create a mirror consistency group to establish the mirroring relationship between the local storage array and the remote storage array. The asynchronous mirroring relationship consists of a mirrored pair: a primary volume on one storage array and a secondary volume on another storage array.

The storage array containing the primary volume is usually located at the primary site and serves the active hosts. The storage array containing the secondary volume is usually located at a secondary site and holds a replica of the data. The secondary volume typically contains a backup copy of the data and is used for disaster recovery.

Synchronization settings

When you create a mirrored pair, you also define the synchronization priority and resynchronization policy that the mirrored pair uses to complete the resynchronization operation after a communication interruption.

When you create a mirror consistency group, you also define the synchronization priority and resynchronization policy for all mirrored pairs within the group. The mirrored pairs use the synchronization priority and resynchronization policy to complete the resynchronization operation after a communication interruption.

The primary and secondary volumes in a mirrored pair can become unsynchronized when the primary volume's storage array is unable to write data to the secondary volume. This condition can be caused by the following issues:

- Network problems between the local and remote storage arrays.
- A failed secondary volume.
- Synchronization being manually suspended on the mirrored pair.
- Mirror group role conflict.

You can synchronize data on the remote storage array either manually or automatically.

Reserved capacity and asynchronous mirroring

Reserved capacity is used to keep track of differences between the primary and secondary volume when synchronization is not occurring. It also keeps track of synchronization statistics for each mirrored pair.

Each volume in a mirrored pair requires its own reserved capacity.

Configuration and management

To enable and configure mirroring between two arrays, you must use the Unified Manager interface. Once mirroring is enabled, you can manage mirrored pairs and synchronization settings in System Manager.

Asynchronous mirroring terminology

Learn how the asynchronous mirroring terms apply to your storage array.

Term	Description
Local storage array	<p>The local storage array is the storage array that you are acting upon.</p> <p>When you see Primary in the Local Role column, it indicates that the storage array contains the volume that holds the primary role in the mirror relationship. When you see Secondary in the Local Role column, it indicates that the storage array contains the volume that holds the secondary role in the mirror relationship.</p>
Mirror consistency group	A mirror consistency group is a container for one or more mirrored pairs. For asynchronous mirroring operations, you must create a mirror consistency group.
Mirrored pair	<p>A mirrored pair is comprised of two volumes, a primary volume and a secondary volume.</p> <p>In asynchronous mirroring, a mirrored pair always belongs to a mirror consistency group. Write operations are performed first to the primary volume and then replicated to the secondary volume. Each mirrored pair in a mirror consistency group share the same synchronization settings.</p>
Primary volume	The primary volume of a mirrored pair is the source volume to be mirrored.
Remote storage array	The remote storage array is usually designated as the secondary site, which usually holds a replica of the data in a mirroring configuration.
Reserved capacity	Reserved capacity is the physical allocated capacity that is used for any copy service operation and storage object. It is not directly readable by the host.
Role change	Role change is assigning the primary role to the secondary volume and vice versa.
Secondary volume	The secondary volume of a mirrored pair is usually located at a secondary site and holds a replica of the data.
Synchronization	Synchronization occurs at initial synchronization between the local storage array and the remote storage array. Synchronization also occurs when the primary and secondary volumes become unsynchronized after a communication interruption. When the communication link is working again, any unreplicated data is synchronized to the secondary volume's storage array.

Workflow for mirroring a volume asynchronously

You configure asynchronous mirroring using the following workflow.



This feature is not available on the EF600 storage system.

1. Perform the initial configuration in Unified Manager:
 - a. Select the local storage array as the source for the data transfer.
 - b. Create or select an existing mirror consistency group, which is a container for the primary volume on

the local array and the secondary volume on the remote array. The primary and secondary volumes are referred to as the "mirrored pair." If you are creating the mirror consistency group for the first time, you specify whether you want to perform manual or scheduled synchronizations.

- c. Select a primary volume from the local storage array, and then determine its reserved capacity. Reserved capacity is the physical allocated capacity to be used for the copy operation.
 - d. Select a remote storage array as the destination of the transfer, a secondary volume, and then determine its reserved capacity.
 - e. Begin the initial data transfer from the primary volume to the secondary volume. Depending on the volume size, this initial transfer could take several hours.
2. Check the progress of the initial synchronization:
 - a. In Unified Manager, launch System Manager for the local array.
 - b. In System Manager, view the status of the mirroring operation. When mirroring is complete, the status of the mirrored pair is "Optimal."
 3. **Optional:** You can reschedule or manually perform subsequent data transfers in System Manager. Only new and changed blocks are transferred from the primary volume to the secondary volume.



Because asynchronous replication is periodic, the system can consolidate the changed blocks and conserve network bandwidth. There is minimal impact on write throughput and write latency.

Requirements for using asynchronous mirroring

If you plan to use asynchronous mirroring, keep the following requirements in mind.

SANtricity Unified Manager

To enable and configure mirroring between two arrays, you must use the Unified Manager interface. Unified Manager is installed on a host system along with the Web Services Proxy.

- The Web Services Proxy service must be running.
- Unified Manager must be running on your local host through an HTTPS connection.
- Unified Manager must be showing valid SSL certificates for the storage array. You can accept a self-signed certificate or install your own security certificate using Unified Manager and navigating to **Certificate > Certificate Management**.

Storage arrays



Mirroring is not available on the EF600 storage array.

- You must have two storage arrays.
- Each storage array must have two controllers.
- The two storage arrays must be discovered in Unified Manager.
- Each controller in both the primary array and secondary array must have an Ethernet management port configured and must be connected to your network.
- The storage arrays have a minimum firmware version of 7.84. (They can each run different OS versions.)
- You must know the password for the local and remote storage arrays.

- You must have enough free capacity on the remote storage array to create a secondary volume equal to or greater than the primary volume that you want to mirror.
- Your local and remote storage arrays are connected through a Fibre Channel fabric or iSCSI interface.

Supported connections

Asynchronous mirroring can use either FC or iSCSI connections, or both for communication between local and remote storage systems. At the time of creating a mirror consistency group, the administrator can select either FC or iSCSI for that group if both are connected to the remote storage array. There is no failover from one channel type to the other.

Asynchronous mirroring uses the storage array's host-side I/O ports to convey mirrored data from the primary side to the secondary side.

• Mirroring through a Fibre Channel (FC) interface

Each controller of the storage array dedicates its highest numbered FC host port to mirroring operations.

If the controller has both base FC ports and host interface card (HIC) FC ports, the highest numbered port is on an HIC. Any host logged on to the dedicated port is logged out, and no host login requests are accepted. I/O requests on this port are accepted only from controllers that are participating in mirroring operations.

The dedicated mirroring ports must be attached to an FC fabric environment that supports the directory service and name service interfaces. In particular, FC-AL and point-to-point are not supported as connectivity options between the controllers that are participating in mirror relationships.

• Mirroring through an iSCSI interface

Unlike FC, iSCSI does not require a dedicated port. When asynchronous mirroring is used in iSCSI environments, it is not necessary to dedicate any of the storage array's front-end iSCSI ports for use with asynchronous mirroring; those ports are shared for both asynchronous mirror traffic and host-to-array I/O connections.

The controller maintains a list of remote storage systems with which the iSCSI initiator attempts to establish a session. The first port that successfully establishes an iSCSI connection is used for all subsequent communication with that remote storage array. If communication fails, a new session is attempted using all available ports.

iSCSI ports are configured at the array level on a port-by-port basis. Intercontroller communication for configuration messaging and data transfer uses the global settings, including settings for:

- VLAN: Both local and remote systems must have the same VLAN setting to communicate
- iSCSI listening port
- Jumbo frames
- Ethernet priority



The iSCSI intercontroller communication must use a host connect port and not the management Ethernet port.

Asynchronous mirroring uses the storage array's host-side I/O ports to convey mirrored data from the primary side to the secondary side. Because asynchronous mirroring is intended for higher-latency, lower-cost networks, iSCSI (and thus TCP/IP-based) connections are a good fit for it. When asynchronous mirroring is

used in iSCSI environments, it is not necessary to dedicate any of the array's front-end iSCSI ports for use with asynchronous mirroring; those ports are shared for both asynchronous mirror traffic and host-to-array I/O connections

Mirrored volume candidates

- RAID level, caching parameters, and segment size can be different on the primary and secondary volumes of an asynchronous mirrored pair.
- The secondary volume must be at least as large as the primary volume.
- A volume can participate in only one mirror relationship.

Reserved capacity


- A reserved capacity volume is required for a primary volume and for a secondary volume in a mirrored pair for logging write information to recover from controller resets and other temporary interruptions.
- Because both the primary volume and the secondary volume in a mirrored pair require additional reserved capacity, you must ensure that you have free capacity available on both storage arrays in the mirror relationship.

Drive Security feature

- If you are using secure-capable drives, the primary volume and the secondary volume must have compatible security settings. This restriction is not enforced; therefore, you must verify it yourself.
- If you are using secure-capable drives, the primary volume and the secondary volume should use the same drive type. This restriction is not enforced; therefore, you must verify it yourself.
- If you are using Data Assurance (DA), the primary volume and the secondary volume must have the same DA settings.

Asynchronous mirror status

The mirror status defines the state of mirror consistency groups and mirrored volume pairs.



This feature is not available on the EF600 storage system.

Status for mirror consistency groups

Status	Description
Synchronizing (initial sync)	<p>The progress of the initial data synchronization that has been completed between the mirrored volume pairs.</p> <p>During an initial synchronization, the volumes can transition to the following states: Degraded/Failed/Optimal/Unknown.</p>
Synchronizing (interval sync)	<p>The progress of the periodic data synchronization that has been completed between the mirrored volume pairs.</p>

Status	Description
System suspended	<p>Storage system-suspended synchronization of data on all mirrored pairs at the mirror consistency group level.</p> <p>At least one mirrored pair in the mirror consistency group is in a Stopped or Failed state.</p>
User suspended	<p>User-suspended synchronization of data on all mirrored pairs at the mirror consistency group level.</p> <p>This state helps to reduce any performance impact to the host application that might occur while any changed data on the local storage array is copied to the remote storage array.</p>
Paused	<p>Data synchronization process has temporarily paused due to an error accessing the remote storage array.</p>
Orphan	<p>An orphaned mirrored pair volume exists when a member volume in a consistency mirror group has been removed on one side of the consistency mirror group (either the primary side or secondary side) but not on the other side.</p> <p>Orphaned mirrored pair volumes are detected when inter-array communication is restored and the two sides of the mirror configuration reconcile mirror parameters.</p> <p>You can remove a mirrored pair to correct an orphaned mirrored pair state.</p>
Role change pending/in-progress	<p>A role change between the mirror consistency groups is pending or in progress.</p> <p>The role reversal change (to either a primary role or secondary role) affects all asynchronous mirrored pairs within the selected mirror consistency group.</p> <p>You can cancel a pending role change, but not an in-progress role change.</p>
Role conflict	<p>A role conflict occurred between mirror consistency groups due to a communication problem between the local storage array and the remote storage array during a role change operation.</p> <p>When the communication problem has been resolved, a Role Conflict occurs. Use the Recovery Guru to recover from this error.</p> <p>A forced promotion is not allowed when resolving a role conflict.</p>

Status for mirrored pairs

A mirrored pair's status indicates whether the data on the primary volume and on the secondary volume is synchronized.

Status	Description
Synchronizing	<p>The progress of initial or periodic data synchronization that has been completed between the mirrored pairs.</p> <p>There are two types of synchronization: initial synchronization and periodic synchronization. The initial synchronization progress is also displayed in the Long Running Operations dialog box.</p>
Optimal	The volumes in the mirrored pair are synchronized, which indicates that the connection between the storage arrays is operational and each volume is in the desired working condition.
Incomplete	<p>The asynchronous mirrored pair is incomplete on the remote storage array because the mirrored pair creation sequence was initiated on a storage array that is not supported with SANtricity System Manager and the mirrored pair has not been completed on the secondary.</p> <p>The mirrored pair creation process is complete when a volume is added to the mirror consistency group on the remote storage array. This volume becomes the secondary volume in the asynchronous mirrored pair.</p> <p>The mirrored pair completes automatically if the remote storage array is managed by SANtricity System Manager.</p>
Failed	The asynchronous mirroring operation is unable to operate normally due to a failure with the primary volumes, secondary volumes, or the mirror reserved capacity.
Orphan	<p>An orphaned mirrored pair volume exists when a member volume in a consistency mirror group has been removed on one side of the consistency mirror group (either the primary side or secondary side) but not on the other side.</p> <p>Orphaned mirrored pair volumes are detected when communication is restored between the two storage arrays and the two sides of the mirror configuration reconcile mirror parameters.</p> <p>You can remove a mirrored pair to correct an orphaned mirrored pair state.</p>
Stopped	The mirrored pair is in a Stopped state because the mirror consistency group is in a system-suspended state.

Volume ownership

You can change the preferred controller owner in a mirrored pair.



This feature is not available on the EF600 storage system.

If the primary volume of the mirrored pair is owned by controller A, then the secondary volume will also be owned by controller A of the remote storage array. Changing the primary volume's owner will automatically change the owner of the secondary volume to ensure that both volumes are owned by the same controller.

Current ownership changes on the primary side automatically propagate to corresponding current ownership changes on the secondary side.

For example, a primary volume is owned by controller A, and then you change the controller owner to controller B. In this case, the next remote write changes the controller owner of the secondary volume from controller A to B. Because controller ownership changes on the secondary side are controlled by the primary side, they do not require any special intervention by the storage administrator.

Controller resets

A controller reset causes a volume ownership change on the primary side from the preferred controller owner to the alternate controller in the storage array.

Sometimes a remote write is interrupted by a controller reset or a storage array power cycle before it can be written to the secondary volume. The controller does not need to perform a full synchronization of the mirrored pair in this case.

When a remote write has been interrupted during a controller reset, the new controller owner on the primary side reads information stored in a log file in the reserved capacity volume of the preferred controller owner. The new controller owner then copies the affected data blocks from the primary volume to the secondary volume, eliminating the need for a full synchronization of the mirrored volumes.

Role change of a mirror consistency group

You can change the role between mirrored pairs in a mirror consistency group. You can do this by demoting the primary mirror consistency group to the secondary role, or by promoting the secondary mirror consistency group to the primary role.



This feature is not available on the EF600 storage system.

Review the following information about the role change operation:

- The role change affects all mirrored pairs within the selected mirror consistency group.
- When a mirror consistency group is demoted to the secondary role, all the mirrored pairs within that mirror consistency group are also demoted to the secondary role and vice versa.
- When the primary mirror consistency group is demoted to the secondary role, hosts that have been assigned to the member volumes within that group no longer have write access to them.
- When a mirror consistency group is promoted to the primary role, any hosts that are accessing the member volumes within that group are now able to write to them.
- If the local storage array is unable to communicate with the remote storage array, you can force the role change on the local storage array.

Force role change

You can force a role change between mirror consistency groups when a communication problem between the local storage array and the remote storage array is preventing the promotion of the member volumes within the secondary mirror consistency group or the demotion of the member volumes within the primary mirror consistency group.

You can force the mirror consistency group on the secondary side to transition to the primary role. Then the recovery host is able to access the newly promoted member volumes within that mirror consistency group, and business operations can continue.

When is a forced promotion allowed and not allowed?

Forced promotion of a mirror consistency group is allowed only if all member volumes of the mirror consistency group have been synchronized and have consistent recovery points.

Forced promotion of a mirror consistency group is not allowed under the following conditions:

- Any of the member volumes of a mirror consistency group are in the process of an initial synchronization.
- Any of the member volumes of a mirror consistency group do not have a point-in-time image of the recovery point (for example, due to a full reserved capacity error).
- The mirror consistency group does not contain member volumes.
- The mirror consistency group is in the Failed, Role-Change-Pending, or Role-Change-In-Progress states, or if any of the associated member volumes or reserved capacity volumes are failed.

Mirror group role conflict

When a communication problem between the local and remote storage arrays has been resolved, a Mirror Group Role Conflict condition occurs. Use the Recovery Guru to recover from this error. A forced promotion is not allowed when resolving a dual-role conflict.

To avoid the Mirror Group Role Conflict condition and subsequent recovery steps, wait until the connection between the storage arrays is operational to force the role change.

Role change in-progress state

If two storage arrays in a mirroring configuration become disconnected, and the primary side of a mirror consistency group is force demoted to a secondary role, and the secondary side of a mirror consistency group is force promoted to a primary role, then when communication is restored, the mirror consistency groups on both storage arrays are placed in the Role-Change-In-Progress state.

The system will complete the role change process by transferring the change logs, re-synchronizing, setting the mirror consistency group state back to a normal operating state, and continuing with periodic synchronizations.

How to

Manage mirror consistency groups

Test communication for mirror consistency groups

You can test the communication link to diagnose possible communication problems between the local storage array and the remote storage array associated with a mirror consistency group.

Before you begin

The mirror consistency group that you want to test must exist on the local and remote storage arrays.

About this task

You can run four different tests:

- **Connectivity** — Verifies that the two controllers have a communication path. The connectivity test sends an inter-array message between the storage arrays, and then validates that the corresponding mirror

consistency group on the remote storage array exists. It also validates that the member volumes of the mirror consistency group on the remote storage array match the member volumes of the mirror consistency group on the local storage array.

- **Latency** — Sends a SCSI Test Unit command to each mirrored volume on the remote storage array associated with the mirror consistency group to test the minimum, average, and maximum latency.
- **Bandwidth** — Sends two inter-array messages to the remote storage array to test the minimum, average, and maximum bandwidth as well as the negotiated link speed of the port on the array performing the test.
- **Port connections** — Shows the port that is being used for mirroring on the local storage array and the port that is receiving the mirrored data on the remote storage array.

Steps

1. Select **Storage > Asynchronous Mirroring**.
2. Select the **Mirror Consistency Groups** tab, and then select the mirror consistency group that you want to test.
3. Select **Test Communication**.

The **Test Communication** dialog box appears.

4. Select one or more communication tests to perform between the local and remote storage arrays associated with the selected mirror consistency group, and then click **Test**.
5. Review the information displayed in the Results window.

Communication Test Status	Description
Normal with no errors	The mirror consistency group is communicating correctly.
Passed status (but not normal)	Check possible network or connection problems and retry the test.
Failed status	The reason for the failure is indicated. Refer to the Recovery Guru to correct the problem.
Port connection error	The reason may be that the local storage array is not connected or the remote storage array cannot be contacted. Refer to the Recovery Guru to correct the problem.

After you finish

After the communication test completes, this dialog box shows a Normal status, a Passed status, or a Failed status.

If the communication test returns a Failed status, the test continues to run after you close this dialog box until communication between the mirror consistency groups is restored.

Suspend or resume synchronization for mirror consistency group

You can suspend or resume the synchronization of data on all mirrored pairs within a mirror consistency group, which is more efficient than suspending or resuming synchronization on individual mirrored pairs.

About this task

Suspending and resuming synchronization on groups helps to reduce any performance impact to the host application, which might occur while any changed data on the local storage array is copied to the remote storage array.

The state of the mirror consistency group and its mirrored pairs stay suspended until you use the Resume option to resume synchronization activity.

Steps

1. Select **Storage › Asynchronous Mirroring**.
2. Select the **Mirror Consistency Groups** tab.

The Mirrored Consistency Group table appears and displays all the mirror consistency groups associated with the storage array.

3. Select the mirror consistency group that you want to suspend or resume, and then select either **More › Suspend** or **More › Resume**.

The system displays a confirmation.

4. Select **Yes** to confirm.

Results

System Manager performs the following actions:

- Either suspends or resumes data transfer between all mirrored pairs in a mirror consistency group without removing the mirror relationship.
- Logs any data that was written to the primary side of the mirror consistency group while the mirror group is suspended and writes the data automatically to the secondary side of the mirror consistency group when the mirror group is resumed. A full synchronization is not required.
- For a *suspended* mirror consistency groups, displays **user-suspended** in the Mirror Consistency Groups table.
- For a *resumed* mirror consistency group, data written to the primary volumes while the mirror consistency group was suspended is written to the secondary volumes immediately. Periodic synchronization resumes if an automatic synchronization interval has been set.

Change synchronization settings for a mirror consistency group

You can change the synchronization settings and warning thresholds that the mirror consistency group on the local storage array uses when data is initially synchronized or when data is re-synchronized during asynchronous mirroring operations.

About this task

Changing the synchronization settings affects the synchronization operations of all mirrored pairs within the mirror consistency group.

Steps

1. Select **Storage › Asynchronous Mirroring**.
2. Select the **Mirror Consistency Groups** tab.

The Mirrored Consistency Group table appears and displays all the mirror consistency groups associated with the storage array.

3. Select the mirror consistency group that you want to edit, and then select **More › Edit settings**.

The system displays the Edit Settings dialog box.

4. Edit the synchronization and alert settings as appropriate, and then click **Save**.

Field details

Field	Description
Synchronize the mirrored pairs...	<p>Specify whether you want to synchronize the mirrored pairs on the remote storage array either manually or automatically.</p> <ul style="list-style-type: none">• Manually – Select this option to manually synchronize the mirrored pairs on the remote storage array.• Automatically, every – Select this option to automatically synchronize the mirrored pairs on the remote storage array by specifying the time interval from the beginning of the previous update to the beginning of the next update. The default interval is 10 minutes.
Alert me...	<p>If you set the synchronization method to occur automatically, set the following alerts:</p> <ul style="list-style-type: none">• Synchronization – Set the length of time after which System Manager sends an alert that synchronization has not completed.• Remote recovery point – Set a time limit after which System Manager sends an alert indicating that the recovery point data on the remote storage array is older than your defined time limit. Define the time limit from the end of the previous update.• Reserved capacity threshold – Define a reserved capacity amount at which System Manager sends an alert that you are nearing the reserved capacity threshold. Define the threshold by percentage of the capacity remaining.

Results

System Manager changes the synchronization settings for every mirrored pair in the mirror consistency group.

Re-synchronize mirror consistency group manually

You can manually start re-synchronization for all mirrored pairs within a mirror consistency group.

Steps

1. Select **Storage › Synchronous Mirroring**.
2. Select the **Mirror Consistency Groups** tab.

The Mirror Consistency Group table appears and displays all the mirror consistency groups associated with the storage array.

3. Select the mirror consistency group that you want to re-synchronize, and then select **More › Manually resynchronize**.

The system displays a confirmation.

4. Select **Yes** to confirm.

Results

The system performs the following actions:

- Initiates re-synchronization of data on all of the mirrored pairs within the selected mirror consistency group.
- Updates modified data from the local storage array to the remote storage array.

View unsynchronized data amount between mirror consistency groups

You can view the amount of unsynchronized data between the mirror consistency groups on the local storage array and on the remote storage array. While the mirror consistency group is in an Unsynchronized status, no mirroring activity takes place.

About this task

You can perform this task when the selected mirror consistency group contains mirrored pairs and when synchronization is not currently in-progress.

Steps

1. Select **Storage › Asynchronous Mirroring**.
2. Select the **Mirror Consistency Groups** tab.

The Mirror Consistency Group table appears and displays all the mirror consistency groups associated with the storage array.

3. Click **More › View unsynchronized data amount**.

If unsynchronized data exists, the table values reflect this. The data amount column lists the unsynchronized data amount in MiB.

Update remote IP address

You can update the iSCSI IP address for your remote storage array to re-establish connection with the local storage array.

Before you begin

Both the local storage array and the remote storage array must be configured for asynchronous mirroring using an iSCSI connection.

Steps

1. Select **Storage › Asynchronous Mirroring**.
2. Select the **Mirror Consistency Groups** tab.

The Mirror Consistency Group table displays all the mirror consistency groups associated with the storage array.

3. Select the mirror consistency group that you want to update, and then select **More › Update remote IP address**.

The system displays the **Update Remote IP Address** dialog box.

4. Select **Update** to update the iSCSI IP address for your remote storage array.

Results

The system resets the IP address of the remote storage array to re-establish connection with the local storage array.

Change mirror consistency group role to primary or secondary

You can change the role between mirror consistency groups for administrative purposes or in the event of a disaster on the local storage array.

About this task

Mirror consistency groups created on the local storage array hold the primary role. Mirror consistency groups created on the remote storage array hold the secondary role. You can either demote the local mirror consistency group to a secondary role or promote the remote mirror consistency group to a primary role.

Steps

1. Select **Storage › Asynchronous Mirroring**.
2. Select the **Mirror Consistency Groups** tab.

The Mirror Consistency Group table appears and displays all the mirror consistency groups associated with the storage array.

3. Select the mirror consistency group for which you want to change the role, and then select **More › Change role to <Primary | Secondary>**.

The system displays a confirmation.

4. Confirm that you want to change the role of the mirror consistency group, and then click **Change Role**.



The system displays the **Cannot Contact Storage Array** dialog box when a role change is requested, but the remote storage array cannot be contacted. Click **Yes** to force the role change.

Results

System Manager performs the following actions:

- The Mirror Consistency Group table displays the status "pending" or "in-progress" next to the mirror consistency group undergoing the role change. You can cancel a Role Change operation that is pending by clicking the **Cancel** link found within the table cell.
- If the associated mirror consistency group can be contacted, the roles between the mirror consistency groups change. System Manager promotes the secondary mirror consistency group to a primary role or demotes the primary mirror consistency group to a secondary role (depending on your selection). The role change affects all mirrored pairs within the selected mirror consistency group.

Delete mirror consistency group

You can delete mirror consistency groups that are no longer needed on the local storage array and on the remote storage array.

Before you begin

All mirrored pairs must be removed from the mirror consistency group.

Steps

1. Select **Storage › Asynchronous Mirroring**.
2. Select the **Mirror Consistency Groups** tab.

The Mirror Consistency Group table appears and displays all the mirror consistency groups associated with the storage array.

3. Select the mirror consistency group that you want to delete, and then select **Uncommon Tasks › Delete**.

The system displays a confirmation.

4. Select **Yes** to delete the mirror consistency group.

Results

System Manager performs the following actions:

- Deletes the mirror consistency group on the local storage array first, and then deletes the mirror consistency group on the remote storage array.
- Removes the mirror consistency group from the Mirror Consistency Group table.

After you finish

Occasionally, there may be instances where the mirror consistency group is successfully deleted from the local storage array, but a communication error prevents the mirror consistency group from being deleted from the remote storage array. In this case, you must access the remote storage array to delete the corresponding mirror consistency group.

Manage asynchronous mirrored pairs

Remove asynchronous mirror relationship

You remove a mirrored pair to remove the mirror relationship from the primary volume on the local storage array and the secondary volume on the remote storage array.

About this task

Review the following information about orphaned mirrored pairs:

- An orphaned mirrored pair exists when a member volume in a consistency mirror group has been removed on one side (either the local storage array side or the remote storage array side) but not on the other side.
- Orphaned mirrored pairs are detected when inter-array communication is restored and the two sides of the mirror configuration reconcile mirror parameters.
- You can remove a mirrored pair to correct an orphaned mirrored pair state.

Steps

1. Select **Storage › Asynchronous Mirroring**.
2. Select the **Mirrored Pair** tab.

The Mirrored Pairs table appears and displays all the mirrored pairs associated with the storage array.

3. Select the mirrored pair that you want to remove, and then click **Remove**.
4. Confirm that you want to remove the mirrored pair, and then click **Remove**.

Results

SANtricity System Manager performs the following actions:

- Removes the mirror relationship from the mirror consistency group on the local storage array and on the remote storage array, and deletes the reserved capacity.
- Returns the primary volume and the secondary volume to host-accessible, non-mirrored volumes.
- Updates the Asynchronous Mirroring tile with the removal of the asynchronous mirrored pair.

Increase reserved capacity

You can increase reserved capacity, which is the physically allocated capacity used for any copy service operation on a storage object. For snapshot operations, it is typically 40 percent of the base volume; for asynchronous mirroring operations, it is typically 20 percent of the base volume. Typically, you increase reserved capacity when you receive a warning that the storage object's reserved capacity is becoming full.

Before you begin

- The volume in the pool or volume group must have an Optimal status and must not be in any state of modification.
- Free capacity must exist in the pool or volume group that you want to use to increase capacity.

If no free capacity exists on any pool or volume group, you can add unassigned capacity in the form of unused drives to a pool or volume group.

About this task

You can increase reserved capacity only in increments of 8 GiB for the following storage objects:

- Snapshot group
- Snapshot volume
- Consistency group member volume
- Mirrored pair volume

Use a high percentage if you believe the primary volume will undergo many changes or if the lifespan of a particular copy service operation will be very long.



You cannot increase reserved capacity for a snapshot volume that is read-only. Only snapshot volumes that are read-write require reserved capacity.

Steps

1. Select **Storage › Pools & Volume Groups**.

2. Select the **Reserved Capacity** tab.
3. Select the storage object for which you want to increase reserved capacity, and then click **Increase Capacity**.

The **Increase Reserved Capacity** dialog box appears.

4. Use the spinner box to adjust the capacity percentage.

If free capacity does not exist on the pool or volume group that contains the storage object you selected, and the storage array has Unassigned Capacity, you can create a new pool or volume group. You can then retry this operation using the new free capacity on that pool or volume group.

5. Click **Increase**.

Results

System Manager performs the following actions:

- Increases the reserved capacity for the storage object.
- Displays the newly-added reserved capacity.

Change the reserved capacity settings for a mirrored pair volume

You can change the settings for a mirrored pair volume to adjust the percentage point at which System Manager sends an alert notification when the reserved capacity for a mirrored pair volume is nearing full.


Steps

1. Select **Storage > Pools & Volume Groups**.
2. Select the **Reserved Capacity** tab.
3. Select the mirrored pair volume that you want to edit, and then click **View/Edit Settings**.

The **Mirrored Pair Volume Reserved Capacity Settings** dialog box appears.

4. Change the reserved capacity settings for the mirrored pair volume as appropriate.

Field details

Setting	Description
Alert me when...	<p>Use the spinner box to adjust the percentage point at which System Manager sends an alert notification when the reserved capacity for a mirrored pair is nearing full.</p> <p>When the reserved capacity for the mirrored pair exceeds the specified threshold, System Manager sends an alert, allowing you time to increase reserved capacity.</p> <div><p>Changing the Alert setting for one mirrored pair changes the Alert setting for all mirrored pairs that belong to the same mirror consistency group.</p></div>

5. Click **Save** to apply your changes.

Complete mirrored pair for primary volumes created on legacy system

If you created a primary volume on a legacy storage array that cannot be managed by SANtricity System Manager, you can create the secondary volume on this array with SANtricity System Manager.

About this task

You can perform asynchronous mirroring between legacy arrays that use a different interface and newer arrays that can be managed by SANtricity System Manager.

- If you are mirroring between two storage arrays that use SANtricity System Manager, you can skip this task because you already completed the mirrored pair in the mirrored pair creation sequence.
- Perform this task on the remote storage array.

Steps

1. Select **Storage > Asynchronous Mirroring**.
2. Select the **Mirrored Pair** tab.

The Mirrored Pairs table appears and displays all the mirrored pairs associated with the storage array.

3. Find the mirrored pair volume with a status of Incomplete, and then click the **Complete mirrored pair** link displayed in the mirrored pair column.
4. Choose whether you want to complete the mirrored pair creation sequence automatically or manually by selecting one of the following radio buttons:
 - **Automatic** — Create new secondary volume.

Accept the default settings for the remote side of the mirrored pair by selecting an existing pool or volume group where you want to create the secondary volume. Use this recommended option to allocate the reserved capacity for the secondary volume with the default settings.

- **Manual** — Select an existing volume.

Define your own parameters for the secondary volume.

- i. Click **Next** to select the secondary volume.
- ii. Select an existing volume that you want to use as the secondary volume and then click **Next** to allocate the reserved capacity.
- iii. Allocate the reserved capacity. Do one of the following:

- Accept the default settings.

The default setting for reserved capacity is 20% of the capacity of the base volume, and usually this capacity is sufficient.

- Allocate your own reserved capacity settings to meet your data storage needs related to asynchronous mirroring.

The capacity needed varies, depending on the frequency and size of I/O writes to the primary volume and how long you need to keep the capacity. In general, choose a larger capacity for reserved capacity if one or both of these conditions exist:

- You intend to keep the mirrored pair for a long period of time.
- A large percentage of data blocks will change on the primary volume due to heavy I/O activity. Use historical performance data or other operating system utilities to help you determine typical I/O activity to the primary volume.

5. Select **Complete**.

Results

SANtricity System Manager performs the following actions:

- Creates the secondary volume on the remote storage array and allocates reserved capacity for the remote side of the mirrored pair.
- Begins initial synchronization between the local storage array and the remote storage array.
- If the volume being mirrored is a thin volume, only the allocated blocks are transferred to the secondary volume during the initial synchronization. This transfer reduces the amount of data that must be transferred to complete the initial synchronization.
- Creates the reserved capacity for the mirrored pair on the local storage array and on the remote storage array.

Deactivate asynchronous mirroring

You can deactivate asynchronous mirroring on the local and remote storage arrays to re-establish normal use of dedicated ports on the storage arrays.

Before you begin

- You must have deleted all mirror relationships. Verify that all mirror consistency groups and mirrored pairs have been deleted from the local and remote storage arrays.
- The local storage array and the remote storage array must be connected through a Fibre Channel fabric or iSCSI interface.

About this task

When you deactivate asynchronous mirroring, no mirror activity can occur on the local and remote storage arrays.

Steps

1. Select **Storage › Asynchronous Mirroring**.
2. Select **Uncommon Tasks › Deactivate**.

The system displays a confirmation.

3. Select **Yes** to confirm.

Results

- The controller's HBA host channels that were dedicated for asynchronous mirroring communication can now accept host read and write requests.
- None of the volumes in this storage array are able to participate in mirror relationships as either primary volumes or secondary volumes.

FAQs

How does asynchronous mirroring differ from synchronous mirroring?

The Asynchronous Mirroring feature differs from the Synchronous Mirroring feature in one essential way: it captures the state of the source volume at a particular point in time and copies just the data that has changed since the last image capture.

With synchronous mirroring, the state of the primary volume is not captured at some point in time, but rather reflects all changes that were made on the primary volume to the secondary volume. The secondary volume is identical to the primary volume at every moment because, with this type of mirror, each time a write is done to the primary volume, a write is done to the secondary volume. The host does not receive an acknowledgment that the write was successful until the secondary volume is successfully updated with the changes that were made on the primary volume.

With asynchronous mirroring, the remote storage array is not fully synchronized with the local storage array, so if the application needs to transition to the remote storage array due to a loss of the local storage array, some transactions could be lost.

Comparison between mirroring features

Asynchronous Mirroring	Synchronous Mirroring
Replication method	
<ul style="list-style-type: none">• Point-in-Time <p>Mirroring is done on demand or automatically according to a user-defined schedule. Schedules can be defined at the granularity of minutes. The minimum time between syncs is 10 minutes.</p>	<ul style="list-style-type: none">• Continuous <p>Mirroring is automatically executed continuously, copying data from every host write.</p>
Reserved capacity	
<ul style="list-style-type: none">• Multiple <p>A reserved capacity volume is required for each mirrored pair.</p>	<ul style="list-style-type: none">• Single <p>Single reserved capacity volume is required for all mirrored volumes.</p>
Communication	
<ul style="list-style-type: none">• iSCSI and Fibre Channel <p>Supports iSCSI and Fibre Channel interfaces between storage arrays.</p>	<ul style="list-style-type: none">• Fibre Channel <p>Supports only Fibre Channel interfaces between storage arrays.</p>
Distance	
<ul style="list-style-type: none">• Unlimited <p>Support for virtually unlimited distances between the local storage array and the remote storage array, with the distance typically limited only by the capabilities of the network and the channel extension technology.</p>	<ul style="list-style-type: none">• Restricted <p>Typically must be within about 10 km (6.2 miles), of the local storage array to meet the latency and application performance requirements.</p>

Why can't I access my chosen mirroring feature?

Mirroring is configured in the Unified Manager interface.

To enable and configure mirroring between two arrays, verify the following:

- The Web Services Proxy service must be running. (Unified Manager is installed on a host system along with the Web Services Proxy.)
- Unified Manager must be running on your local host through an HTTPS connection.

- The two storage arrays you want to use for mirroring must be discovered in Unified Manager.
- Unified Manager must have valid SSL certificates for the storage arrays. You can accept a self-signed certificate or install CA-signed certificates from Unified Manager.

For more information about requirements, see the online help for Unified Manager.



Mirroring is not available on the EF600 storage array.

What do I need to know before creating a mirror consistency group?

Follow these guidelines before you create a mirror consistency group.

You create a consistency group in Unified Manager in the Create Mirrored Pairs wizard.

Meet the following requirements for Unified Manager:

- The Web Services Proxy service must be running.
- Unified Manager must be running on your local host through an HTTPS connection.
- Unified Manager must be showing valid SSL certificates for the storage array. You can accept a self-signed certificate or install your own security certificate using Unified Manager and navigating to **Certificate > Certificate Management**.

Also be sure to meet the following requirements for storage arrays:

- The two storage arrays must be discovered in Unified Manager.
- Each storage array must have two controllers.
- Each controller in both the primary array and secondary array must have an Ethernet management port configured and must be connected to your network.
- The storage arrays have a minimum firmware version of 7.84. (They can each run different OS versions.)
- You must know the password for the local and remote storage arrays.
- Your local and remote storage arrays are connected through a Fibre Channel fabric or iSCSI interface.



This feature is not available on the EF600 storage system.

Asynchronous mirroring - What do I need to know before creating a mirrored pair?

You configure mirrored pairs in the Unified Manager interface, and then manage the pairs in System Manager.

Before creating a mirrored pair, follow these guidelines.

- You must have two storage arrays.
- Each storage array must have two controllers.
- Each controller in both the primary array and secondary array must have an Ethernet management port configured and must be connected to your network.
- Your local and remote storage arrays are connected through a Fibre Channel fabric or iSCSI interface.
- The storage arrays have a minimum firmware version of 7.84. (They can each run different OS versions.)

- You must know the password for the local and remote storage arrays.
- You must have enough free capacity on the remote storage array to create a secondary volume equal to or greater than the primary volume that you want to mirror.
- You have installed the Web Services Proxy and Unified Manager. Mirrored pairs are configured in the Unified Manager interface.
- The two storage arrays are discovered in Unified Manager.
- Your storage array must contain at least one mirror consistency group. You create a consistency group in Unified Manager in the Create Mirrored Pairs wizard.

What do I need to know before increasing my reserved capacity on a mirrored pair volume?

Typically, you should increase reserved capacity when you receive a warning that the reserved capacity for a mirrored pair is becoming full. You can increase reserved capacity only in increments of 8 GiB.

For asynchronous mirroring operations, reserved capacity is typically 20 percent of the base volume. Choose a larger capacity for reserved capacity if one or both of these conditions exist:

- You intend to keep the mirrored pair for a long period of time.
- A large percentage of data blocks will change on the primary volume due to heavy I/O activity. Use historical performance data or other operating system utilities to help you determine typical I/O activity to the primary volume.

You can increase the reserved capacity for a mirrored pair by performing one of these actions:

- Adjust the capacity percentage for a mirrored pair volume by selecting **Storage > Pools and Volume Groups** and then clicking the **Reserved Capacity** tab.
- Create a new volume using free capacity that is available on a pool or volume group.

If no free capacity exists on any pool or volume group, you can add unconfigured capacity in the form of unused drives to a pool or volume group.

Why can't I increase reserved capacity with my requested amount?

You can increase reserved capacity only in increments of 4 GiB.

Review the following guidelines:

- You must have sufficient free capacity in the pool or volume group so it can be expanded if necessary.
If no free capacity exists on any pool or volume group, you can add unassigned capacity in the form of unused drives to a pool or volume group.
- The volume in the pool or volume group must have an Optimal status and must not be in any state of modification.
- Free capacity must exist in the pool or volume group that you want to use to increase capacity.

For asynchronous mirroring operations reserved capacity is typically 20 percent of the base volume. Use a higher percentage if you believe the base volume will undergo many changes or if the estimated life expectancy of a storage object's copy service operation will be very long.

Why would I change this percentage?

Reserved capacity is typically 40 percent of the base volume for snapshot operations and 20 percent of the base volume for asynchronous mirroring operations. Usually this capacity is sufficient. The capacity needed varies, depending on the frequency and size of I/O writes to the base volume and how long you intend to use the storage object's copy service operation.

In general, choose a larger percentage for reserved capacity if one or both of these conditions exist:

- If the lifespan of a particular storage object's copy service operation will be very long.
- If a large percentage of data blocks will change on the base volume due to heavy I/O activity. Use historical performance data or other operating system utilities to help you determine typical I/O activity to the base volume.

Why do I see more than one reserved capacity candidate?

If there is more than one volume in a pool or volume group that meets the capacity percentage amount you selected for the storage object, you will see multiple candidates.

You can refresh the list of recommended candidates by changing the percentage of physical drive space that you want to reserve on the base volume for copy service operations. The best candidates are displayed based on your selection.

Why do I see Not Available values displayed in the table?

The table lists Not Available values when the data located on the remote storage array is not available to be displayed.

To display the remote storage array data, launch System Manager from Unified Manager.

Why don't I see all of my pools and volume groups?

When you create a secondary volume for the asynchronous mirrored pair, the system displays a list of all the eligible pools and volume groups for that asynchronous mirrored pair. Any pool or volume group that is not eligible to be used does not display in that list.

Pools or volume groups may not be eligible for any of the following reasons.

- The security capabilities of a pool or volume group do not match.
- A pool or volume group is in a non-optimal state.
- The capacity of a pool or volume group is too small.

Asynchronous mirroring - Why don't I see all my volumes?

When you are selecting a primary volume for a mirrored pair, a list shows all the eligible volumes.

Any volumes that are not eligible to be used do not display in that list. Volumes may not be eligible for any of the following reasons:

- The volume is not optimal.
- The volume is already participating in a mirroring relationship.
- For thin volumes, auto-expansion must be enabled.

Asynchronous mirroring - Why don't I see all the volumes on the remote storage array?

When you are selecting a secondary volume on the remote storage array, a list shows all the eligible volumes for that mirrored pair.

Any volumes that are not eligible to be used, do not display in that list. Volumes might not be eligible for any of the following reasons:

- The volume is not optimal.
- The volume is already participating in a mirroring relationship.
- The thin volume attributes between the primary volume and the secondary volume do not match.
- If you are using Data Assurance (DA), the primary volume and the secondary volume must have the same DA settings.
 - If the primary volume is DA enabled, the secondary volume must be DA enabled.
 - If the primary volume is not DA enabled, the secondary volume must not be DA enabled.

Why would I update my remote storage array's IP address?

You update your remote storage array's IP address when the IP address of an iSCSI port changes and the local storage array is unable to communicate with the remote storage array.

More about the remote IP address

When establishing an asynchronous mirroring relationship with an iSCSI connection, both the local and the remote storage arrays store a record of the IP address of the remote storage array in the asynchronous mirroring configuration. If the IP address of an iSCSI port changes, the remote storage array that is attempting to use that port encounters a communication error.

The storage array with the changed IP address sends a message to each remote storage array associated with the mirror consistency groups that are configured to mirror over an iSCSI connection. Storage arrays that receive this message automatically update their remote-target IP address.

If the storage array with the changed IP address is unable to send its inter-array message to a remote storage array, the system sends you an alert of the connectivity issue. Use the **Update Remote IP Address** option to re-establish connection with the local storage array.

Synchronous mirroring

Concepts

How synchronous mirroring works

Synchronous mirroring replicates data volumes in real time to ensure continuous availability.



Mirroring is not available on the EF600 storage array.

Synchronous mirroring achieves a recovery point objective (RPO) of zero lost data by having a copy of important data available if a disaster happens on one of the two storage arrays. The copy is identical to production data at every moment because each time a write is done to the primary volume, a write is done to the secondary volume. The host does not receive an acknowledgment that the write was successful until the secondary volume is successfully updated with the changes that were made on the primary volume.


This type of mirroring is ideal for business continuity purposes such as disaster recovery.

Synchronous mirroring relationship

A synchronous mirroring relationship consists of a primary volume and a secondary volume on separate storage arrays. The storage array containing the primary volume is usually located at the primary site and serves the active hosts. The storage array containing the secondary volume is usually located at a secondary site and holds a replica of the data. The secondary volume is used if the primary volume's storage array is unavailable because of, for example, a complete power outage, a fire, or a hardware failure at the primary site.

Synchronous mirroring session

The synchronous mirroring configuration process involves configuring volumes into pairs. After you create a mirrored pair, which consists of a primary volume on one storage array and a secondary volume on another storage array, you can start synchronous mirroring. The steps in synchronous mirroring are depicted below.



1. A write comes in from the host.
2. The write is committed to the primary volume, propagated to the remote system, and then committed to the secondary volume.
3. The primary volume's storage array sends an I/O completion message to the host system *after* both write operations have been successfully completed.

Reserved capacity is used to log information about the incoming write request from a host.

When the current controller owner of the primary volume receives a write request from a host, the controller first logs information about the write to the primary volume's reserved capacity. It then writes the data to the primary volume. Next, the controller initiates a remote write operation to copy the affected data blocks to the secondary volume at the remote storage array.

Because the host application must wait for the write to occur on the local storage array and across the network on the remote storage array, a very fast connection between the local storage array and remote storage array is required to maintain the mirror relationship without overly reducing local I/O performance.

Disaster recovery

Synchronous mirroring maintains a copy of data that is physically distant from the site where the data resides. If a disaster occurs at the primary site, such as a power outage or a flood, the data can be quickly accessed from the secondary site.

The secondary volume is unavailable to host applications while the synchronous mirroring operation is in progress, so, in the event of a disaster at the local storage array, you can fail over to the remote storage array. To fail over, promote the secondary volume to the primary role. Then the recovery host is able to access the newly promoted volume, and business operations can continue.

Synchronization settings

When you create a mirrored pair, you also define the synchronization priority and resynchronization policy that the mirrored pair uses to complete the resynchronization operation after a communication interruption.

If the communication link between the two storage arrays stops working, hosts continue to receive acknowledgements from the local storage array, preventing an access loss. When the communication link is working again, any unreplicated data can be automatically or manually resynced to the remote storage array.

Whether data is resynchronized automatically depends on the mirrored pair's resynchronization policy. An automatic resynchronization policy allows the mirrored pair to resynchronize automatically when the link is working again. A manual resynchronization policy requires you to manually resume synchronization after a communication problem. Manual resynchronization is the recommended policy.

You can edit the synchronization settings for a mirrored pair only on the storage array that contains the primary volume.

Unsynchronized data

The primary and secondary volumes become unsynchronized when the primary volume's storage array is unable to write data to the secondary volume. This can be caused by the following issues:

- Network problems between the local and remote storage arrays
- A failed secondary volume
- Synchronization being manually suspended on the mirrored pair

Orphaned mirrored pair

An orphaned mirrored pair volume exists when a member volume has been removed on one side (either the primary side or secondary side) but not on the other side.

Orphaned mirrored pair volumes are detected when inter-array communication is restored and the two sides of the mirror configuration reconcile mirror parameters.

You can remove a mirrored pair to correct an orphaned mirrored pair state.

Configuration and management

To enable and configure mirroring between two arrays, you must use the Unified Manager interface. Once mirroring is enabled, you can manage mirrored pairs and synchronization settings in System Manager.

Synchronous mirroring terminology

Learn how the synchronous mirroring terms apply to your storage array.

Term	Description
Local storage array	<p>The local storage array is the storage array that you are acting upon.</p> <p>When you see Primary in the Local Role column, it indicates that the storage array contains the volume that holds the primary role in the mirror relationship. When you see Secondary in the Local Role column, it indicates that the storage array contains the volume that holds the secondary role in the mirror relationship.</p>
Mirrored pair	A mirrored pair is comprised of two volumes, a primary volume and a secondary volume.
Primary volume	The primary volume of a mirrored pair is the source volume to be mirrored.
Recovery point objective (RPO)	Recovery Point Objective (RPO) represents an objective that indicates the difference considered acceptable between the primary volume and secondary volume in a mirrored pair. An RPO of zero indicates that no difference between the primary volume and secondary volume can be tolerated. An RPO greater than zero indicates that the secondary volume is less current or lags behind the primary volume.
Remote storage array	The remote storage array is usually designated as the secondary site, which usually holds a replica of the data in a mirroring configuration.
Reserved capacity	Reserved capacity is the physical allocated capacity that is used for any copy service operation and storage object. It is not directly readable by the host.
Role change	Role change is assigning the primary role to the secondary volume and vice versa.
Secondary volume	The secondary volume of a mirrored pair is usually located at a secondary site and holds a replica of the data.
Synchronization	Synchronization occurs at initial synchronization between the local storage array and the remote storage array. Synchronization also occurs when the primary and secondary volumes become unsynchronized after a communication interruption. When the communication link is working again, any unreplicated data is synchronized to the secondary volume's storage array.

Workflow for mirroring a volume synchronously

You configure synchronous mirroring using the following workflow.



This feature is not available on the EF600 storage system.

1. Perform the initial configuration in Unified Manager:
 - a. Select a local storage array as the source for the data transfer.
 - b. Select a primary volume from the local storage array.

- c. Select a remote storage array as the destination for the data transfer, and then select a secondary volume.
 - d. Select synchronization and resynchronization priorities.
 - e. Begin the initial data transfer from the primary volume to the secondary volume. Depending on the volume size, this initial transfer could take several hours.
2. Check the progress of the initial synchronization:
 - a. In Unified Manager, launch System Manager for the local array.
 - b. In System Manager, view the status of the mirroring operation. When mirroring is complete, the status of the mirrored pair is "Optimal." The two arrays attempt to stay synchronized through normal operations. Only new and changed blocks are transferred from the primary volume to the secondary volume.
3. **Optional:** You can change synchronization settings in System Manager.



Because synchronous replication is continuous, the replication link between the two sites must provide sufficient bandwidth capabilities.

Requirements for using synchronous mirroring

If you plan to use synchronous mirroring, keep the following requirements in mind.

SANtricity Unified Manager

To enable and configure mirroring between two arrays, you must use the Unified Manager interface. Unified Manager is installed on a host system along with the Web Services Proxy.

- The Web Services Proxy service must be running.
- Unified Manager must be running on your local host through an HTTPS connection.
- Unified Manager must be showing valid SSL certificates for the storage array. You can accept a self-signed certificate or install your own security certificate using Unified Manager and navigating to **Certificate > Certificate Management**.

Storage arrays



Mirroring is not available on the EF600 storage array.

- You must have two storage arrays.
- Each storage array must have two controllers.
- The two storage arrays must be discovered in Unified Manager.
- Each controller in both the primary array and secondary array must have an Ethernet management port configured and must be connected to your network.
- The storage arrays have a minimum firmware version of 7.84. (They can each run different OS versions.)
- You must know the password for the local and remote storage arrays.
- You must have enough free capacity on the remote storage array to create a secondary volume equal to or greater than the primary volume that you want to mirror.
- Your local and remote storage arrays are connected through a Fibre Channel fabric.

Supported connections

Communication for synchronous mirroring is supported only on controllers with Fibre Channel (FC) host ports.

Synchronous mirroring uses the highest numbered host port on each controller on both the local storage array and the remote storage array. Controller host bus adapter (HBA) host port 4 is typically reserved for mirror data transmission.

Mirrored volume candidates

- RAID level, caching parameters, and segment size can be different on the primary and secondary volumes of a synchronous mirrored pair.
- The primary and secondary volumes in a synchronous mirrored pair must be standard volumes. They cannot be thin volumes or snapshot volumes.
- The secondary volume must be at least as large as the primary volume.
- Only the primary volume may have snapshots associated with it and/or be the source or target volume in a volume copy operation.
- A volume can participate in only one mirror relationship.
- There are limits to the number of volumes that are supported on a given storage array. Make sure that the number of configured volumes on your storage array is less than the supported limit. When synchronous mirroring is active, the two reserved capacity volumes that are created count against the volume limit.

Reserved capacity

- Reserved capacity is required for a primary volume and for a secondary volume for logging write information to recover from controller resets and other temporary interruptions.
- The reserved capacity volumes are created automatically when synchronous mirroring is activated. Because both the primary volume and the secondary volume in a mirrored pair require reserved capacity, you must ensure that you have enough free capacity available on both storage arrays that are participating in the synchronous mirror relationship.

Drive Security feature

- If you are using secure-capable drives, the primary volume and the secondary volume must have compatible security settings. This restriction is not enforced; therefore, you must verify it yourself.
- If you are using secure-capable drives, the primary volume and the secondary volume should use the same drive type. This restriction is not enforced; therefore, you must verify it yourself.
 - If the primary volume uses Full Disk Encryption (FDE) drives, the secondary volume should use FDE drives.
 - If the primary volume uses Federal Information Processing Standards 140-2 (FIPS) validated drives, the secondary volume should use FIPS 140-2 validated drives.
- If you are using Data Assurance (DA), the primary volume and the secondary volume must have the same DA settings.

Synchronous mirroring status

A synchronous mirrored pair's status indicates whether the data on the primary volume and on the secondary volume is synchronized. A mirror status is independent of the component status of the volumes in the mirrored pair.



This feature is not available on the EF600 storage system.

Synchronous mirrored pairs can have one of the following statuses:

- **Optimal**

Indicates that the volumes in the mirrored pair are synchronized, which means that the fabric connection between the storage arrays is operational and each volume is in the desired working condition.

- **Synchronizing**

Shows the progress of the data synchronization between the mirrored pairs. This status will also be shown during the initial synchronization.

After a communication link interruption, only the blocks of data that have changed on the primary volume during the link interruption are copied to the secondary volume.

- **Unsynchronized**

Indicates that the primary volume's storage array is unable to write incoming data to the remote array. The local host can continue to write to the primary volume, but remote writes do not take place. Different conditions can prevent the primary volume's storage array from writing incoming data to the secondary volume, such as:

- The secondary volume is not accessible.
- The remote storage array is not accessible.
- The fabric connection between the storage arrays is not accessible.
- The secondary volume cannot be updated with a new World Wide Identifier (WWID).

- **Suspended**

Indicates that the synchronous mirroring operation has been suspended by the user. When a mirrored pair is suspended, no attempt is made to contact the secondary volume. Any writes to the primary volume are persistently logged in the mirror reserved capacity volumes.

- **Failed**

Indicates that the synchronous mirroring operation is unable to operate normally due to a failure with the primary volume, secondary volume, or the mirror reserved capacity.

Volume ownership

You can change the preferred controller owner in a mirrored pair.



This feature is not available on the EF600 storage system.

If the primary volume of the mirrored pair is owned by controller A, then the secondary volume will also be owned by controller A of the remote storage array. Changing the primary volume's owner will automatically change the owner of the secondary volume to ensure that both volumes are owned by the same controller. Current ownership changes on the primary side automatically propagate to corresponding current ownership changes on the secondary side.

For example, a primary volume is owned by controller A, and then you change the controller owner to

controller B. In this case, the next remote write changes the controller owner of the secondary volume from controller A to B. Because controller ownership changes on the secondary side are controlled by the primary side, they do not require any special intervention by the storage administrator.

Controller resets

A controller reset causes a volume ownership change on the primary side from the preferred controller owner to the alternate controller in the storage array.

Sometimes a remote write is interrupted by a controller reset or a storage array power cycle before it can be written to the secondary volume. The controller does not need to perform a full synchronization of the mirrored pair in this case.

When a remote write has been interrupted during a controller reset, the new controller owner on the primary side reads information stored in a log file in the reserved capacity volume of the preferred controller owner. The new controller owner then copies the affected data blocks from the primary volume to the secondary volume, eliminating the need for a full synchronization of the mirrored volumes.

Role change between volumes in a mirrored pair

You can change the role between volumes in a mirrored pair. You can do this by demoting the primary volume to the secondary role or promoting the secondary volume to the primary role.



This feature is not available on the EF600 storage system.

Review the following information about the role change operation:

- When a primary volume is demoted to the secondary role, the secondary volume in that mirrored pair is promoted to the primary role and vice versa.
- When the primary volume is demoted to the secondary role, hosts that have been assigned to that volume no longer have write access to it.
- When the secondary volume is promoted to the primary role, any hosts that are accessing that volume are now able to write to it.
- If the local storage array is unable to communicate with the remote storage array, you can force the role change on the local storage array.

Force role change

You can force a role change between volumes in a mirrored pair when a communication problem between the local storage array and the remote storage array is preventing the promotion of the secondary volume or the demotion of the primary volume.

You can force the volume on the secondary side to transition to the primary role. Then the recovery host can access the newly promoted volume, and business operations can continue.



When the remote storage array has recovered and any communication problems have been resolved, a Synchronous Mirroring - Primary Volume Conflict condition occurs. The recovery steps include resynchronizing the volumes. Use the Recovery Guru to recover from this error.

When is a forced promotion allowed and not allowed?

Forced promotion of a volume in a mirrored pair is not allowed under the following conditions:

- Any of the volumes in a mirrored pair are in the process of an initial synchronization.
- The mirrored pair is in the Failed, Role-Change-Pending, or Role-Change-In-Progress states or if any of the associated reserved capacity volumes are failed.

Role change in-progress state

If two storage arrays in a mirroring configuration become disconnected, and the primary volume of a mirrored pair is force demoted to a secondary role, and the secondary volume of a mirrored pair is force promoted to a primary role, then when communication is restored, the volumes on both storage arrays are placed in the Role-Change-In-Progress state.

The system will complete the role change process by transferring the change logs, re-synchronizing, setting the mirrored pair state back to a normal operating state, and continuing with synchronizations.

How to

Manage synchronous mirrored pairs

Test communication for synchronous mirroring

You can test the communication between a local storage array and a remote storage array to diagnose possible communication problems for a mirrored pair that is participating in synchronous mirroring.

About this task

Two different tests are run:

- **Communication** — Verifies that the two storage arrays have a communication path. The communication test validates that the local storage array can communicate with the remote storage array and that the secondary volume associated with the mirrored pair exists on the remote storage array.
- **Latency** — Sends a SCSI test unit command to the secondary volume on the remote storage array associated with the mirrored pair to test the minimum, average, and maximum latency.

Steps

1. Select **Storage > Synchronous Mirroring**.
2. Select the mirrored pair that you want to test, and then select **Test Communication**.
3. Review the information displayed in the Results window, and, if necessary, follow the corrective action indicated.



If the communication test fails, the test continues to run after you close this dialog until communication between the mirrored pair is restored.

Suspend and resume synchronization for a mirrored pair

You can use the Suspend option and Resume option to control when to synchronize the data on the primary volume and the secondary volume in a mirrored pair.

About this task

If a mirrored pair is manually suspended, the mirrored pair will not synchronize until it is manually resumed.

Steps

1. Select **Storage › Synchronous Mirroring**.
2. Select the mirrored pair that you want to suspend or resume, and then select either **More › Suspend** or **More › Resume**.

The system displays a confirmation.

3. Select **Yes** to confirm.

Results

System Manager performs the following actions:

- Either suspends or resumes data transfer between the mirrored pair without removing the mirror relationship.
- For a *suspended* mirrored pair:
 - Displays **Suspended** in the Mirrored Pair table.
 - Logs any data that was written to the primary volume of the mirrored pair while synchronization is suspended.
- For a *resumed* mirrored pair, writes the data automatically to the secondary volume of the mirrored pair when synchronization is resumed. A full synchronization is not required.

Change role between volumes in a mirrored pair

You can perform a role reversal between the two volumes in a mirrored pair that are participating in synchronous mirroring. You change the role between volumes in a mirrored pair for administrative purposes or in the event of a disaster on the local storage array.

About this task

You can either demote the primary volume to the secondary role or promote the secondary volume to the primary role. Any hosts that are accessing the primary volume have read/write access to the volume. When the primary volume becomes a secondary volume, only remote writes initiated by the primary controller are written to the volume.

Steps

1. Select **Storage › Synchronous Mirroring**.
2. Select the mirrored pair that contains the volumes for which you want to change the role, and then select **More › Change role**.

The system displays a confirmation.

3. Confirm that you want to change the role of the volumes, and then select **Change Role**.



If the local storage array cannot communicate with the remote storage array, the system displays the **Cannot Contact Storage Array** dialog box when a role change is requested, but the remote storage array cannot be contacted. Click **Yes** to force the role change.

Results

System Manager performs the following action:

- If the associated volume in the mirrored pair can be contacted, the roles between the volumes change. System Manager promotes the secondary volume in the mirrored pair to the primary role or demotes the primary volume in the mirrored pair to the secondary role (depending on your selection).

Change synchronization settings for a mirrored pair

You can change the synchronization priority and resynchronization policy that the mirrored pair uses to complete the resynchronization operation after a communication interruption.

About this task

You can edit the synchronization settings for a mirrored pair only on the storage array that contains the primary volume.

Steps

1. Select **Storage > Synchronous Mirroring**.
2. Select the mirrored pair that you want to edit, and then select **More > Edit settings**.

The system displays the View/Edit Settings dialog box.

3. Use the slider bar to edit the synchronization priority.

The synchronization priority determines how much of the system resources are used to complete the resynchronization operation after a communication interruption as compared to service I/O requests.

More about synchronization rates

There are five synchronization priority rates:

- Lowest
- Low
- Medium
- High
- Highest

If the synchronization priority is set to the lowest rate, I/O activity is prioritized, and the resynchronization operation takes longer. If the synchronization priority is set to the highest rate, the resynchronization operation is prioritized, but I/O activity for the storage array might be affected.

4. Edit the resynchronization policy as appropriate.

You can resynchronize the mirrored pairs on the remote storage array either manually or automatically.

- **Manual** (the recommended option) — Select this option to require synchronization to be manually resumed after communication is restored to a mirrored pair. This option provides the best opportunity for recovering data.

- **Automatic** — Select this option to start resynchronization automatically after communication is restored to a mirrored pair.

5. Select **Save**.

Remove synchronous mirror relationship

You remove a mirrored pair to remove the mirror relationship from the primary volume on the local storage array and the secondary volume on the remote storage array.

About this task

You can also remove a mirrored pair to correct an orphaned mirrored pair state. Review the following information about orphaned mirrored pairs:

- An orphaned mirrored pair exists when a member volume has been removed on one side (local/remote) but not on the other side.
- Orphaned mirrored pairs are detected when inter-array communication is restored.

Steps

1. Select **Storage › Synchronous Mirroring**.
2. Select the mirrored pair that you want to remove, and then select the **Uncommon Tasks › Remove**.

The **Remove Mirror Relationship** dialog box appears.

3. Confirm that you want to remove the mirrored pair, and then click **Remove**.

Results

System Manager performs the following actions:

- Removes the mirror relationship from the mirrored pair on the local storage array and on the remote storage array.
- Returns the primary volume and the secondary volume to host-accessible, non-mirrored volumes.
- Updates the Synchronous Mirroring tile with the removal of the synchronous mirrored pair.

Deactivate synchronous mirroring

You can deactivate the Synchronous Mirroring feature on a storage array to re-establish normal use of host bus adapter (HBA) host port 4, which was reserved for mirror data transmission.

Before you begin

You must have deleted all synchronous mirror relationships. Verify that all mirrored pairs have been deleted from the storage array.

Steps

1. Select **Storage › Synchronous Mirroring**.
2. Select **Uncommon Tasks › Deactivate**.

The system displays a confirmation.

3. Select **Yes** to confirm.

Results

- The controller's HBA host port 4, which was dedicated for synchronous mirroring communication, can now accept host read and write requests.
- The reserved capacity volumes on the storage array are deleted.

FAQs

How does synchronous mirroring differ from asynchronous mirroring?

The Asynchronous Mirroring feature differs from the Synchronous Mirroring feature in one essential way: it captures the state of the source volume at a particular point in time and copies just the data that has changed since the last image capture.

With synchronous mirroring, the state of the primary volume is not captured at some point in time, but rather reflects all changes that were made on the primary volume to the secondary volume. The secondary volume is identical to the primary volume at every moment because, with this type of mirror, each time a write is done to the primary volume, a write is done to the secondary volume. The host does not receive an acknowledgment that the write was successful until the secondary volume is successfully updated with the changes that were made on the primary volume.

With asynchronous mirroring, the remote storage array is not fully synchronized with the local storage array, so if the application needs to transition to the remote storage array due to a loss of the local storage array, some transactions could be lost.

Comparison between mirroring features

Asynchronous Mirroring	Synchronous Mirroring
Replication method	
<ul style="list-style-type: none">• Point-in-Time <p>Mirroring is done on demand or automatically according to a user-defined schedule. Schedules can be defined at the granularity of minutes. The minimum time between syncs is 10 minutes.</p>	<ul style="list-style-type: none">• Continuous <p>Mirroring is automatically executed continuously, copying data from every host write.</p>
Reserved capacity	
<ul style="list-style-type: none">• Multiple <p>A reserved capacity volume is required for each mirrored pair.</p>	<ul style="list-style-type: none">• Single <p>Single reserved capacity volume is required for all mirrored volumes.</p>
Communication	
<ul style="list-style-type: none">• iSCSI and Fibre Channel <p>Supports iSCSI and Fibre Channel interfaces between storage arrays.</p>	<ul style="list-style-type: none">• Fibre Channel <p>Supports only Fibre Channel interfaces between storage arrays.</p>
Distance	
<ul style="list-style-type: none">• Unlimited <p>Support for virtually unlimited distances between the local storage array and the remote storage array, with the distance typically limited only by the capabilities of the network and the channel extension technology.</p>	<ul style="list-style-type: none">• Restricted <p>Typically must be within about 10 km (6.2 miles), of the local storage array to meet the latency and application performance requirements.</p>

Synchronous mirroring - Why don't I see all my volumes?

When you are selecting a primary volume for a mirrored pair, a list shows all the eligible volumes.

Any volumes that are not eligible to be used do not display in that list. Volumes might not be eligible for any of the following reasons:

- The volume is a non-standard volume, such as a snapshot volume or thin volume.

- The volume is not optimal.
- The volume is already participating in a mirroring relationship.

Synchronous mirroring - Why don't I see all the volumes on the remote storage array?

When you are selecting a secondary volume on the remote storage array, a list shows all the eligible volumes for that mirrored pair.

Any volumes that are not eligible to be used, do not display in that list. Volumes might not be eligible for any of the following reasons:

- The volume is a non-standard volume, such as a snapshot volume or thin volume.
- The volume is not optimal.
- The volume is already participating in a mirroring relationship.
- If you are using Data Assurance (DA), the primary volume and the secondary volume must have the same DA settings.
 - If the primary volume is DA enabled, the secondary volume must be DA enabled.
 - If the primary volume is not DA enabled, the secondary volume must not be DA enabled.

Synchronous mirroring - What do I need to know before creating a mirrored pair?

You configure mirrored pairs in the Unified Manager interface, and then manage the pairs in System Manager.

Before creating a mirrored pair, follow these guidelines:

- You must have two storage arrays.
- Each storage array must have two controllers.
- Each controller in both the primary array and secondary array must have an Ethernet management port configured and must be connected to your network.
- Your local and remote storage arrays are connected through a Fibre Channel fabric.
- The storage arrays have a minimum firmware version of 7.84. (They can each run different OS versions.)
- You must know the password for the local and remote storage arrays.
- You must have enough free capacity on the remote storage array to create a secondary volume equal to or greater than the primary volume that you want to mirror.
- You have installed the Web Services Proxy and Unified Manager. Mirrored pairs are configured in the Unified Manager interface.
- The two storage arrays are discovered in Unified Manager.

What impact does synchronization priority have on synchronization rates?

The synchronization priority defines how much processing time is allocated for synchronization activities relative to system performance.

The controller owner of the primary volume performs this operation in the background. At the same time, the controller owner processes local I/O writes to the primary volume and associated remote writes to the secondary volume. Because the resynchronization diverts controller processing resources from I/O activity,

resynchronization can have a performance impact to the host application.

Keep these guidelines in mind to help you determine how long a synchronization priority might take and how the synchronization priorities can affect system performance.

About synchronization priority rates

These priority rates are available:

- Lowest
- Low
- Medium
- High
- Highest

The lowest priority rate supports system performance, but the resynchronization takes longer. The highest priority rate supports resynchronization, but system performance might be compromised.

These guidelines roughly approximate the differences between the priorities.

Priority rate for full synchronization	Time elapsed compared to highest synchronization rate
Lowest	Approximately eight times as long as at the highest priority rate.
Low	Approximately six times as long as at the highest priority rate.
Medium	Approximately three-and-a-half times as long as at the highest priority rate.
High	Approximately twice as long as at the highest priority rate.

Volume size and host I/O rate loads affect the synchronization time comparisons.

Why is it recommended to use a manual synchronization policy?

Manual resynchronization is recommended because it lets you manage the resynchronization process in a way that provides the best opportunity for recovering data.

If you use an Automatic resynchronization policy and intermittent communication problems occur during resynchronization, data on the secondary volume could be temporarily corrupted. When resynchronization is complete, the data is corrected.

Hardware

Shelves

Concepts

Hardware page overview

The Hardware page provides a graphical depiction of the storage array's physical components. From here, you can check component status and perform some functions related to those components.

Shelves

A shelf is a component that contains the hardware for the storage array (controllers, power/fan canisters, and drives). There are two types of shelves:

- **Controller shelf** — Contains the drives, power/fan canisters, and controllers.
- **Drive shelf** (or **expansion shelf**) — Contains drives, power/fan canisters, and two input/output modules (IOMs). The IOMs, also known as environmental service modules (ESMs), include SAS ports that connect the drive shelf to the controller shelf.

Shelves are available in three sizes for housing up to 12, 24, or 60 drives. Each shelf includes an ID number, which is assigned by the controller firmware. The ID appears on the upper left of the shelf view.

The shelf view on the Hardware page shows the front or back components. You can switch between the two views by selecting either **Show back of shelf** or **Show front of shelf** from the upper right of the shelf view. You can also select **Show all front** or **Show all back** from the bottom of the page. The front and back views show the following:

- **Front components** — Drives and empty drive bays.
- **Back components** — Controllers and power/fan canisters (for controller shelves) or the IOMs and power/fan canisters (for drive shelves).

You can perform the following functions related to shelves:

- Turn on the shelf's locator light, so you can find the physical location of the shelf in the cabinet or rack.
- Change the ID number shown in the upper left of the shelf view.
- View the shelf settings, such as the types of drives installed and the serial number.
- Move the shelf views up or down to match the physical layout in the storage array.

Controllers

A controller is the combined hardware and firmware that implements storage array and management functions. It includes the cache memory, drive support, and host-interface support.

You can perform the following functions related to controllers:

- Configure the management ports for IP addresses and speed.
- Configure iSCSI host connections (if you have iSCSI hosts).
- Configure a Network Time Protocol (NTP) server and a Domain Name System (DNS) server.
- View controller status and settings.
- Allow users from outside the local area network to start an SSH session and change settings on the controller.

- Place the controller offline, online, or in service mode.

Drives


The storage array can include hard disk drives (HDDs) or solid state drives (SSDs). Depending on the shelf size, up to 12, 24, or 60 drives can be installed in the shelf.



You can perform the following functions related to drives:

- Turn on the drive's locator light, so you can find the physical location of the drive in the shelf.
- View drive status and settings.
- Re-assign a drive (logically replace a failed drive with an unassigned drive), and manually reconstruct the drive if necessary.
- Manually fail a drive so you can replace it. (Failing a drive allows you to copy the drive's contents before you replace it.)
- Assign or unassign hot spares.
- Erase secure-enabled drives.

Hardware terminology

Learn how the hardware terms apply to your storage array.

Component	Description
Storage array	A storage array includes the shelves, controllers, drives, software, and firmware.
Shelf	A shelf is an enclosure installed in a cabinet or rack. It contains the hardware components for the storage array. There are two types of shelves: a controller shelf and a drive shelf. A controller shelf includes controllers and drives. A drive shelf includes input/output modules (IOMs) and drives.
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the System Manager functions.
Drive	A drive is an electromagnetic mechanical device or solid state memory device that provides the physical storage media for data.
Bay	A bay is a slot in the shelf where a drive or other component is installed.
Drive shelf	<p>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.</p> <div>  <p>Drive shelves are not available for EF600 storage systems.</p> </div>
Controller shelf	A controller shelf contains a set of drives and one or more controller canisters. A controller canister holds the controllers, host interface cards (HICs), and batteries.

Component	Description
Power/fan canister	A power/fan canister is an assembly that slides into a shelf. It includes a power supply and an integrated fan.
IOM (ESM)	<p>An IOM is an input/output module that includes SAS ports for connecting the drive shelf to the controller shelf.</p> <div>  Drive shelves are not available for EF600 storage systems. </div> <div>  In previous controller models, the IOM was referred to as an environmental service module (ESM). </div>
SFP	An SFP is a Small Form-factor Pluggable (SFP) transceiver.

How to


View hardware components

The Hardware page provides sorting and filtering functions that make it easier to find components.

Steps

1. Select **Hardware**.
2. Use the functions described in the following table to view hardware components.

Function	Description
Front and back shelf views	To switch between front and back shelf views, select either Show back of shelf or Show front of shelf from the far right (the link that appears depends on the current view). The front view shows drives and any empty drive bays. The back view shows the controllers, and any IOM (ESM) modules, power/fan canisters, or empty controller bays. At the bottom of the page, you can also select Show all front or Show all back .

Function	Description
Drive view filters	<p>If the storage array contains drives with different types of physical and logical attributes, the Hardware page includes drive view filters. These filter fields help you quickly locate specific drives by limiting the drive types displayed on the page. Under Show drives that are..., click the filter field on the left (by default, shows Any drive type) to see a drop-down list of physical attributes (for example, capacity and speed). Click the filter field on the right (by default, shows Anywhere in the storage array) to see a drop-down list of logical attributes (for example, volume group assignment). You can use these filters together or separately.</p> <div>  <p>If the storage array contains drives that all share the same physical attributes, the Any drive type field on the left does not appear. If the drives are all in the same logical location, the Anywhere in the storage array field on the right does not appear.</p> </div>
Legend	The components are displayed in certain colors to depict their role states. To expand and collapse the descriptions of these states, click Legend .
Show status icon details	The status indicators can include text descriptions for availability states. Click Show status icon details to show or hide this status text.
Shelf/shelf icons	Each shelf view provides a list of related commands, along with properties and status. Click Shelf to see a drop-down list of commands. You can also select one of the icons along the top to see status and properties for individual components: controllers, IOMs (ESMs), power supplies, fans, temperature, batteries, and SFPs.
Shelf order	The shelves can be rearranged on the Hardware page. Use the up and down arrows on the top right of each shelf view to change the top/bottom order of shelves.

Show or hide component status

You can display status descriptions for drives, controllers, fans, and power supplies.

Steps

1. Select **Hardware**.
2. To see either the back or front components:
 - If you want to see the controller and power/fan canister components, but the drives are displayed, click **Show back of shelf**.
 - If you want to see the drives, but the controller and power/fan canister components are displayed, click **Show front of shelf**.
3. To view or hide pop-over status descriptions:
 - If you want to see a pop-over description of the status icons, click **Show status icon details** at the upper right of the shelf view (select the check box).

- To hide the pop-over descriptions, click **Show status icon details** again (clear the check box).
- 4. If you want to see full status details, select the component in the shelf view, and then select **View settings**.
- 5. If you want to view the descriptions of the colored components, select **Legend**.

Switch between front and back views

The Hardware page can show either the front view or the back view of the shelves.

About this task

The back view shows the controllers/IOMs and the power-fan canisters. The front view shows the drives.

Steps

1. Select **Hardware**.
2. If the graphic shows the drives, click **Show back of shelf**.

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

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

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

4. **Optional:** You can select **Show all front** or **Show all back**, located at the bottom of the page.

Change view order of shelves

You can change the order of shelves displayed on the Hardware page to match the physical order of shelves in a cabinet.

Steps

1. Select **Hardware**.
2. From the top right of a shelf view, select the up or down arrows to rearrange the order of shelves shown on the **Hardware** page.

Turn on shelf locator light

To find the physical location of a shelf shown on the Hardware page, you can turn on the shelf's locator light.

Steps

1. Select **Hardware**.
2. Select the drop-down list for the Controller Shelf or Drive Shelf, and then select **Turn on locator light**.

The locator light for the shelf turns on.

3. When you have physically located the shelf, return to the dialog box and select **Turn off**.

Change shelf IDs

The shelf ID is a number that uniquely identifies a shelf in the storage array. Shelves are numbered consecutively, beginning with either 00 or 01, on the top left of each shelf view.

About this task

The controller firmware automatically assigns the shelf ID, but you can change that number if you want to create a different ordering scheme.

Steps

1. Select **Hardware**.
2. Select the drop-down list for the Controller Shelf or Drive Shelf, and then select **Change ID**.
3. In the **Change Shelf ID** dialog box, select the drop-down list to display available numbers.

This dialog box does not display IDs currently assigned to active shelves.

4. Select an available number, and then click **Save**.

Depending on the number you selected, the shelf order may be rearranged on the Hardware page. If desired, you can use the up/down arrows on the top right of each shelf to readjust the order.

View shelf component status and settings

The Hardware page provides status and settings for shelf components, including the power supplies, fans, and batteries.

About this task







The available components depend on the type of shelf:







- **Drive shelf** — Contains a set of drives, power/fan canisters, input/output modules (IOMs), and other supporting components in a single shelf.
- **Controller shelf** — Contains a set of drives, one or two controller canisters, power/fan canisters, and other supporting components in a single shelf.

Steps

1. Select **Hardware**.
2. Select the drop-down list for the Controller Shelf or Drive Shelf, and then select **View Settings**.

The Shelf Components Settings dialog box opens, with tabs that show the status and settings related to the shelf components. Depending on the type of shelf selected, some tabs described in the table might not appear.

Tab	Description
Shelf	<p>The Shelf tab shows the following properties:</p> <ul style="list-style-type: none"> • Shelf ID — Uniquely identifies a shelf in the storage array. The controller firmware assigns this number, but you can change it by selecting Shelf > Change ID. • Shelf path redundancy — Specifies whether connections between the shelf and the controller have alternate methods in place (Yes) or not (No). • Current drive types — Shows the type of technology built into the drives (for example, a SAS drive that is secure-capable). If there is more than one drive type, both technologies are shown. • Serial number — Shows the serial number of the shelf.
IOMs (ESMs)	<p>The IOMs (ESMs) tab shows status of the input/output module (IOM), which is also called an environmental service module (ESM). It monitors the status of the components in a drive shelf and serves as the connection point between the drive tray and the controller.</p> <p>Status can be Optimal, Failed, Optimal (Miswire), or Uncertified. Other information includes the firmware version and the configuration settings version.</p> <p>Select Show more settings to see the maximum and current data rates, and the state of the card communication (either Yes or No).</p> <div data-bbox="654 1087 711 1140">  </div> <div data-bbox="768 1087 1450 1150"> <p>You can also view this status by selecting the IOM icon , next to the Shelf drop-down list.</p> </div>
Power Supplies	<p>The Power Supplies tab shows the status of the power supply canister and the power supply itself. Status can be Optimal, Failed, Removed, or Unknown. It also shows the part number of the power supply.</p> <div data-bbox="654 1392 711 1444">  </div> <div data-bbox="768 1392 1417 1455"> <p>You can also view this status by selecting the Power Supply icon , next to the Shelf drop-down list.</p> </div>
Fans	<p>The Fans tab shows the status of the fan canister and the fan itself. Status can be Optimal, Failed, Removed, or Unknown.</p> <div data-bbox="654 1633 711 1686">  </div> <div data-bbox="768 1633 1450 1696"> <p>You can also view this status by selecting the Fan icon , next to the Shelf drop-down list.</p> </div>

Tab	Description
Temperature	<p>The Temperature tab shows the temperature status of the shelf components, such as the sensors, controllers, and power/fan canisters. Status can be Optimal, Nominal temperature exceeded, Maximum temperature exceeded, or Unknown.</p> <p> You can also view this status by selecting the Temperature icon , next to the Shelf drop-down list.</p>
Batteries	<p>The Batteries tab shows the status of the controller batteries. Status can be Optimal, Failed, Removed or Unknown. Other information includes the battery age, days until replacement, learn cycles, and weeks between learn cycles.</p> <p> You can also view this status by selecting the Batteries icon , next to the Shelf drop-down list.</p>
SFPs	<p>The SFPs tab shows status of Small Form-factor Pluggable (SFP) transceivers on the controllers. Status can be Optimal, Failed, or Unknown.</p> <p>Select Show more settings to see the part number, the serial number, and the vendor of the SFPs.</p> <p> You can also view this status by selecting the SFP icon , next to the Shelf drop-down list.</p>

3. Click **Close**.

Update battery learn cycles

A learn cycle is an automatic cycle for calibrating the smart battery gauge. The cycles are scheduled to start automatically, at the same day and time, in 8-week intervals (per controller). If you want to set a different schedule, you can adjust the learn cycles.

About this task

Updating the learn cycles affect both controller batteries.

Steps

1. Select **Hardware**.
2. Select the drop-down list for the Controller Shelf, and then select **View settings**.
3. Select the **Batteries** tab.
4. Select **Update battery learn cycles**.

The Update Battery Learn Cycles dialog box opens.

5. From the drop-down lists, select a new day and time.

6. Click **Save**.

FAQs

What is shelf loss protection and drawer loss protection?

Shelf loss protection and drawer loss protection are attributes of pools and volume groups that allow you to maintain data access in the event of a single shelf or drawer failure.

Shelf loss protection

A shelf is the enclosure that contains either the drives or the drives and the controller. Shelf loss protection guarantees accessibility to the data on the volumes in a pool or volume group if a total loss of communication occurs with a single drive shelf. An example of total loss of communication might be loss of power to the drive shelf or failure of both I/O modules (IOMs).



Shelf loss protection is not guaranteed if a drive has already failed in the pool or volume group. In this situation, losing access to a drive shelf and consequently another drive in the pool or volume group causes loss of data.

The criteria for shelf loss protection depends on the protection method, as described in the following table:

Level	Criteria for Shelf Loss Protection	Minimum number of shelves required
Pool	The pool must include drives from at least five shelves and there must be an equal number of drives in each shelf. Shelf loss protection is not applicable to high-capacity shelves; if your system contains high-capacity shelves, refer to Drawer Loss Protection.	5
RAID 6	The volume group contains no more than two drives in a single shelf.	3
RAID 3 or RAID 5	Each drive in the volume group is located in a separate shelf.	3
RAID 1	Each drive in a RAID 1 pair must be located in a separate shelf.	2
RAID 0	Cannot achieve Shelf Loss Protection.	Not applicable

Drawer loss protection

A drawer is one of the compartments of a shelf that you pull out to access the drives. Only the high-capacity shelves have drawers. Drawer loss protection guarantees accessibility to the data on the volumes in a pool or volume group if a total loss of communication occurs with a single drawer. An example of total loss of communication might be loss of power to the drawer or failure of an internal component within the drawer.



Drawer loss protection is not guaranteed if a drive has already failed in the pool or volume group. In this situation, losing access to a drawer (and consequently another drive in the pool or volume group) causes loss of data.

The criteria for drawer loss protection depends on the protection method, as described in the following table:

Level	Criteria for drawer loss protection	Minimum number of drawers required
Pool	Pool candidates must include drives from all drawers, and there must be an equal number of drives in each drawer. The pool must include drives from at least five drawers and there must be an equal number of drives in each drawer. A 60-drive shelf can achieve Drawer Loss Protection when the pool contains 15, 20, 25, 30, 35, 40, 45, 50, 55, or 60 drives. Increments in multiples of 5 can be added to the pool after initial creation.	5
RAID 6	The volume group contains no more than two drives in a single drawer.	3
RAID 3 or RAID 5	Each drive in the volume group is located in a separate drawer.	3
RAID 1	Each drive in a mirrored pair must be located in a separate drawer.	2
RAID 0	Cannot achieve Drawer Loss Protection.	Not applicable

What are battery learn cycles?

A learn cycle is an automatic cycle for calibrating the smart battery gauge.

A learn cycle consists of these phases:

- Controlled battery discharge
- Rest period
- Charge

The batteries are discharged to a predetermined threshold. During this phase, the battery gauge is calibrated.

A learn cycle requires these parameters:

- Fully charged batteries
- No overheated batteries

Learn cycles for duplex controller systems occur simultaneously. For controllers having backup power from more than one battery or set of battery cells, learn cycles occur sequentially.

Learn cycles are scheduled to start automatically at regular intervals, at the same time and on the same day of the week. The interval between cycles is described in weeks.



A learn cycle might take several hours to complete.

Controllers

Concepts

Controller terminology

Learn how the controller terms apply to your storage array.

Component	Description
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the System Manager functions.
Controller shelf	A controller shelf contains a set of drives and one or more controller canisters. A controller canister holds the controllers, host interface cards (HICs), and batteries.
DHCP	Dynamic Host Configuration Protocol (DHCP) is a protocol used on Internet Protocol (IP) networks for dynamically distributing network configuration parameters, such as IP addresses.
DNS	Domain Name System (DNS) is a naming system for devices connected to the Internet or a private network. The DNS server maintains a directory of domain names and translates them to Internet Protocol (IP) addresses.
Duplex configurations	Duplex is a two-controller module configuration within the storage array. Duplex systems are fully redundant with respect to controllers, logical volume paths, and disk paths. If one controller fails, the other controller takes over its I/O to maintain availability. Duplex systems also have redundant fans and power supplies.
Full-duplex / half-duplex connections	Full-duplex and half-duplex refer to connection modes. In full-duplex mode, two devices can communicate simultaneously in both directions. In half-duplex mode, devices can communicate in one direction at a time (one device sends a message, while the other device receives it).
HIC	A host interface card (HIC) can optionally be installed within a controller canister. Host ports that are built into the controller are called baseboard host ports. Host ports that are built into the HIC are called HIC ports.
ICMP PING response	Internet Control Message Protocol (ICMP) is a protocol used by operating systems of networked computers to send messages. ICMP messages determine whether a host is reachable and how long it takes to get packets to and from that host.

Component	Description
MAC address	Media access control identifiers (MAC addresses) are used by Ethernet to distinguish between separate logical channels connecting two ports on the same physical transport network interface.
management client	A management client is the computer where a browser is installed for accessing System Manager.
MTU	A Maximum Transmission Unit (MTU) is the largest size packet or frame that can be sent in a network.
NTP	Network Time Protocol (NTP) is a networking protocol for clock synchronization between computer systems in data networks.
Simplex configurations	Simplex is a single-controller module configuration within the storage array. A simplex system does not offer controller or disk-path redundancy, but does have redundant fans and power supplies.
VLAN	A virtual local area network (VLAN) is a logical network that behaves like it is physically separate from other networks supported by the same devices (switches, routers, etc.).

Controller states

You can place a controller into three different states: online, offline, and service mode.

Online state

The online state is the normal operating state of the controller. It means that the controller is operating normally and is available for I/O operations.

When you place a controller online, its status is set to optimal.

Offline state

The offline state is typically used to prepare a controller for replacement when there are two controllers in the storage array. A controller can enter the offline state in two ways: you can issue an explicit command or the controller can fail. A controller can exit the offline state only by issuing another explicit command or by replacing the failed controller. You can place a controller offline only if there are two controllers in the storage array.

When a controller is in the offline state, the following conditions are true:

- The controller is not available for I/O.
- You cannot manage the storage array through that controller.
- Any volumes currently owned by that controller are moved to the other controller.

Service mode

Service Mode is typically used only by technical support to move all storage array volumes to one controller so that the other controller can be diagnosed. A controller must be manually placed in service mode and must be manually placed back online after the service operation is completed.

When a controller is in service mode, the following conditions are true:

- The controller is not available for I/O.
- Technical support can access the controller through the serial port or network connection to analyze potential problems.
- Any volumes currently owned by that controller are moved to the other controller.
- Cache mirroring is disabled and all volumes are changed to write through cache mode.

Considerations for assigning IP addresses

By default, controllers ship with DHCP enabled on both network ports. You can assign static IP addresses, use the default static IP addresses, or use DHCP-assigned IP addresses. You also can use IPv6 stateless auto-configuration.



IPv6 is disabled by default on new controllers, but you can configure the management port IP addresses using an alternate method, and then enable IPv6 on the management ports using SANtricity System Manager.

When the network port is in a "link down" state, that is, disconnected from a LAN, the system reports its configuration as either static, displaying an IP address of 0.0.0.0 (earlier releases), or DHCP enabled with no IP address reported (later releases). After the network port is in a "link up" state (that is, connected to a LAN), it attempts to obtain an IP address through DHCP.

If the controller is unable to obtain a DHCP address on a given network port, it reverts to a default IP address, which might take up to three minutes. The default IP addresses are as follows:

```
Controller 1 (port 1): IP Address: 169.254.128.101
```

```
Controller 1 (port 2): IP Address: 169.254.129.101
```

```
Controller 2 (port 1): IP Address: 169.254.128.102
```

```
Controller 2 (port 2): IP Address: 169.254.129.102
```

When assigning IP addresses:

- Reserve Port 2 on the controllers for technical support usage. Do not change the default network settings (DHCP enabled).
- To set static IP addresses for E2800 and E5700 controllers, use SANtricity System Manager. After a static

IP address is configured, it remains set through all link down/up events.

- To set static IP addresses for E2700 and E5600 controllers, use SANtricity Storage Manager. After a static IP address is configured, it remains set through all link down/up events.
- To use DHCP to assign the IP address of the controller, connect the controller to a network that can process DHCP requests. Use a permanent DHCP lease.



The default addresses are not persisted across link down events. When a network port on a controller is set to use DHCP, the controller attempts to obtain a DHCP address on every link up event, including cable insertions, reboots, and power cycles. Any time a DHCP attempt fails, the default static IP address for that port is used.

How to

Configure management port

The controller includes an Ethernet port used for system management. If necessary, you can change its transmission parameters and IP addresses.

About this task

During this procedure, you select port 1 and then determine the speed and port addressing method. Port 1 connects to the network where the management client can access the controller and System Manager.



Do not use port 2 on either controller. Port 2 is reserved for use by technical support.

Steps

1. Select **Hardware**.
2. If the graphic shows the drives, click **Show back of shelf**.

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

3. Click the controller with the management port you want to configure.

The controller's context menu appears.

4. Select **Configure management ports**.

The Configure Management Ports dialog box opens.

5. Make sure port 1 is displayed, and then click **Next**.
6. Select the configuration port settings, and then click **Next**.


Field Details

Field	Description
Speed and duplex mode	Keep the Auto-negotiate setting if you want System Manager to determine the transmission parameters between the storage array and the network; or if you know the speed and mode of your network, select the parameters from the drop-down list. Only the valid speed and duplex combinations appear in the list.
Enable IPv4 / Enable IPv6	Select one or both options to enable support for IPv4 and IPv6 networks.

If you select **Enable IPv4**, a dialog box opens for selecting IPv4 settings after you click **Next**. If you select **Enable IPv6**, a dialog box opens for selecting IPv6 settings after you click **Next**. If you select both options, the dialog box for IPv4 settings opens first, and then after you click **Next**, the dialog box for IPv6 settings opens.

7. Configure the IPv4 and/or IPv6 settings, either automatically or manually.

Field Details

Field	Description
Automatically obtain configuration from DHCP server	Select this option to obtain the configuration automatically.
Manually specify static configuration	<div>Select this option, and then enter the controller's IP address. (If desired, you can cut and paste addresses into the fields.) For IPv4, include the network subnet mask and gateway. For IPv6, include the routable IP address and router IP address.</div> <div> If you change the IP address configuration, you lose the management path to the storage array. If you use SANtricity Unified Manager to globally manage arrays in your network, open the user interface and go to Manage > Discover. If you use SANtricity Storage Manager, you must remove the device from the Enterprise Management Window (EMW), add it back to the EMW by selecting Edit > Add Storage Array, and then enter the new IP address.</div>

8. Click **Finish**.

Results

The management port configuration is displayed in the controller settings, Management Ports tab.

Configure iSCSI ports

If your controller includes an iSCSI host connection, you can configure the iSCSI port settings from the Hardware page.

Before you begin

- Your controller must include iSCSI ports; otherwise, the iSCSI settings are not available.
- You must know the network speed (the data transfer rate between the ports and the host).



The iSCSI settings and functions only appear if your storage array supports iSCSI.

Steps

1. Select **Hardware**.
2. If the graphic shows the drives, click **Show back of shelf**.

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

3. Click the controller with the iSCSI ports you want to configure.

The controller's context menu appears.

4. Select **Configure iSCSI ports**.




The **Configure iSCSI ports** option appears only if System Manager detects iSCSI ports on the controller.

The Configure iSCSI Ports dialog box opens.

5. In the drop-down list, select the port you want to configure, and then click **Next**.
6. Select the configuration port settings, and then click **Next**.

To see all port settings, click the **Show more port settings** link on the right of the dialog box.

Field Details

Port Setting	Description
Enable IPv4 / Enable IPv6	<p>Select one or both options to enable support for IPv4 and IPv6 networks.</p> <div> If you want to disable port access, deselect both check boxes.</div>
TCP listening port (Available by clicking Show more port settings .)	<p>If necessary, enter a new port number.</p> <p>The listening port is the TCP port number that the controller uses to listen for iSCSI logins from host iSCSI initiators. The default listening port is 3260. You must enter 3260 or a value between 49152 and 65535.</p>
MTU size (Available by clicking Show more port settings .)	<p>If necessary, enter a new size in bytes for the Maximum Transmission Unit (MTU).</p> <p>The default Maximum Transmission Unit (MTU) size is 1500 bytes per frame. You must enter a value between 1500 and 9000.</p>
Enable ICMP PING responses	<p>Select this option to enable the Internet Control Message Protocol (ICMP). The operating systems of networked computers use this protocol to send messages. These ICMP messages determine whether a host is reachable and how long it takes to get packets to and from that host.</p>

If you selected **Enable IPv4**, a dialog box opens for selecting IPv4 settings after you click **Next**. If you selected **Enable IPv6**, a dialog box opens for selecting IPv6 settings after you click **Next**. If you selected both options, the dialog box for IPv4 settings opens first, and then after you click **Next**, the dialog box for IPv6 settings opens.

7. Configure the IPv4 and/or IPv6 settings, either automatically or manually. To see all port settings, click the **Show more settings** link on the right of the dialog box.

Field Details

Port setting	Description
Automatically obtain configuration	Select this option to obtain the configuration automatically.
Manually specify static configuration	Select this option, and then enter a static address in the fields. (If desired, you can cut and paste addresses into the fields.) For IPv4, include the network subnet mask and gateway. For IPv6, include the routable IP address and router IP address.
Enable VLAN support (Available by clicking Show more settings.)	Select this option to enable a VLAN and enter its ID. A VLAN is a logical network that behaves like it is physically separate from other physical and virtual local area networks (LANs) supported by the same switches, the same routers, or both.
Enable ethernet priority (Available by clicking Show more settings.)	<p>Select this option to enable the parameter that determines the priority of accessing the network. Use the slider to select a priority between 1 (lowest) and 7 (highest).</p> <p>In a shared local area network (LAN) environment, such as Ethernet, many stations might contend for access to the network. Access is on a first-come, first-served basis. Two stations might try to access the network at the same time, which causes both stations to back off and wait before trying again. This process is minimized for switched Ethernet, where only one station is connected to a switch port.</p>

8. Click **Finish**.

Configure iSER over InfiniBand ports

If your controller includes an iSER over InfiniBand port, you can configure the network connection to the host.

Before you begin

- Your controller must include an iSER over InfiniBand port; otherwise, the iSER over InfiniBand settings are not available in System Manager.
- You must know the IP address of the host connection.

Steps

1. Select **Hardware**.
2. If the graphic shows the drives, click **Show back of shelf**.

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

3. Click the controller with the iSER over InfiniBand port you want to configure.

The controller's context menu appears.

4. Select **Configure iSER over InfiniBand ports**.

The Configure iSER over InfiniBand Ports dialog box opens.

5. In the drop-down list, select the HIC port you want to configure, and then enter the IP address of the host.

6. Click **Configure**.

7. Complete the configuration, and then reset the iSER over InfiniBand port by clicking **Yes**.

Configure NVMe over InfiniBand ports

If your controller includes an NVMe over InfiniBand connection, you can configure the NVMe port settings from the Hardware page.

Before you begin

- Your controller must include an NVMe over InfiniBand host port; otherwise, the NVMe over InfiniBand settings are not available in System Manager.
- You must know the IP address of the host connection.



The NVMe over InfiniBand settings and functions appear only if your storage array's controller includes an NVMe over InfiniBand port.

Steps

1. Select **Hardware**.

2. If the graphic shows the drives, click **Show back of shelf**.

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

3. Click the controller with the NVMe over InfiniBand port you want to configure.

The controller's context menu appears.

4. Select **Configure NVMe over InfiniBand ports**.

The Configure NVMe over InfiniBand Ports dialog box opens.

5. Select the HIC port you want to configure from the drop-down list, and then enter the IP address.

If you are configuring an EF600 storage array with a 200Gb-capable HIC, this dialog box displays two IP Address fields, one for a physical port (external) and one for a virtual port (internal). You should assign a unique IP address for both ports. These settings allow the host to establish a path between each port, and for the HIC to achieve maximum performance. If you do not assign an IP address to the virtual port, the HIC will run at approximately half its capable speed.

6. Click **Configure**.

7. Complete the configuration, and then reset the NVMe over InfiniBand port by clicking **Yes**.

Configure NVMe over RoCE ports

If your controller includes a connection for NVMe over RoCE (RDMA over Converged Ethernet), you can configure the NVMe port settings from the Hardware page.

Before you begin

- Your controller must include an NVMe over RoCE host port; otherwise, the NVMe over RoCE settings are not available in System Manager.
- You must know the IP address of the host connection.

Steps

1. Select **Hardware**.
2. If the graphic shows the drives, click **Show back of shelf**.

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

3. Click the controller with the NVMe over RoCE port you want to configure.

The controller’s context menu appears.


4. Select **Configure NVMe over RoCE ports**.

The Configure NVMe over RoCE Ports dialog box opens.

5. In the drop-down list, select the HIC port you want to configure.
6. Click **Next**.

To see all port settings, click the **Show more port settings** link on the right of the dialog box.

Field Details

Port Setting	Description
Configured ethernet port speed	Select the speed that matches the speed capability of the SFP on the port.
Enable IPv4 / Enable IPv6	<div>Select one or both options to enable support for IPv4 and IPv6 networks.</div> <div> If you want to disable port access, deselect both check boxes.</div>
MTU size (Available by clicking Show more port settings .)	<div>If necessary, enter a new size in bytes for the Maximum Transmission Unit (MTU).</div> <div>The default Maximum Transmission Unit (MTU) size is 1500 bytes per frame. You must enter a value between 1500 and 9000.</div>

If you selected **Enable IPv4**, a dialog box opens for selecting IPv4 settings after you click **Next**. If you selected **Enable IPv6**, a dialog box opens for selecting IPv6 settings after you click **Next**. If you selected both options, the dialog box for IPv4 settings opens first, and then after you click **Next**, the dialog box for IPv6 settings opens.

7. Configure the IPv4 and/or IPv6 settings, either automatically or manually.

Field Details

Port setting	Description
Automatically obtain configuration	Select this option to obtain the configuration automatically.
Manually specify static configuration	<p>Select this option, and then enter a static address in the fields. (If desired, you can cut and paste addresses into the fields.) For IPv4, include the network subnet mask and gateway. For IPv6, include the routable IP address and router IP address.</p> <p>If you are configuring an EF600 storage array with a 200Gb-capable HIC, this dialog box displays two sets of fields for network parameters, one for a physical port (external) and one for a virtual port (internal). You should assign unique parameters for both ports. These settings allow the host to establish a path between each port, and for the HIC to achieve maximum performance. If you do not assign an IP address to the virtual port, the HIC will run at approximately half its capable speed.</p>

8. Click **Finish**.

Configure NTP server addresses

You can configure a connection to the Network Time Protocol (NTP) server so that the controller periodically queries the NTP server to update its internal time-of-day clock.

Before you begin

- An NTP server must be installed and configured in your network.
- You must know the address of the primary NTP server and an optional backup NTP server. These addresses can be fully qualified domain names, IPv4 addresses, or IPv6 addresses.



If you enter one or more domain names for the NTP servers, you must also configure a DNS server to resolve the NTP server address. You need to configure the DNS server only on those controllers where you configured NTP and provided a domain name.

About this task

NTP enables the storage array to automatically synchronize the controller's clocks with an external host using Simple Network Time Protocol (SNTP). The controller periodically queries the configured NTP server, and then uses the results to update its internal time-of-day clock. If only one controller has NTP enabled, the alternate controller periodically synchronizes its clock with the controller that has NTP enabled. If neither controller has NTP enabled, the controllers periodically synchronize their clocks with each other.



You do not need to configure NTP on both controllers; however, doing so improves the storage array's ability to stay synchronized during hardware or communication failures.

Steps

1. Select **Hardware**.

2. If the graphic shows the drives, click **Show back of shelf**.

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

3. Click the controller you want to configure.

The controller's context menu appears.

4. Select **Configure NTP server**.

The Configure Network Time Protocol (NTP) Server dialog box opens.

5. Select **I want to enable NTP on Controller (A or B)**.

Additional selections appear in the dialog box.

6. Select one of the following options:

- **Automatically obtain NTP server addresses from DHCP server** — The detected NTP server addresses are shown.



If the storage array is set to use a static NTP address, no NTP servers appear.

- **Manually specify NTP server addresses** — Enter the primary NTP server address and a backup NTP server address. The backup server is optional. (These address fields appear after you select the radio button.) The server address can be a fully qualified domain name, IPv4 address, or IPv6 address.

7. **Optional:** Enter server information and authentication credentials for a backup NTP server.

8. Click **Save**.

Results

The NTP server configuration is displayed in the controller settings, **DNS / NTP** tab.

Configure DNS server addresses

Domain Name System (DNS) is used to resolve fully qualified domain names for the controllers and a Network Time Protocol (NTP) server. The management ports on the storage array can support IPv4 or IPv6 protocols simultaneously.

Before you begin

- A DNS server must be installed and configured in your network.
- You know the address of the primary DNS server and an optional backup DNS server. These addresses can be IPv4 addresses or IPv6 addresses.

About this task

This procedure describes how to specify a primary and backup DNS server address. The backup DNS server can be optionally configured to use if a primary DNS server fails.



If you already configured the storage array's management ports with Dynamic Host Configuration Protocol (DHCP), and you have one or more DNS or NTP servers associated with the DHCP setup, then you do not need to manually configure DNS or NTP. In this case, the storage array should have already obtained the DNS/NTP server addresses automatically. However, you should still follow the instructions below to open the dialog box and make sure that the correct addresses are detected.

Steps

1. Select **Hardware**.
2. If the graphic shows the drives, click **Show back of shelf**.

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

3. Select the controller to configure.

The controller's context menu appears.

4. Select **Configure DNS server**.

The **Configure Domain Name System (DNS) Server** dialog box opens.

5. Select one of the following options:

- **Automatically obtain DNS server addresses from DHCP server** — The detected DNS server addresses are shown.



If the storage array is set to use a static DNS address, no DNS servers appear.

- **Manually specify DNS server addresses** — Enter a primary DNS server address and a backup DNS server address. The backup server is optional. (These address fields appear after you select the radio button.) These addresses can be IPv4 addresses or IPv6 addresses.

6. Click **Save**.
7. Repeat these steps for the other controller.

Results

The DNS configuration is displayed in the controller settings, **DNS / NTP** tab.

View controller settings

You can view information about a controller, such as the status of the host interfaces, drive interfaces, and management ports.

Steps

1. Select **Hardware**.
2. If the graphic shows the drives, click **Show back of shelf**.

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


3. Do one of the following actions to display the controller settings:
 - Click the controller to display the context menu, and then select **View settings**.

- Select the controller icon (next to the **Shelf** drop-down list). For duplex configurations, select either **Controller A** or **Controller B** from the dialog box, and then click **Next**. The Controller Settings dialog box opens.

4. Select the tabs to move between property settings.

Some tabs have a link for **Show more settings** at the top right.

Field Details

Tab	Description
Base	Shows the controller status, model name, replacement part number, current firmware version, and the non-volatile static random access memory (NVSRAM) version.
Cache	Shows the cache settings of the controller, which include the data cache, processor cache, and the cache backup device. The cache backup device is used to back up data in the cache if you lose power to the controller. Status can be Optimal, Failed, Removed, Unknown, Write Protected, or Incompatible.
Host Interfaces	<p>Shows the host interface information and the link status of each port. The host interface is the connection between the controller and the host, such as Fibre Channel or iSCSI.</p> <div>  <p>The host interface card (HIC) location is either in the baseboard or in a slot (bay). "Baseboard" indicates that the HIC ports are built into the controller. "Slot" ports are on the optional HIC.</p> </div>
Drive Interfaces	Shows the drive interface information and the link status of each port. The drive interface is the connection between the controller and the drives, such as SAS.
Management Ports	Shows the management port details, such as the host name used to access the controller and whether a remote login has been enabled. The management port connects the controller and the management client, which is where a browser is installed for accessing System Manager.
DNS / NTP	<p>Shows the addressing method and IP addresses for the DNS server and the NTP server, if these servers have been configured in System Manager. Domain Name System (DNS) is a naming system for devices connected to the Internet or a private network. The DNS server maintains a directory of domain names and translates them to Internet Protocol (IP) addresses.</p> <p>Network Time Protocol (NTP) is a networking protocol for clock synchronization between computer systems in data networks.</p>

5. Click **Close**.

Allow remote login

Remote login lets users from outside of the local area network start an SSH session and change settings on the controller.



Security risk — For security reasons, enable remote login for use only by technical support.

Steps

1. Select **Hardware**.
2. If the graphic shows the drives, click **Show back of shelf**.

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

3. Click the controller for which you want to enable remote login.

The controller's context menu appears.

4. Select **Change remote login**, and confirm that you want to perform the operation.

Results

After technical support is finished troubleshooting, follow this same procedure to disable remote login. Disabling remote login terminates any current SSH sessions and rejects any new login requests.

Place controller online

If a controller is in the offline state or in service mode, you can place it back online.

Steps

1. Select **Hardware**.
2. If the graphic shows the drives, click **Show back of shelf**.

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

3. Click a controller that is in either the offline state or service mode.

The controller's context menu appears.

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

Results

Detection of a restored preferred path by the multipath driver can take up to 10 minutes.

Any volumes originally owned by this controller are automatically moved back to the controller as I/O requests are received for each volume. In some cases, you might need to manually redistribute the volumes with the **Redistribute volumes** command.

Place controller offline

If you are instructed to do so, you can place a controller offline.

Before you begin

- Your storage array must have two controllers. The controller that you are not placing offline must be online

(in the optimal state).

- Make sure that no volumes are in use or that you have a multipath driver installed on all hosts using these volumes.

About this task



Do not place a controller offline unless you are instructed to do so by the Recovery Guru or technical support.

Steps

1. Select **Hardware**.
2. If the graphic shows the drives, click **Show back of shelf**.

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

3. Click the controller that you want to place offline.

The controller's context menu appears.

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

Results

It might take several minutes for System Manager to update the controller's status to offline. Do not begin any other operations until after the status has been updated.

Place controller in service mode

If you are instructed to do so, you can place a controller in service mode.

Before you begin

- The storage array must have two controllers. The controller that you are not placing in service mode must be online (in the optimal state).
- Make sure that no volumes are in use or that you have a multipath driver installed on all hosts using these volumes.

About this task



Placing a controller in service mode might significantly reduce performance. Do not place a controller in service mode unless you are instructed to do so by technical support.

Steps

1. Select **Hardware**.
2. If the graphic shows the drives, click **Show back of shelf**.

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

3. Click the controller that you want to place into service mode.

The controller's context menu appears.

4. Select **Place in service mode**, and confirm that you want to perform the operation.

Reset (reboot) controller

Some issues require a controller reset (reboot). You can reset the controller even if you don't have physical access to it.

Before you begin

- The storage array must have two controllers. The controller that you are not resetting must be online (in the optimal state).
- Make sure that no volumes are in use or that you have a multipath driver installed on all hosts using these volumes.

Steps

1. Select **Hardware**.
2. If the graphic shows the drives, click **Show back of shelf**.

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

3. Click the controller that you want to reset.

The controller's context menu appears.

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

FAQs

What is auto-negotiation?

Auto-negotiation is the ability of a network interface to automatically coordinate its own connection parameters (speed and duplex) with another network interface.

Auto-negotiation is usually the preferred setting for configuring management ports; however, if the negotiation fails, mismatched network interface settings can severely impact network performance. In cases where that condition is unacceptable, you should manually set the network interface settings to a correct configuration. Auto-negotiation is performed by the controller's Ethernet management ports. Auto-negotiation is not performed by iSCSI host bus adapters.



If auto-negotiation fails, the controller attempts to establish a connection at 10BASE-T, half-duplex, which is the lowest common denominator.

What is IPv6 stateless address auto-configuration?

With stateless auto-configuration, hosts do not obtain addresses and other configuration information from a server. Stateless auto-configuration in IPv6 features link-local addresses, multicasting, and the Neighbor Discovery (ND) protocol. IPv6 can generate the interface ID of an address from the underlying data link layer address.

Stateless auto-configuration and stateful auto-configuration complement each other. For example, the host can use stateless auto-configuration to configure its own addresses, but use stateful auto-configuration to obtain other information. Stateful auto-configuration allows hosts to obtain addresses and other configuration

information from a server. Internet Protocol version 6 (IPv6) also defines a method whereby all of the IP addresses on a network can be renumbered at one time. IPv6 defines a method for devices on the network to automatically configure their IP address and other parameters without the need for a server.

Devices perform these steps when using stateless auto-configuration:

1. **Generate a link-local address** — The device generates a link-local address, which has 10 bits, followed by 54 zeros, and followed by the 64-bit interface ID.
2. **Test the uniqueness of a link-local address** — The node tests to make sure that the link-local address that it generates is not already in use on the local network. The node sends a neighbor solicitation message by using the ND protocol. In response, the local network listens for a neighbor advertisement message, which indicates that another device is already using the link-local address. If so, either a new link-local address must be generated or auto-configuration fails, and another method must be used.
3. **Assign a link-local address** — If the device passes the uniqueness test, the device assigns the link-local address to its IP interface. The link-local address can be used for communication on the local network but not over the Internet.
4. **Contact the router** — The node tries to contact a local router for more information about continuing the configuration. This contact is performed either by listening for router advertisement messages sent periodically by the routers or by sending a specific router solicitation message to ask a router for information about what to do next.
5. **Provide direction to the node** — The router provides direction to the node about how to proceed with auto-configuration. Alternatively, the router tells the host how to determine the global Internet address.
6. **Configure the global address** — The host configures itself with its globally unique Internet address. This address is generally formed from a network prefix provided to the host by the router.

Which do I choose — DHCP or manual configuration?

The default method for network configuration is Dynamic Host Configuration Protocol (DHCP). Always use this option unless your network does not have a DHCP server.

What is a DHCP server?

Dynamic Host Configuration Protocol (DHCP) is a protocol that automates the task of assigning an Internet Protocol (IP) address.

Each device that is connected to a TCP/IP network must be assigned a unique IP address. These devices include the controllers in your storage array.

Without DHCP, a network administrator enters these IP addresses manually. With DHCP, when a client needs to start TCP/IP operations, the client broadcasts a request for address information. The DHCP server receives the request, assigns a new address for a specified amount of time called a lease period, and sends the address to the client. With DHCP, a device can have a different IP address each time it connects to the network. In some systems, the IP address for the device can change even while the device is still connected.

How do I configure my DHCP server?

You must configure a Dynamic Host Configuration Protocol (DHCP) server to use static Internet Protocol (IP) addresses for the controllers in your storage array.

The IP addresses that your DHCP server assigns are generally dynamic and can change because they have a lease period that expires. Some devices, for example, servers and routers, need to use static addresses. The

controllers in your storage array also need static IP addresses.

For information about how to assign static addresses, see the documentation for your DHCP server.

Why do I need to change the controller network configuration?

You must set the network configuration for each controller—its Internet Protocol (IP) address, subnetwork mask (subnet mask), and gateway—when you use out-of-band management.

You can set the network configuration by using a Dynamic Host Configuration Protocol (DHCP) server. If you are not using a DHCP server, you must enter the network configuration manually.

Where do I get the network configuration?

You can get the Internet Protocol (IP) address, subnetwork mask (subnet mask), and gateway information from your network administrator.

You need this information when you are configuring ports on the controllers.

What are ICMP PING responses?

Internet Control Message Protocol (ICMP) is one of the protocols of the TCP/IP suite.

The `ICMP echo request` and the `ICMP echo reply` messages are commonly known as `ping` messages. `Ping` is a troubleshooting tool used by system administrators to manually test for connectivity between network devices, and also to test for network delay and packet loss. The `ping` command sends an `ICMP echo request` to a device on the network, and the device immediately responds with an `ICMP echo reply`. Sometimes, a company's network security policy requires `ping` (`ICMP echo reply`) to be disabled on all devices to make them more difficult to be discovered by unauthorized persons.

When should I refresh the port configuration or the iSNS server from the DHCP server?

Refresh the DHCP server any time the server is modified or upgraded, and the DHCP information relevant to the current storage array and the storage array that you want to use has changed.

Specifically, refresh the port configuration or the iSNS server from the DHCP server when you know that the DHCP server will be assigning different addresses.



Refreshing a port configuration is destructive to all of the iSCSI connections on that port.

What should I do after configuring the management ports?

If you changed the IP address for the storage array, you might want to update the global array view in SANtricity Unified Manager.

To update the global array view in Unified Manager, open the interface and go to **Manage > Discover**.

If you are still using the SANtricity Storage Manager, go to the Enterprise Management Window (EMW), where you must remove and re-add the new IP address.

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 <i>secure-capable</i> 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 <i>secure-enabled</i> .
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 <i>secure-capable</i> drives, the drives become <i>secure-enabled</i> . 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: <ul style="list-style-type: none"> • 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.



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.



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.



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.



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	<p>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.</p> <p>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.</p> <p>If the drive is unassigned, this field does not appear.</p>
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. <div>  Not available for NVMe drives. </div>
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 <i>secure-capable</i> 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 <i>secure-enabled</i> .
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 <i>secure-capable</i> drives, the drives become <i>secure-enabled</i> . 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**.



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.



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.

Settings

Alerts

Concepts

How alerts work

Alerts notify administrators about important events that occur on the storage array. Alerts can be sent through email, SNMP traps, and syslog.

The alerts process works as follows:

1. An administrator configures one or more of the following alerting methods in System Manager:
 - **Email** — Messages are sent to email addresses.
 - **SNMP** — SNMP traps are sent to an SNMP server.
 - **Syslog** — Messages are sent to a syslog server.
2. When the storage array's event monitor detects an issue, it writes information about that issue to the event log (available from **Support > Event Log**). For example, issues can include such events as a battery failure, a component moving from Optimal to Offline, or redundancy errors in the controller.
3. If the event monitor determines that the event is "alertable," it then sends a notification using the configured alerting methods (email, SNMP, and/or syslog). All Critical events are considered "alertable," along with some Warning and Informational events.

Alerts configuration

You can configure alerts from the Initial Setup wizard (for email alerts only) or from the Alerts page. To check the current configuration, go to **Settings > Alerts**.

The Alerts tile displays the alerts configuration, which can be one of the following:

- Not configured.

- Configured; at least one alerting method is set up. To determine which alerting methods are configured, point the cursor at the tile.

Alerts information

Alerts can include the following types of information:

- Name of the storage array.
- Event error type related to an event log entry.
- Date and time when the event occurred.
- Brief description of the event.



Syslog alerts follow the RFC 3164 messaging standard.

Alerts terminology

Learn how the alerts terms apply to your storage array.

Component	Description
Event monitor	The event monitor resides on the storage array and runs as a background task. When the event monitor detects anomalies on the storage array, it writes information about the issues to the event log. Issues can include such events as a battery failure, a component moving from Optimal to Offline, or redundancy errors in the controller. If the event monitor determines that the event is "alertable," it then sends a notification using the configured alerting methods (email, SNMP, and/or syslog). All Critical events are considered "alertable," along with some Warning and Informational events.
Mail server	The mail server is used for sending and receiving email alerts. The server uses Simple Mail Transfer Protocol (SMTP).
SNMP	Simple Network Management Protocol (SNMP) is an Internet-standard protocol used for managing and sharing information between devices on IP networks.
SNMP trap	An SNMP trap is a notification sent to an SNMP server. The trap contains information about significant issues with the storage array.
SNMP trap destination	An SNMP trap destination is an IPv4 or IPv6 address of the server running an SNMP service.
Community name	A community name is a string that acts like a password for the network server(s) in a SNMP environment.
MIB file	The management information base (MIB) file defines the data being monitored and managed in the storage array. It must be copied and compiled on the server with the SNMP service application. This MIB file is available with the System Manager software on the Support site.

Component	Description
MIB variables	Management Information Base (MIB) variables can return values such as the storage array name, array location, and a contact person in response to SNMP GetRequests.
Syslog	Syslog is a protocol used by network devices for sending event messages to a logging server.
UDP	User Datagram Protocol (UDP) is a transport layer protocol that specifies a source and destination port number in their packet headers.

How to

Manage email alerts

Configure mail server and recipients for alerts

To configure email alerts, you must specify a mail server address and the email addresses of the alert recipients. Up to 20 email addresses are allowed.

Before you begin

- The address of the mail server must be available. The address can be an IPv4 or IPv6 address, or a fully qualified domain name.



To use a fully qualified domain name, you must configure a DNS server on both controllers. You can configure a DNS server from the Hardware page.

- Email address to be used as the alert sender must be available. This is the address that appears in the "From" field of the alert message. A sender address is required in the SMTP protocol; without it, an error results.
- Email address(es) of the alert recipient(s) must be available. The recipient is typically an address for a network administrator or storage administrator. You can enter up to 20 email addresses.

About this task

This task describes how to configure the mail server, enter email addresses for the sender and recipients, and test all the email addresses entered from the Alerts page.



Email alerts can also be configured from the Initial Setup wizard.

Steps

1. Select **Settings > Alerts**.
2. Select the **Email** tab.

If an email server is not yet configured, the Email tab displays "Configure Mail Server."

3. Select **Configure Mail Server**.

The **Configure Mail Server** dialog box opens.

4. Enter the mail server information, and then click **Save**.

- **Mail server address** — Enter a fully qualified domain name, IPv4 address, or IPv6 address of the mail server.



To use a fully qualified domain name, you must configure a DNS server on both controllers. You can configure a DNS server from the **Hardware** page.

- **Email sender address** — Enter a valid email address to be used as the sender of the email. This address appears in the "From" field of the email message.
- **Include contact information in email** — To include the sender's contact information with the alert message, select this option, and then enter a name and phone number. After you click **Save**, the email addresses appear in the **Email** tab of the **Alerts** page.

5. Select **Add Emails**.

The Add Emails dialog box opens.

6. Enter one or more email addresses for the alert recipients, and then click **Add**.

The email addresses appear on the Alerts page.

7. If you want to make sure the email addresses are valid, click **Test All Emails** to send test messages to the recipients.

Results

After you configure email alerts, the event monitor sends email messages to the specified recipients whenever an alertable event occurs.

Edit email addresses for alerts

You can change the email addresses of the recipients who receive email alerts.

Before you begin

The email address you intend to edit must be defined in the Email tab of the Alerts page.

Steps

1. Select **Settings > Alerts**.
2. Select the **Email** tab.
3. From the **Email Address** table, select the address you want to change, and then click the **Edit** (pencil) icon on the far right.

The row becomes an editable field.

4. Enter a new address, and then click the **Save** (checkmark) icon.



If you want to cancel changes, select the **Cancel** (X) icon.

Results

The Email tab of the Alerts page displays the updated email addresses.

Add email addresses for alerts

You can add up to 20 recipients for email alerts.

Steps

1. Select **Settings > Alerts**.
2. Select the **Email** tab.
3. Select **Add Emails**.

The **Add Emails** dialog box opens.

4. In the empty field, enter a new email address. If you want to add more than one address, select **Add another email** to open another field.
5. Click **Add**.

Results

The **Email** tab of the **Alerts** page displays the new email addresses.

Delete mail server or email addresses for alerts

You can remove the previously defined mail server so that alerts are no longer sent to the email addresses, or you can remove individual email addresses.

Steps

1. Select **Settings > Alerts**.
2. Select the **Email** tab.
3. From the table, do one of the following:
 - To remove a mail server so that alerts are no longer sent to the email addresses, select the row for the mail server.
 - To remove an email address so that alerts are no longer sent to this address, select the row for the email address you want to delete. The **Delete** button in the upper right of the table becomes available for selection.
4. Click **Delete**, and confirm the operation.

Edit mail server for alerts

You can change the mail server address and email sender address used for email alerts.

Before you begin

The address of the mail server you are changing must be available. The address can be an IPv4 or IPv6 address, or a fully qualified domain name.



To use a fully qualified domain name, you must configure a DNS server on both controllers. You can configure a DNS server from the Hardware page.

Steps

1. Select **Settings > Alerts**.

2. Select the **Email** tab.
3. Select **Configure Mail Server**.

The Configure Mail Server dialog box opens.

4. Edit the mail server address, sender information, and contact information.

- **Mail server address** — Edit the fully qualified domain name, IPv4 address, or IPv6 address of the mail server.



To use a fully qualified domain name, you must configure a DNS server on both controllers. You can configure a DNS server from the Hardware page.

- **Email sender address** — Edit the email address to be used as the sender of the email. This address appears in the "From" field of the email message.
- **Include contact information in email** — To edit the sender's contact information, select this option, and then edit the name and phone number.

5. Click **Save**.

Manage SNMP alerts

Configure communities and destinations for SNMP alerts

To configure Simple Network Management Protocol (SNMP) alerts, you must identify at least one server where the storage array's event monitor can send SNMP traps. The configuration requires a community name and IP address for the server.

Before you begin

- A network server must be configured with an SNMP service application. You need the network address of this server (either an IPv4 or an IPv6 address), so the event monitor can send trap messages to that address. You can use more than one server (up to 10 servers are allowed).
- A community name must be created, consisting of only printable ASCII characters. The community name, which is a string that acts like a password for the network servers, is typically created by a network administrator. Up to 256 communities can be created.
- The management information base (MIB) file has been copied and compiled on the server with the SNMP service application. This MIB file defines the data being monitored and managed.

If you do not have the MIB file, you can obtain it from the NetApp Support site:

- Go to [NetApp Support](#).
- Click **Downloads**.
- Click **Software**.
- Find your management software (for example, SANtricity System Manager), and then click **Go!** on the right.
- Click **View & Download** on the latest version.
- Click **Continue** at the bottom of the page.
- Accept the EULA.
- Scroll down until you see **MIB file for SNMP traps**, and then click the link to download the file.

About this task

This task describes how to identify the SNMP server for trap destinations, and then test your configuration.

Steps

1. Select **Settings > Alerts**.
2. Select the **SNMP** tab.

If a community is not yet configured, the SNMP tab displays "Configure Communities."

3. Select **Configure Communities**.

The **Configure Communities** dialog box opens.

4. In the **Community Name** field, enter one or more community strings for the network servers, and then click **Save**.

The Alerts page displays "Add Trap Destinations."

5. Select **Add Trap Destinations**.

The **Add Trap Destinations** dialog box opens.

6. Enter one or more trap destinations, select their associated community names, and then click **Add**.
 - **Trap Destination** — Enter an IPv4 or IPv6 address of the server running an SNMP service.
 - **Community name** — From the drop-down, select the community name for this trap destination. (If you defined only one community name, the name already appears in this field.)
 - **Send Authentication Failure Trap** — Select this option (the checkbox) if you want to alert the trap destination whenever an SNMP request is rejected because of an unrecognized community name. After you click **Add**, the trap destinations and associated community names appear in the **SNMP** tab of the **Alerts** page.
7. To make sure a trap is valid, select a trap destination from the table, and then click **Test Trap Destination** to send a test trap to the configured address.

Results

The event monitor sends SNMP traps to the server(s) whenever an alertable event occurs.

Edit community names for SNMP traps

You can edit community names for SNMP traps, and also associate a different community name to an SNMP trap destination.

Before you begin

A community name must be created, consisting of only printable ASCII characters. The community name, which is a string that acts like a password for the network servers, is created by a network administrator.

Steps

1. Select **Settings > Alerts**.
2. Select the **SNMP** tab.

The trap destinations and community names appear in the table.

3. Edit community names as follows:

- To edit a community name, select **Configure Communities**. Enter the new community name, and then click **Save**. Community names can consist of only printable ASCII characters.
- To associate a community name to a new trap destination, select the community name from the table, and then click the **Edit** (pencil) icon on the far right. From the Community Name drop-down, select a new community name for an SNMP trap destination, and then click the **Save** (checkmark) icon.



If you want to cancel changes, select the **Cancel** (X) icon.

Results

The **SNMP** tab of the **Alerts** page displays the updated communities.

Add community names for SNMP traps

You can add up to 256 community names for SNMP traps.

Before you begin

The community name(s) must be created. The community name, which is a string that acts like a password for the network servers, is typically created by a network administrator. It consists of only printable ASCII characters.

Steps

1. Select **Settings > Alerts**.
2. Select the **SNMP** tab.

The trap destinations and community names appear in the table.

3. Select **Configure Communities**.

The Configure Communities dialog box opens.

4. Select **Add another community**.
5. Enter the new community name, and then click **Save**.

Results

The new community name appears in the **SNMP** tab of the **Alerts** page.

Remove community name for SNMP traps

You can remove a community name for SNMP traps.

Steps

1. Select **Settings > Alerts**.
2. Select the **SNMP** tab.

The trap destinations and community names appear on the Alerts page.

3. Select **Configure Communities**.

The **Configure Communities** dialog box opens.

4. Select the community name you want to delete, and then click the **Remove** (X) icon on the far right.

If trap destinations are associated with this community name, the **Confirm Remove Community** dialog box shows the affected trap destination addresses.

5. Confirm the operation, and then click **Remove**.

Results

The community name and its associated trap destination are removed from the **Alerts** page.

Configure SNMP MIB variables

For SNMP alerts, you can optionally configure Management Information Base (MIB) variables that appear in SNMP traps. These variables can return the storage array name, array location, and a contact person.

Before you begin

The MIB file must be copied and compiled on the server with the SNMP service application.

If you do not have a MIB file, you can obtain it as follows:

- Go to [NetApp Support](#).
- Click **Downloads**.
- Click **Software**.
- Find your management software (for example, SANtricity System Manager), and then click **Go!** on the right.
- Click **View & Download** on the latest version.
- Click **Continue** at the bottom of the page.
- Accept the EULA.
- Scroll down until you see **MIB file for SNMP traps**, and then click the link to download the file.

About this task

This task describes how to define MIB variables for SNMP traps. These variables can return the following values in response to SNMP GetRequests:

- *sysName* (name for the storage array)
- *sysLocation* (location of the storage array)
- *sysContact* (name of an administrator)

Steps

1. Select **Settings > Alerts**.
2. Select the **SNMP** tab.
3. Select **Configure SNMP MIB Variables**.

The Configure SNMP MIB Variables dialog box opens.

4. Enter one or more of the following values, and then click **Save**.

- **Name** — The value for the MIB variable *sysName*. For example, enter a name for the storage array.
- **Location** — The value for the MIB variable *sysLocation*. For example, enter a location of the storage array.
- **Contact** — The value for the MIB variable *sysContact*. For example, enter an administrator responsible for the storage array.

Results

These values appear in SNMP trap messages for storage array alerts.

Add trap destinations for SNMP alerts

You can add up to 10 servers for sending SNMP traps.

Before you begin

- The network server you want to add must be configured with an SNMP service application. You need the network address of this server (either an IPv4 or an IPv6 address), so the event monitor can send trap messages to that address. You can use more than one server (up to 10 servers are allowed).
- A community name must be created, consisting of only printable ASCII characters. The community name, which is a string that acts like a password for the network servers, is typically created by a network administrator. Up to 256 communities can be created.
- The management information base (MIB) file has been copied and compiled on the server with the SNMP service application. This MIB file defines the data being monitored and managed.

If you do not have the MIB file, you can obtain it from the NetApp Support site:

- Go to [NetApp Support](#).
- Click **Downloads**.
- Click **Software**.
- Find your management software (for example, SANtricity System Manager), and then click **Go!** on the right.
- Click **View & Download** on the latest version.
- Click **Continue** at the bottom of the page.
- Accept the EULA.
- Scroll down until you see **MIB file for SNMP traps**, and then click the link to download the file.

Steps

1. Select **Settings > Alerts**.
2. Select the **SNMP** tab.

The currently defined trap destinations appear in the table.

3. Select **Add Trap Destinations**.

The Add Trap Destinations dialog box opens.

4. Enter one or more trap destinations, select their associated community names, and then click **Add**.
 - **Trap Destination** — Enter an IPv4 or IPv6 address of the server running an SNMP service.

- **Community name** — From the drop-down, select the community name for this trap destination. (If you defined only one community name, the name already appears in this field.)
 - **Send Authentication Failure Trap** — Select this option (the checkbox) if you want to alert the trap destination whenever an SNMP request is rejected because of an unrecognized community name. After you click **Add**, the trap destinations and associated community names appear in the table.
5. To make sure a trap is valid, select a trap destination from the table, and then click **Test Trap Destination** to send a test trap to the configured address.

Results

The event monitor sends SNMP traps to the server(s) whenever an alertable event occurs.

Delete trap destinations

You can delete a trap destination address so that the storage array's event monitor no longer sends SNMP traps to that address.

Steps

1. Select **Settings > Alerts**.
2. Select the **SNMP** tab.

The trap destination addresses appear in the table.

3. Select a trap destination, and then click **Delete** in the upper right of the page.
4. Confirm the operation, and then click **Delete**.

The destination address no longer appears on the **Alerts** page.

Results

The deleted trap destination no longer receives SNMP traps from the storage array's event monitor.

Manage syslog alerts

Configure syslog server for alerts

To configure syslog alerts, you must enter a syslog server address and a UDP port. Up to five syslog servers are allowed.

Before you begin

- The syslog server address must be available. This address can be a fully qualified domain name, an IPv4 address, or an IPv6 address.
- UDP port number of the syslog server must be available. This port is typically 514.

About this task

This task describes how to enter the address and port for the syslog server, and then test the address you entered.

Steps

1. Select **Settings > Alerts**.
2. Select the **Syslog** tab.

If a syslog server is not yet defined, the **Alerts** page displays "Add Syslog Servers."

3. Click **Add Syslog Servers**.

The **Add Syslog Server** dialog box opens.

4. Enter information for one or more syslog servers (maximum of five), and then click **Add**.
- **Server Address** — Enter a fully qualified domain name, an IPv4 address, or an IPv6 address.
 - **UDP Port** — Typically, the UDP port for syslog is 514. The table displays the configured syslog servers.
5. To send a test alert to the server addresses, select **Test All Syslog Servers**.

Results

The event monitor sends alerts to the syslog server whenever an alertable event occurs.

Edit syslog servers for alerts

You can edit the server address used for receiving syslog alerts.

Steps

1. Select **Settings > Alerts**.
2. Select the **Syslog** tab.
3. From the table, select a syslog server address, and then click the **Edit** (pencil) icon from on the far right.

The row becomes an editable field.

4. Edit the server address and UDP port number, and then click the **Save** (checkmark) icon.

Results

The updated server address appears in the table.

Add syslog servers for alerts

You can add a maximum of five servers for syslog alerts.

Before you begin

- The syslog server address must be available. This address can be a fully qualified domain name, an IPv4 address, or an IPv6 address.
- The UDP port number of the syslog server must be available. This port is typically 514.

Steps

1. Select **Settings > Alerts**.
2. Select the **Syslog** tab.
3. Select **Add Syslog Servers**.

The Add Syslog Server dialog box opens.

4. Select **Add another syslog server**.
5. Enter information for the syslog server, and then click **Add**.

- **Syslog Server Address** — Enter a fully qualified domain name, an IPv4 address, or an IPv6 address.
- **UDP Port** — Typically, the UDP port for syslog is 514.



You can configure up to five syslog servers.

Results

The syslog server addresses appear in the table.

Delete syslog servers for alerts

You can delete a syslog server so it no longer receives alerts.

Steps

1. Select **Settings > Alerts**.
2. Select the **Syslog** tab.
3. Select a syslog server address, and then click **Remove** from the top right.

The Confirm Delete Syslog Server dialog box opens.

4. Confirm the operation, and then click **Delete**.

Results

The server you removed no longer receives alerts from the event monitor.

FAQs

What if alerts are disabled?

If you want administrators to receive notifications about important events that occur in the storage array, you must configure an alerting method.

For storage arrays managed with SANtricity System Manager, you configure alerts from the Alerts page. Alert notifications can be sent through email, SNMP traps, or syslog messages. In addition, email alerts can be configured from the Initial Setup Wizard.

How do I configure SNMP or syslog alerts?

In addition to email alerts, you can configure alerts to be sent by Simple Network Management Protocol (SNMP) traps or by syslog messages.

To configure SNMP or syslog alerts, go to **Settings > Alerts**.

Why are timestamps inconsistent between the array and alerts?

When the storage array sends alerts, it does not correct for the time zone of the target server or host that receives the alerts. Instead, the storage array uses the local time (GMT) to create the timestamp used for the alert record. As a result, you might see inconsistencies between the timestamps for the storage array and the server or host receiving an alert.

Because the storage array does not correct for time zone when sending alerts, the timestamp on the alerts is GMT-relative, which has a time-zone offset of zero. To calculate a timestamp appropriate to your local time zone, you should determine your hour offset from GMT, and then add or subtract that value from the timestamps.



To avoid this issue, configure Network Time Protocol (NTP) on your storage array controllers. NTP ensures that the controllers are always synced to the correct time.

System: Storage array settings

Concepts

Cache settings and performance

Cache memory is an area of temporary volatile storage on the controller that has a faster access time than the drive media.

With caching, overall I/O performance can be increased as follows:

- Data requested from the host for a read might already be in the cache from a previous operation, thus eliminating the need for drive access.
- Write data is written initially to the cache, which frees the application to continue instead of waiting for the data to be written to the drive.

The default cache settings meet the requirements for most environments, but you can change them if you want.

Storage array cache settings

For all volumes in the storage array, you can specify the following values from the System page:

- **Start value for flushing** — the percentage of unwritten data in the cache that triggers a cache flush (write to disk). When the cache holds the specified start percentage of unwritten data, a flush is triggered. By default, the controller starts flushing the cache when the cache reaches 80 percent full.
- **Cache block size** — the maximum size of each cache block, which is an organizational unit for cache management. The cache block size is by default 8 KiB, but can be set to 4, 8, 16, or 32 KiB. Ideally the cache block size should be set to the predominant I/O size of your applications. File systems or database applications generally use smaller sizes, while a larger size is good for applications requiring large data transfer or sequential I/O.

Volume cache settings

For individual volumes in a storage array, you can specify the following values from the Volumes page (**Storage > Volumes**):

- **Read caching** — The read cache is a buffer that stores data that has been read from the drives. The data for a read operation might already be in the cache from a previous operation, which eliminates the need to access the drives. The data stays in the read cache until it is flushed.
 - **Dynamic read cache prefetch** — Dynamic cache read prefetch allows the controller to copy additional sequential data blocks into the cache while it is reading data blocks from a drive to the cache. This caching increases the chance that future requests for data can be filled from the cache. Dynamic cache read prefetch is important for multimedia applications that use sequential I/O. The rate and amount of

data that is prefetched into cache is self-adjusting based on the rate and request size of the host reads. Random access does not cause data to be prefetched into cache. This feature does not apply when read caching is disabled.

- **Write caching** — The write cache is a buffer that stores data from the host that has not yet been written to the drives. The data stays in the write cache until it is written to the drives. Write caching can increase I/O performance.



Possible loss of data — If you enable the Write caching without batteries option and do not have a universal power supply for protection, you could lose data. In addition, you could lose data if you do not have controller batteries and you enable the Write caching without batteries option.

- **Write caching without batteries** — The write caching without batteries setting lets write caching continue even when the batteries are missing, failed, discharged completely, or not fully charged. Choosing write caching without batteries is not typically recommended, because data might be lost if power is lost. Typically, write caching is turned off temporarily by the controller until the batteries are charged or a failed battery is replaced.
- **Write caching with mirroring** — Write caching with mirroring occurs when the data written to the cache memory of one controller is also written to the cache memory of the other controller. Therefore, if one controller fails, the other can complete all outstanding write operations. Write cache mirroring is available only if write caching is enabled and two controllers are present. Write caching with mirroring is the default setting at volume creation.

Automatic load balancing overview

Automatic load balancing provides improved I/O resource management by reacting dynamically to load changes over time and automatically adjusting volume controller ownership to correct any load imbalance issues when workloads shift across the controllers.

The workload of each controller is continually monitored and, with cooperation from the multipath drivers installed on the hosts, can be automatically brought into balance whenever necessary. When workload is automatically re-balanced across the controllers, the storage administrator is relieved of the burden of manually adjusting volume controller ownership to accommodate load changes on the storage array.

When Automatic Load Balancing is enabled, it performs the following functions:

- Automatically monitors and balances controller resource utilization.
- Automatically adjusts volume controller ownership when needed, thereby optimizing I/O bandwidth between the hosts and the storage array.

Enabling and disabling Automatic Load Balancing

Automatic Load Balancing is enabled by default on all storage arrays.

You might want to disable Automatic Load Balancing on your storage array for the following reasons:

- You do not want to automatically change a particular volume's controller ownership to balance workload.
- You are operating in a highly tuned environment where load distribution is purposefully set up to achieve a specific distribution between the controllers.

Host types that support the Automatic Load Balancing feature

Even though Automatic Load Balancing is enabled at the storage array level, the host type you select for a host or host cluster has a direct influence on how the feature operates.

When balancing the storage array's workload across controllers, the Automatic Load Balancing feature attempts to move volumes that are accessible by both controllers and that are mapped only to a host or host cluster capable of supporting the Automatic Load Balancing feature.

This behavior prevents a host from losing access to a volume due to the load balancing process; however, the presence of volumes mapped to hosts that do not support Automatic Load Balancing affects the storage array's ability to balance workload. For Automatic Load Balancing to balance the workload, the multipath driver must support TPGS and the host type must be included in the following table.



For a host cluster to be considered capable of Automatic Load Balancing, all hosts in that group must be capable of supporting Automatic Load Balancing.

Host type supporting Automatic Load Balancing	With this multipath driver
Windows or Windows Clustered	MPIO with NetApp E-Series DSM
Linux DM-MP (Kernel 3.10 or later)	DM-MP with <code>scsi_dh_alua</code> device handler
VMware	Native Multipathing Plugin (NMP) with <code>VMW_SATP_ALUA</code> Storage Array Type plug-in



With minor exceptions, host types that do not support Automatic Load Balancing continue to operate normally whether or not the feature is enabled. One exception is that if a system has a failover, storage arrays move unmapped or unassigned volumes back to the owning controller when the data path returns. Any volumes that are mapped or assigned to non-Automatic Load Balancing hosts are not moved.

See the [Interoperability Matrix Tool](#) for compatibility information for specific multipath driver, OS level, and controller-drive tray support.

Verifying OS compatibility with the Automatic Load Balancing feature

Verify OS compatibility with the Automatic Load Balancing feature before setting up a new (or migrating an existing) system.

1. Go to the [Interoperability Matrix Tool](#) to find your solution and verify support.

If your system is running Red Hat Enterprise Linux 6 or SUSE Linux Enterprise Server 11, contact technical support.

2. Update and configure the `/etc/multipath.conf` file.
3. Ensure that both `retain_attached_device_handler` and `detect_prio` are set to `yes` for the applicable vendor and product, or use default settings.

Default host operating system type

The default host type is used by the storage array when hosts are initially connected. It defines how the controllers in the storage array work with the host's operating system when volumes are accessed. You can change the host type if there is a need to change how the storage array operates, relative to the hosts that are connected to it.

Generally, you will change the default host type before you connect hosts to the storage array or when you connect additional hosts.

Keep these guidelines in mind:

- If all of the hosts you plan to connect to the storage array have the same operating system (homogenous host environment), then change the host type to match the operating system.
- If there are hosts with different operating systems that you plan to connect to the storage array (heterogeneous host environment), change the host type to match the majority of the hosts' operating systems.

For example, if you are connecting eight different hosts to the storage array, and six of those hosts are running a Windows operating system, you must select Windows as the default host operating system type.

- If the majority of the connected hosts have a mix of different operating systems, change the host type to Factory Default.

For example, if you are connecting eight different hosts to the storage array, and two of those hosts are running a Windows operating system, three are running a VMware operating system, and another three are running a Linux operating system, you must select Factory Default as the default host operating system type.

How to

Edit storage array name

You can change the storage array name that appears in the title bar of SANtricity System Manager.

Steps

1. Select **Settings** > **System**.
2. Under **General**, look for the **Name:** field.

If a storage array name has not been defined, this field displays "Unknown."

3. Click the **Edit** (pencil) icon next to the storage array name.

The field becomes editable.

4. Enter a new name.

A name can contain letters, numbers, and the special characters underscore (_), dash (-), and hash sign (#). A name cannot contain spaces. A name can have a maximum length of 30 characters. The name must be unique.

5. Click the **Save** (check mark) icon.



If you want to close the editable field without making changes, click the **Cancel** (X) icon.

Results

The new name appears in the title bar of SANtricity System Manager.

Turn on storage array locator lights

To find the physical location of a storage array in a cabinet, you can turn on its locator (LED) lights.

Steps

1. Select **Settings > System**.
2. Under **General**, click **Turn on Storage Array Locator Lights**.

The **Turn On Storage Array Locator Lights** dialog box opens, and the corresponding storage array's locator lights turn on.

3. When you have physically located the storage array, return to the dialog box and select **Turn Off**.

Results

The locator lights turn off, and the dialog box closes.

Synchronize storage array clocks

If Network Time Protocol (NTP) is not enabled, you can manually set the clocks on the controllers so they are synchronized with the management client (the system used to run the browser that accesses SANtricity System Manager).

About this task

Synchronization ensures that event time stamps in the event log match time stamps written to the host log files. During the synchronization process, the controllers remain available and operational.



If NTP is enabled in System Manager, do not use this option to synchronize clocks. Instead, NTP automatically synchronizes the clocks with an external host using SNTP (Simple Network Time Protocol).



After synchronization, you might notice that performance statistics are lost or skewed, schedules are impacted (ASUP, snapshots, etc.), and time stamps in log data are skewed. Using NTP avoids this problem.

Steps

1. Select **Settings > System**.
2. Under **General**, click **Synchronize Storage Array Clocks**.

The Synchronize Storage Array Clocks dialog box opens. It shows the current date and time for the controller(s) and the computer used as the management client.



For simplex storage arrays, only one controller is shown.

3. If the times shown in the dialog box do not match, click **Synchronize**.

Results

After synchronization is successful, event time stamps are the same for the event log and host logs.

Save storage array configuration

You can save a storage array's configuration information in a script file to save time setting up additional storage arrays with the same configuration.

Before you begin

The storage array must not be undergoing any operation that changes its logical configuration settings. Examples of these operations include creating or deleting volumes, downloading controller firmware, assigning or modifying hot spare drives, or adding capacity (drives) to a volume group.

About this task

Saving the storage array configuration generates a command line interface (CLI) script that contains storage array settings, volume configuration, host configuration, or host-to-volume assignments for a storage array. You can use this generated CLI script to replicate a configuration to another storage array with the exact same hardware configuration.

However, you should not use this generated CLI script for disaster recovery. Instead, to do a system restore, use the configuration database backup file that you create manually or contact technical support to get this data from the latest Auto-Support data.

This operation *does not* save these settings:

- The life of the battery
- The controller time-of-day
- The nonvolatile static random access memory (NVSRAM) settings
- Any premium features
- The storage array password
- The operating status and states of the hardware components
- The operating status (except Optimal) and states of the volume groups
- Copy services, such as mirroring and volume copy



Risk of application errors — Do not use this option if the storage array is undergoing an operation that will change any logical configuration setting. Examples of these operations include creating or deleting volumes, downloading controller firmware, assigning or modifying hot spare drives, or adding capacity (drives) to a volume group.

Steps

1. Select **Settings > System**.
2. Select **Save Storage Array Configuration**.
3. Select the items of the configuration that you want to save:

- **Storage array settings**
- **Volume configuration**
- **Host configuration**
- **Host-to-volume assignments**



If you select the **Host-to-volume assignments** item, the **Volume configuration** item and the **Host configuration** item are also selected by default. You cannot save **Host-to-volume assignments** without also saving **Volume configuration** and **Host configuration**.

4. Click **Save**.

The file is saved in the Downloads folder for your browser with the name `storage-array-configuration.cfg`.

After you finish

To load the saved storage array configuration onto another storage array, use the SANtricity command line interface (SMcli) with the `-f` option to apply the `.cfg` file.



You can also load a storage array configuration to other storage arrays by using the Unified Manager interface (select **Manage** > **Import Settings**).

Clear storage array configuration

Use the Clear Configuration operation when you want to delete all the pools, volume groups, volumes, host definitions, and host assignments from the storage array.

Before you begin

- Before clearing the storage array configuration, back up the data.

About this task

There are two Clear Storage Array Configuration options:

- **Volume** — Typically, you might use the Volume option to reconfigure a test storage array as a production storage array. For example, you might configure a storage array for testing, and then, when you are done testing, remove the test configuration and set up the storage array for a production environment.
- **Storage Array** — Typically, you might use the Storage Array option to move a storage array to another department or group. For example, you might be using a storage array in Engineering, and now Engineering is getting a new storage array, so you want to move the current storage array to Administration where it will be reconfigured.

The Storage Array option deletes some additional settings.

	Volume	Storage Array
Deletes pools and volume groups	X	X
Deletes volumes	X	X

	Volume	Storage Array
Deletes hosts and host clusters	X	X
Deletes host assignments	X	X
Deletes storage array name		X
Resets storage array cache settings to default		X



Risk of data loss — This operation deletes all data from your storage array. (It does not do a secure erase.) You cannot cancel this operation after it starts. Perform this operation only when instructed to do so by technical support.

Steps

1. Select **Settings > System**.
2. Select **Clear Storage Array Configuration**.
3. In the drop-down list, select either **Volume** or **Storage Array**.
4. **Optional:** If you want to save the configuration (not the data), use the links in the dialog box.
5. Confirm that you want to perform the operation.

Results

- The current configuration is deleted, destroying all existing data on the storage array.
- All drives are unassigned.

Configure login banner

You can create a login banner that is presented to users before they establish sessions in SANtricity System Manager. The banner can include an advisory notice and a consent message.

About this task

When you create a banner, it appears before the login screen in a dialog box.

Steps

1. Select **Settings > System**.
2. Under the **General** section, select **Configure Login Banner**.

The Configure Login Banner dialog box opens.

3. Enter the text you want to appear in the login banner.



Do not use HTML or other markup tags for formatting.

4. Click **Save**.

Results

The next time users log in to System Manager, the text opens in a dialog box. Users must click **OK** to continue to the login screen.

Manage session timeouts

You can configure timeouts in SANtricity System Manager, so that users' inactive sessions are disconnected after a specified time.

About this task

By default, the session timeout for System Manager is 30 minutes. You can adjust that time or you can disable session timeouts altogether.



If Access Management is configured using the Security Assertion Markup Language (SAML) capabilities embedded in the array, a session timeout might occur when the user's SSO session reaches its maximum limit. This might occur before the System Manager session timeout.

Steps

- 1. Select **Settings > System**.
- 2. Under the **General** section, select **Enable/Disable Session Timeout**.

The **Enable/Disable Session Timeout** dialog box opens.

- 3. Use the spinner controls to increase or decrease the time in minutes.

The minimum timeout you can set for System Manager is 15 minutes.



To disable session timeouts, deselect the **Set the length of time...** checkbox.

- 4. Click **Save**.

Change cache settings for the storage array

For all volumes in the storage array, you can adjust the cache memory settings for flushing and block size.

About this task

Cache memory is an area of temporary volatile storage on the controller, which has a faster access time than the drive media. To tune cache performance, you can adjust the following settings:

Cache setting	Description
Start demand cache flushing	Start demand cache flushing specifies the percentage of unwritten data in the cache that triggers a cache flush (write to disk). By default, cache flushing starts when unwritten data reaches 80% capacity. A higher percentage is a good choice for environments with primarily write operations, so new write requests can be processed by cache without having to go to the disk. Lower settings are better in environments where the I/O is erratic (with data bursts), so that the system flushes cache frequently between data bursts. However, a start percentage lower than 80% may cause decreased performance.

Cache setting	Description
Cache block size	The cache block size determines the maximum size of each cache block, which is an organizational unit for cache management. By default, the block size is 32 KiB. System Manager allows the cache block size to be 4, 8, 16, or 32 KiBs. Applications use different block sizes, which have an impact on storage performance. A smaller size is a good choice for file systems or database applications. A larger size is ideal for applications that generate sequential I/O, such as multimedia.

Steps

1. Select **Settings > System**.
2. Scroll down to **Additional Settings**, and then click **Change Cache Settings**.

The Change Cache Settings dialog box opens.

3. Adjust the following values:
 - **Start demand cache flushing** — Choose a percentage that is appropriate for the I/O used in your environment. If you choose a value lower than 80%, you may notice decreased performance.
 - **Cache block size** — Choose a size that is appropriate for your applications.
4. Click **Save**.

Set host connectivity reporting

You can enable host connectivity reporting so the storage array continuously monitors the connection between the controllers and the configured hosts, and then alerts you if the connection is disrupted. This feature is enabled by default.

About this task

If you disable host connectivity reporting, the system no longer monitors connectivity or multipath driver issues with a host connected to the storage array.



Disabling host connectivity reporting also disables automatic load balancing, which monitors and balances controller resource utilization.

Steps

1. Select **Settings > System**.
2. Scroll down to **Additional Settings**, and then click **Enable/Disable Host Connectivity Reporting**.

The text below this option indicates whether it is currently enabled or disabled.

A confirmation dialog box opens.

3. Click **Yes** to continue.

By selecting this option, you toggle the feature between enabled/disabled.

Set automatic load balancing

The Automatic Load Balancing feature ensures that incoming I/O traffic from the hosts is dynamically managed and balanced across both controllers. This feature is enabled by default, but you can disable it from System Manager.

About this task

When Automatic Load Balancing is enabled, it performs the following functions:

- Automatically monitors and balances controller resource utilization.
- Automatically adjusts volume controller ownership when needed, thereby optimizing I/O bandwidth between the hosts and the storage array.

You might want to disable Automatic Load Balancing on your storage array for the following reasons:

- You do not want to automatically change a particular volume's controller ownership to balance workload.
- You are operating in a highly tuned environment where load distribution is purposefully set up to achieve a specific distribution between the controllers.

Steps

1. Select **Settings > System**.
2. Scroll down to **Additional Settings**, and then click **Enable/Disable Automatic Load Balancing**.

The text below this option indicates whether the feature is currently enabled or disabled.

A confirmation dialog box opens.

3. Confirm by clicking **Yes** to continue.

By selecting this option, you toggle the feature between enabled/disabled.



If this feature is moved from disabled to enabled, the Host Connectivity Reporting feature is automatically enabled as well.

Change default host type

Use the Change Default Host Operating System setting to change the default host type at the storage array level. Generally, you will change the default host type before you connect hosts to the storage array or when you connect additional hosts.

About this task

Keep these guidelines in mind:

- If all of the hosts you plan to connect to the storage array have the same operating system (homogenous host environment), then change the host type to match the operating system.
- If there are hosts with different operating systems that you plan to connect to the storage array (heterogeneous host environment), change the host type to match the majority of the hosts' operating systems.

For example, if you are connecting eight different hosts to the storage array, and six of those hosts are running a Windows operating system, you must select Windows as the default host operating system type.

- If the majority of the connected hosts have a mix of different operating systems, change the host type to Factory Default.

For example, if you are connecting eight different hosts to the storage array, and two of those hosts are running a Windows operating system, three are running a VMware operating system, and another three are running a Linux operating system, you must select Factory Default as the default host operating system type.

Steps

1. Select **Settings > System**.
2. Scroll down to **Additional Settings**, and then click **Change Default Host Operating System Type**.
3. Select the host operating system type that you want to use as the default.
4. Click **Change**.

Enable or disable legacy management interface

You can enable or disable the legacy management interface (SYMBOL), which is a method of communication between the storage array and the management client.

About this task

By default, the legacy management interface is on. If you disable it, the storage array and management client will use a more secure method of communication (REST API over https); however, certain tools and tasks might be affected if it is disabled.



For the EF600 storage system, this feature is disabled by default.

The setting affects operations as follows:

- **On** (default) — Required setting for configuring mirroring with the CLI and some other tools, such as the OCI adapter.
- **Off** — Required setting to enforce confidentiality in communications between the storage array and the management client, and to access external tools. Recommended setting when configuring a Directory Server (LDAP).

Steps

1. Select **Settings > System**.
2. Scroll down to **Additional Settings**, and then click **Change Management Interface**.
3. In the dialog box, click **Yes** to continue.

FAQs

What is controller cache?

The controller cache is a physical memory space that streamlines two types of I/O (input/output) operations: between the controllers and hosts, and between the controllers and disks.

For read and write data transfers, the hosts and controllers communicate over high-speed connections. However, communications from the back-end of the controller to the disks is slower, because disks are

relatively slow devices.

When the controller cache receives data, the controller acknowledges to the host applications that it is now holding the data. This way, the host applications do not need to wait for the I/O to be written to disk. Instead, applications can continue operations. The cached data is also readily accessible by server applications, eliminating the need for extra disk reads to access the data.

The controller cache affects the overall performance of the storage array in several ways:

- The cache acts as a buffer, so that host and disk data transfers do not need to be synchronized.
- The data for a read or write operation from the host might be in cache from a previous operation, which eliminates the need to access the disk.
- If write caching is used, the host can send subsequent write commands before the data from a previous write operation is written to disk.
- If cache prefetch is enabled, sequential read access is optimized. Cache prefetch makes a read operation more likely to find its data in the cache, instead of reading the data from disk.



Possible loss of data — If you enable the **Write caching without batteries** option and do not have a universal power supply for protection, you could lose data. In addition, you could lose data if you do not have controller batteries and you enable the **Write caching without batteries** option.

What is cache flushing?

When the amount of unwritten data in the cache reaches a certain level, the controller periodically writes cached data to a drive. This write process is called "flushing."

The controller uses two algorithms for flushing cache: demand-based and age-based. The controller uses a demand-based algorithm until the amount of cached data drops below the cache flush threshold. By default, a flush begins when 80 percent of the cache is in use.

In System Manager, you can set the "Start demand cache flushing" threshold to best support the type of I/O used in your environment. In an environment that is primarily write operations, you should set the "Start demand cache flushing" percentage high to increase the probability that any new write requests can be processed by cache without having to go to the disk. A high percentage setting limits the number of cache flushes so that more data remains in cache, which increases the chance of more cache hits.

In an environment where the I/O is erratic (with data bursts), you can use low cache flushing so that the system flushes cache frequently between data bursts. In a diverse I/O environment that processes a variety of loads, or when the type of loads are unknown, set the threshold at 50 percent as a good middle ground. Be aware that if you choose a start percentage lower than 80 percent, you might see decreased performance because data needed for a host read might not be available. Choosing a lower percentage also increases the number of disk writes necessary to maintain the cache level, which increases system overhead.

The age-based algorithm specifies the period of time during which write data can remain in the cache before it is eligible to be flushed to the disks. The controllers use the age-based algorithm until the cache flush threshold is reached. The default is 10 seconds, but this time period is counted only during periods of inactivity. You cannot modify the flush timing in System Manager; instead, you must use the **Set Storage Array** command in the command-line interface (CLI).



Possible loss of data — If you enable the **Write caching without batteries** option and do not have a universal power supply for protection, you could lose data. In addition, you could lose data if you do not have controller batteries and you enable the **Write caching without batteries** option.

What is cache block size?

The storage array's controller organizes its cache into "blocks," which are chunks of memory that can be 8, 16, 32 KiB in size. All volumes on the storage system share the same cache space; therefore, the volumes can have only one cache block size.

Applications use different block sizes, which can have an impact on storage performance. By default, the block size in System Manager is 32 KiB, but you can set the value to 8, 16, 32 KiBs. A smaller size is a good choice for file systems or database applications. A larger size is a good choice for applications that require large data transfer, sequential I/O, or high bandwidth, such as multimedia.

When should I synchronize storage array clocks?

You should manually synchronize the controller clocks in the storage array if you notice that the time stamps shown in System Manager are not aligned with time stamps shown in your management client (the computer that is accessing System Manager through the browser). This task is only necessary if NTP (Network Time Protocol) is not enabled in System Manager.



We highly recommend that you use an NTP server instead of manually synchronizing the clocks. NTP automatically synchronizes the clocks with an external server using SNTP (Simple Network Time Protocol).

You can check synchronization status from the Synchronize Storage Array Clocks dialog box, which is available from the System page. If the times shown in the dialog box do not match, run a synchronization. You can periodically view this dialog box, which indicates whether the controller clocks' time displays have drifted apart and are no longer synchronized.

What is host connectivity reporting?

When host connectivity reporting is enabled, the storage array continuously monitors the connection between the controllers and the configured hosts, and then alerts you if the connection is disrupted.

Disruptions to the connection might occur if there is a loose, damaged, or missing cable, or another problem with the host. In these situations, the system might open a Recovery Guru message:

- **Host Redundancy Lost** — Opens if either controller cannot communicate with the host.
- **Host Type Incorrect** — Opens if the host's type is incorrectly specified on the storage array, which could result in failover problems.

You might want to disable host connectivity reporting in situations where rebooting a controller might take longer than the connection timeout. Disabling this feature suppresses Recovery Gurus messages.



Disabling host connectivity reporting also disables automatic load balancing, which monitors and balances controller resource use. However, if you re-enable host connectivity reporting, the automatic load balancing feature is not automatically re-enabled.

System: iSCSI settings

Concepts

iSCSI terminology

Learn how the iSCSI terms apply to your storage array.

Term	Description
CHAP	The Challenge Handshake Authentication Protocol (CHAP) method validates the identity of targets and initiators during the initial link. Authentication is based on a shared security key called a <i>CHAPsecret</i> .
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the System Manager functions.
DHCP	Dynamic Host Configuration Protocol (DHCP) is a protocol used on Internet Protocol (IP) networks for dynamically distributing network configuration parameters, such as IP addresses.
IB	InfiniBand (IB) is a communications standard for data transmission between high-performance servers and storage systems.
ICMP PING response	Internet Control Message Protocol (ICMP) is a protocol used by operating systems of networked computers to send messages. ICMP messages determine whether a host is reachable and how long it takes to get packets to and from that host.
IQN	An iSCSI Qualified Name (IQN) identifier is a unique name for an iSCSI initiator or iSCSI target.
iSER	iSCSI Extensions for RDMA (iSER) is a protocol that extends the iSCSI protocol for operation over RDMA transports, such as InfiniBand or Ethernet.
iSNS	Internet Storage Name Service (iSNS) is a protocol that allows automated discovery, management, and configuration of iSCSI and Fibre Channel devices on TCP/IP networks.
MAC address	Media access control identifiers (MAC addresses) are used by Ethernet to distinguish between separate logical channels connecting two ports on the same physical transport network interface.
Management client	A management client is the computer where a browser is installed for accessing System Manager.

Term	Description
MTU	A Maximum Transmission Unit (MTU) is the largest size packet or frame that can be sent in a network.
RDMA	Remote Direct Memory Access (RDMA) is a technology that allows network computers to exchange data in main memory without involving the operating system of either computer.
Unnamed discovery session	When the option for unnamed discovery sessions is enabled, iSCSI initiators are not required to specify the target IQN to retrieve the controller's information.

How to

Configure iSCSI ports

If your controller includes an iSCSI host connection, you can configure the iSCSI port settings from the System page.

Before you begin

- Your controller must include iSCSI ports; otherwise, the iSCSI settings are not available.
- You must know the network speed (the data transfer rate between the ports and the host).



The iSCSI settings and functions only appear if your storage array supports iSCSI.

Steps

1. Select **Settings** > **System**.
2. Under **iSCSI Settings**, select **Configure iSCSI Ports**.




The **Configure iSCSI Ports** option appears only if System Manager detects iSCSI ports on the controller.

3. Select the controller with the iSCSI ports you want to configure.
4. In the drop-down list, select the port you want to configure, and then click **Next**.
5. Select the configuration port settings, and then click **Next**.

To see all port settings, click the **Show more port settings** link on the right of the dialog box.

Field Details

Port Setting	Description
Enable IPv4 / Enable IPv6	<p>Select one or both options to enable support for IPv4 and IPv6 networks.</p> <div> If you want to disable port access, deselect both check boxes.</div>
TCP listening port (Available by clicking Show more port settings .)	<p>If necessary, enter a new port number.</p> <p>The listening port is the TCP port number that the controller uses to listen for iSCSI logins from host iSCSI initiators. The default listening port is 3260. You must enter 3260 or a value between 49152 and 65535.</p>
MTU size (Available by clicking Show more port settings .)	<p>If necessary, enter a new size in bytes for the Maximum Transmission Unit (MTU).</p> <p>The default Maximum Transmission Unit (MTU) size is 1500 bytes per frame. You must enter a value between 1500 and 9000.</p>
Enable ICMP PING responses	<p>Select this option to enable the Internet Control Message Protocol (ICMP). The operating systems of networked computers use this protocol to send messages. These ICMP messages determine whether a host is reachable and how long it takes to get packets to and from that host.</p>

If you selected **Enable IPv4**, a dialog box opens for selecting IPv4 settings after you click **Next**. If you selected **Enable IPv6**, a dialog box opens for selecting IPv6 settings after you click **Next**. If you selected both options, the dialog box for IPv4 settings opens first, and then after you click **Next**, the dialog box for IPv6 settings opens.

6. Configure the IPv4 and/or IPv6 settings, either automatically or manually. To see all port settings, click the **Show more settings** link on the right of the dialog box.

Field Details

Port setting	Description
Automatically obtain configuration	Select this option to obtain the configuration automatically.
Manually specify static configuration	Select this option, and then enter a static address in the fields. (If desired, you can cut and paste addresses into the fields.) For IPv4, include the network subnet mask and gateway. For IPv6, include the routable IP address and router IP address.
Enable VLAN support (Available by clicking Show more settings .)	Select this option to enable a VLAN and enter its ID. A VLAN is a logical network that behaves like it is physically separate from other physical and virtual local area networks (LANs) supported by the same switches, the same routers, or both.
Enable ethernet priority (Available by clicking Show more settings .)	<p>Select this option to enable the parameter that determines the priority of accessing the network. Use the slider to select a priority between 1 (lowest) and 7 (highest).</p> <p>In a shared local area network (LAN) environment, such as Ethernet, many stations might contend for access to the network. Access is on a first-come, first-served basis. Two stations might try to access the network at the same time, which causes both stations to back off and wait before trying again. This process is minimized for switched Ethernet, where only one station is connected to a switch port.</p>

7. Click **Finish**.

Configure iSCSI authentication

For extra security in an iSCSI network, you can set authentication between controllers (targets) and hosts (initiators). System Manager uses the Challenge Handshake Authentication Protocol (CHAP) method, which validates the identity of targets and initiators during the initial link. Authentication is based on a shared security key called a *CHAP secret*.

Before you begin

You can set the CHAP secret for the initiators (iSCSI hosts) either before or after you set the CHAP secret for the targets (controllers). Before you follow the instructions in this task, you should wait until the hosts have made an iSCSI connection first, and then set the CHAP secret on the individual hosts. After the connections are made, the IQN names of the hosts and their CHAP secrets are listed in the dialog box for iSCSI authentication (described in this task), and you do not need to manually enter them.

About this task

You can select one of the following authentication methods:

- **One-way authentication** — Use this setting to allow the controller to authenticate the identity of the iSCSI hosts (uni-directional authentication).

- **Two-way authentication** — Use this setting to allow both the controller and the iSCSI hosts to perform authentication (bi-directional authentication). This setting provides a second level of security by enabling the controller to authenticate the identity of the iSCSI hosts; and in turn, the iSCSI hosts to authenticate the identity of the controller.



The iSCSI settings and functions only display on the Settings page if your storage array supports iSCSI.

Steps

1. Select **Settings > System**.
2. Under **iSCSI settings**, click **Configure Authentication**.

The **Configure Authentication** dialog box appears, which shows the currently set method. It also shows if any hosts have CHAP secrets configured.

3. Select one of the following:
 - **No authentication** — If you do not want the controller to authenticate the identity of iSCSI hosts, select this option and click **Finish**. The dialog box closes, and you are done with configuration.
 - **One-way authentication** — To allow the controller to authenticate the identity of the iSCSI hosts, select this option and click **Next** to display the Configure Target CHAP dialog box.
 - **Two-way authentication** — To allow both the controller and the iSCSI hosts to perform authentication, select this option and click **Next** to display the Configure Target CHAP dialog box.
4. For one-way or two-way authentication, enter or confirm the CHAP secret for the controller (the target). The CHAP secret must be between 12 and 57 printable ASCII characters.



If the CHAP secret for the controller was configured previously, the characters in the field are masked. If necessary, you can replace the existing characters (new characters are not masked).

5. Do one of the following:
 - If you are configuring *one-way* authentication, click **Finish**. The dialog box closes, and you are done with configuration.
 - If you are configuring *two-way* authentication, click **Next** to display the Configure Initiator CHAP dialog box.
6. For two-way authentication, enter or confirm a CHAP secret for any of the iSCSI hosts (the initiators), which can be between 12 and 57 printable ASCII characters. If you do not want to configure two-way authentication for a particular host, leave the **Initiator CHAP Secret** field blank.



If the CHAP secret for a host was configured previously, the characters in the field are masked. If necessary, you can replace the existing characters (new characters are not masked).

7. Click **Finish**.

Results

Authentication occurs during the iSCSI login sequence between the controllers and iSCSI hosts, unless you specified no authentication.

Enable iSCSI discovery settings

You can enable settings related to the discovery of storage devices in an iSCSI network. The Target Discovery Settings allow you to register the storage array's iSCSI information using the Internet Storage Name Service (iSNS) protocol, and also determine whether to allow unnamed discovery sessions.

Before you begin

If the iSNS server uses a static IP address, that address must be available for iSNS registration. Both IPv4 and IPv6 are supported.

About this task

You can enable the following settings related to iSCSI discovery:

- **Enable iSNS server to register a target** — When enabled, the storage array registers its iSCSI Qualified Name (IQN) and port information from the iSNS server. This setting allows iSNS discovery, so that an initiator can retrieve the IQN and port information from the iSNS server.
- **Enable unnamed discovery sessions** — When unnamed discovery sessions are enabled, the initiator (iSCSI host) does not need to provide the IQN of the target (controller) during the login sequence for a discovery-type connection. When disabled, the hosts do need to provide the IQN to establish a discovery-session to the controller. However, the target IQN is always required for a normal (I/O bearing) session. Disabling this setting can prevent unauthorized iSCSI hosts from connecting to the controller using only its IP address.



The iSCSI settings and functions only display on the Settings page if your storage array supports iSCSI.

Steps

1. Select **Settings > System**.
2. Under **iSCSI settings**, click **View/Edit Target Discovery Settings**.

The **Target Discovery Settings** dialog box appears. Below the **Enable iSNS server...** field, the dialog box indicates if the controller is already registered.

3. To register the controller, select **Enable iSNS server to register my target**, and then select one of the following:
 - **Automatically obtain configuration from DHCP server** — Select this option if you want to configure the iSNS server using a Dynamic Host Configuration Protocol (DHCP) server. Be aware that if you use this option, all iSCSI ports on the controller must be configured to use DHCP as well. If necessary, update your controller iSCSI port settings to enable this option.



For the DHCP server to provide the iSNS server address, you must configure the DHCP server to use Option 43 — “Vendor Specific Information.” This option needs to contain the iSNS server IPv4 address in data bytes 0xa-0xd (10-13).

- **Manually specify static configuration** — Select this option if you want to enter a static IP address for the iSNS server. (If desired, you can cut and paste addresses into the fields.) In the field, enter either an IPv4 address or an IPv6 address. If you configured both, IPv4 is the default. Also enter a TCP listening port (use the default of 3205 or enter a value between 49152 and 65535).
4. To allow the storage array to participate in unnamed discovery sessions, select **Enable unnamed**

discovery sessions.

- When enabled, iSCSI initiators are not required to specify the target IQN to retrieve the controller's information.
- When disabled, discovery sessions are prevented unless the initiator provides the target IQN. Disabling unnamed discovery sessions provides added security.

5. Click **Save**.

Results

A progress bar appears as System Manager attempts to register the controller with the iSNS server. This process might take up to five minutes.

View iSCSI statistics packages

You can view data about the iSCSI connections to your storage array.

About this task

System Manager shows these types of iSCSI statistics. All statistics are read-only and cannot be set.

- **Ethernet MAC statistics** — Provides statistics for the media access control (MAC). MAC also provides an addressing mechanism called the physical address or the MAC address. The MAC address is a unique address that is assigned to each network adapter. The MAC address helps deliver data packets to a destination within the subnetwork.
- **Ethernet TCP/IP statistics** — Provides statistics for the TCP/IP, which is the Transmission Control Protocol (TCP) and Internet Protocol (IP) for the iSCSI device. With TCP, applications on networked hosts can create connections to one another, over which they can exchange data in packets. The IP is a data-oriented protocol that communicates data across a packet-switched inter-network. The IPv4 statistics and the IPv6 statistics are shown separately.
- **Local Target/Initiator (Protocol) statistics** — Shows statistics for the iSCSI target, which provides block level access to its storage media, and shows the iSCSI statistics for the storage array when used as an initiator in asynchronous mirroring operations.
- **DCBX Operational States statistics** — Displays the operational states of the various Data Center Bridging Exchange (DCBX) features.
- **LLDP TLV statistics** — Displays the Link Layer Discovery Protocol (LLDP) Type Length Value (TLV) statistics.
- **DCBX TLV statistics** — Displays the information that identifies the storage array host ports in a Data Center Bridging (DCB) environment. This information is shared with network peers for identification and capability purposes.

You can view each of these statistics as raw statistics or as baseline statistics. Raw statistics are all of the statistics that have been gathered since the controllers were started. Baseline statistics are point-in-time statistics that have been gathered since you set the baseline time.

Steps

1. Select **Settings > System**.
2. Select **View iSCSI Statistics Packages**.
3. Click a tab to view the different sets of statistics.
4. **Optional:** To set the baseline, click **Set new baseline**.

Setting the baseline sets a new starting point for the collection of the statistics. The same baseline is used

for all iSCSI statistics.

View iSCSI sessions

You can view detailed information about the iSCSI connections to your storage array. iSCSI sessions can occur with hosts or remote storage arrays in an asynchronous mirror relationship.

Steps

1. Select **Settings** > **System**.
2. Select **View/End iSCSI Sessions**.

A list of the current iSCSI sessions appears.

3. To see additional information about a specific iSCSI session, select a session, and then click **View Details**.

Field Details

Item	Description
Session Identifier (SSID)	A hexadecimal string that identifies a session between an iSCSI initiator and an iSCSI target. The SSID is composed of the ISID and the TPGT.
Initiator Session ID (ISID)	The initiator part of the session identifier. The initiator specifies the ISID during login.
Target Portal Group	The iSCSI target.
Target Portal Group Tag (TPGT)	The target part of the session identifier. A 16-bit numerical identifier for an iSCSI target portal group.
Initiator iSCSI name	The worldwide unique name of the initiator.
Initiator iSCSI label	The user label set in System Manager.
Initiator iSCSI alias	A name that also can be associated with an iSCSI node. The alias allows an organization to associate a user-friendly string with the iSCSI name. However, the alias is not a substitute for the iSCSI name. The initiator iSCSI alias only can be set at the host, not in System Manager
Host	A server that sends input and output to the storage array.
Connection ID (CID)	A unique name for a connection within the session between the initiator and the target. The initiator generates this ID and presents it to the target during login requests. The connection ID is also presented during logouts that close connections.
Ethernet port identifier	The controller port associated with the connection.
Initiator IP address	The IP address of the initiator.
Negotiated login parameters	The parameters that are transacted during the login of the iSCSI session.
Authentication method	The technique to authenticate users who want access to the iSCSI network. Valid values are CHAP and None .
Header digest method	The technique to show possible header values for the iSCSI session. HeaderDigest and DataDigest can be either None or CRC32C . The default value for both is None .
Data digest method	The technique to show possible data values for the iSCSI session. HeaderDigest and DataDigest can be either None or CRC32C . The default value for both is None .

Item	Description
Maximum connections	The greatest number of connections allowed for the iSCSI session. The maximum number of connections can be 1 through 4. The default value is 1 .
Target alias	The label associated with the target.
Initiator alias	The label associated with the initiator.
Target IP address	The IP address of the target for the iSCSI session. DNS names are not supported.
Initial R2T	The initial ready to transfer status. The status can be either Yes or No .
Maximum burst length	The maximum SCSI payload in bytes for this iSCSI session. The maximum burst length can be from 512 to 262,144 (256 KB). The default value is 262,144 (256 KB) .
First burst length	The SCSI payload in bytes for unsolicited data for this iSCSI session. The first burst length can be from 512 to 131,072 (128 KB). The default value is 65,536 (64 KB) .
Default time to wait	The minimum number of seconds to wait before you attempt to make a connection after a connection termination or a connection reset. The default time to wait value can be from 0 to 3600. The default is 2 .
Default time to retain	The maximum number of seconds that connection is still possible following a connection termination or a connection reset. The default time to retain can be from 0 to 3600. The default value is 20 .
Maximum outstanding R2T	The maximum number of "ready to transfers" outstanding for this iSCSI session. The maximum outstanding ready to transfer value can be from 1 to 16. The default is 1 .
Error recovery level	The level of error recovery for this iSCSI session. The error recovery level value is always set to 0 .
Maximum receive data segment length	The maximum amount of data that either the initiator or the target can receive in any iSCSI payload data unit (PDU).
Target name	The official name of the target (not the alias). The target name with the <i>iqn</i> format.
Initiator name	The official name of the initiator (not the alias). The initiator name that uses either the <i>iqn</i> or <i>eui</i> format.

4. **Optional:** To save the report to a file, click **Save**.

The file is saved in the Downloads folder for your browser with the filename `iscsi-session-connections.txt`.

End iSCSI session

You can end an iSCSI session that is no longer needed. iSCSI sessions can occur with hosts or remote storage arrays in an asynchronous mirror relationship.

About this task

You might want to end an iSCSI session for these reasons:

- **Unauthorized access** — If an iSCSI initiator is logged on and should not have access, you can end the iSCSI session to force the iSCSI initiator off the storage array. The iSCSI initiator could have logged on because the None authentication method was available.
- **System downtime** — If you need to take down a storage array and you see that iSCSI initiators are still logged on, you can end the iSCSI sessions to get the iSCSI initiators off the storage array.

Steps

1. Select **Settings > System**.
2. Select **View/End iSCSI Sessions**.

A list of the current iSCSI sessions appears.

3. Select the session that you want to end.
4. Click **End Session**, and confirm that you want to perform the operation.

Configure iSER over InfiniBand ports

If your controller includes an iSER over InfiniBand port, you can configure the network connection to the host.

Before you begin

- Your controller must include an iSER over InfiniBand port; otherwise, the iSER over InfiniBand settings are not available in System Manager.
- You must know the IP address of the host connection.

Steps

1. Select **Settings > System**
2. Under **iSER over InfiniBand settings**, select **Configure iSER over InfiniBand ports**.
3. Click the controller with the iSER over InfiniBand port you want to configure. Click **Next**.
4. In the drop-down list, select the HIC port you want to configure, and then enter the IP address of the host.
5. Click **Finish**.
6. Reset the iSER over InfiniBand port by clicking **Yes**.

View iSER over InfiniBand statistics

If your storage array's controller includes an iSER over InfiniBand port, you can view data about the host connections.

About this task

System Manager shows the following types of iSER over InfiniBand statistics. All statistics are read-only and cannot be set.

- **Local Target (Protocol) statistics** — Provides statistics for the iSER over InfiniBand target, which shows block-level access to its storage media.
- **iSER over InfiniBand Interface statistics** — Provides statistics for all iSER ports on the InfiniBand interface, which includes performance statistics and link error information associated with each switch port.

You can view each of these statistics as raw statistics or as baseline statistics. Raw statistics are all of the statistics that have been gathered since the controllers were started. Baseline statistics are point-in-time statistics that have been gathered since you set the baseline time.

Steps

1. Select **Settings > System**.
2. Select **View iSER over InfiniBand Statistics**.
3. Click a tab to view the different sets of statistics.
4. **Optional:** To set the baseline, click **Set new baseline**.

Setting the baseline sets a new starting point for the collection of the statistics. The same baseline is used for all iSER over InfiniBand statistics.

FAQs

What happens when I use an iSNS server for registration?

When Internet Storage Name Service (iSNS) server information is used, the hosts (initiators) can be configured to query the iSNS server to retrieve information from the target (controllers).

This registration provides the iSNS server with the controller's iSCSI Qualified Name (IQN) and port information, and allows for queries between the initiators (iSCSI hosts) and targets (controllers).

Which registration methods are automatically supported for iSCSI?

The iSCSI implementation supports either the Internet Storage Name Service (iSNS) discovery method or the use of the Send Targets command.

The iSNS method allows for iSNS discovery between the initiators (iSCSI hosts) and targets (the controllers). You register the target controller to provide the iSNS server with the controller's iSCSI Qualified Name (IQN) and port information.

If you do not configure iSNS, the iSCSI host can send the Send Targets command during an iSCSI discovery session. In response, the controller returns the port information (for example, the Target IQN, port IP address, listening port, and Target Port Group). This discovery method is not required if you use iSNS, because the host initiator can retrieve the target IPs from the iSNS server.

How do I interpret iSER over InfiniBand statistics?


The View iSER over InfiniBand Statistics dialog box displays local target (protocol) statistics and iSER over InfiniBand (IB) interface statistics. All statistics are read-only, and cannot be set.

- **Local Target (Protocol) statistics** — Provides statistics for the iSER over InfiniBand target, which shows block-level access to its storage media.
- **iSER over InfiniBand Interface statistics** — Provides statistics for all iSER over InfiniBand ports on the InfiniBand interface, which includes performance statistics and link error information associated with each switch port.

You can view each of these statistics as raw statistics or as baseline statistics. Raw statistics are all of the statistics that have been gathered since the controllers were started. Baseline statistics are point-in-time statistics that have been gathered since you set the baseline time.

What else do I need to do to configure or diagnose iSER over InfiniBand?

The following table lists the System Manager functions that you can use to configure and manage iSER over InfiniBand sessions.



The iSER over InfiniBand settings are available only if your storage array’s controller includes an iSER over InfiniBand host management port.

Configure and diagnose iSER over InfiniBand

Action	Location
Configure iSER over InfiniBand ports	<div>1. Select Hardware.</div> <div>2. Select Show back of shelf.</div> <div>3. Select a controller.</div> <div>4. Select Configure iSER over InfiniBand ports.</div> <div>or</div> <div>1. Select Settings > System.</div> <div>2. Scroll down to iSER over InfiniBand settings, and then select Configure iSER over InfiniBand Ports.</div>
View iSER over InfiniBand statistics	<div>1. Select Settings > System.</div> <div>2. Scroll down to iSER over InfiniBand settings, and then select View iSER over InfiniBand Statistics.</div>

System: NVMe settings

Concepts

Some controllers include a port for implementing NVMe (Non-Volatile Memory Express) over fabrics. NVMe allows for high-performance communication between hosts and the storage array.

What is NVMe?

NVM stands for "Non-Volatile Memory" and is persistent memory used in many types of storage devices. NVMe (NVM Express) is a standardized interface or protocol designed specifically for high-performance multi-queue communication with NVM devices.

What is NVMe over Fabrics?

NVMe over Fabrics (NVMe-oF) is a technology specification that enables NVMe message-based commands and data to transfer between a host computer and storage over a network. An NVMe storage array (called a *subsystem*) can be accessed by a host using a fabric. NVMe commands are enabled and encapsulated in transport abstraction layers on both the host side and the subsystem side. This extends the high performance NVMe interface end-to-end from the host to the storage and standardizes and simplifies the command set.

NVMe-oF storage is presented to a host as a local block storage device. The volume (called a *namespace*) can be mounted to a file system as with any other block storage device. You can use the REST API, the SMcli, or SANtricity System Manager to provision your storage as needed.

What is an NVMe Qualified Name (NQN)?

The NVMe Qualified Name (NQN) is used to identify the remote storage target. The NVMe qualified name for the storage array is always assigned by the subsystem and may not be modified. There is only one NVMe Qualified Name for the entire array. The NVMe Qualified Name is limited to 223 characters in length. You can compare it to an iSCSI Qualified Name.

What is a namespace and a namespace ID?

A namespace is the equivalent of a logical unit in SCSI, which relates to a volume in the array. The namespace ID (NSID) is equivalent to a logical unit number (LUN) in SCSI. You create the NSID at namespace creation time, and can set it to a value between 1 and 255.

What is an NVMe controller?

Similar to a SCSI I_T nexus, which represents the path from the host's initiator to the storage system's target, an NVMe controller created during the host connection process provides an access path between a host and the namespaces in the storage array. An NQN for the host plus a host port identifier uniquely identify an NVMe controller. While an NVMe controller can only be associated with a single host, it can access multiple namespaces.

You configure which hosts can access which namespaces and set the namespace ID for the host using SANtricity System Manager. Then, when the NVMe controller is created, the list of namespace IDs accessible by the NVMe controller is created and used to configure the permissible connections.

NVMe terminology

Learn how the NVMe terms apply to your storage array.

Term	Description
InfiniBand	InfiniBand (IB) is a communications standard for data transmission between high-performance servers and storage systems.
Namespace	A namespace is NVM storage that is formatted for block access. It is analogous to a logical unit in SCSI, which relates to a volume in the storage array.
Namespace ID	The namespace ID is the NVMe controller's unique identifier for the namespace, and can be set to a value between 1 and 255. It is analogous to a logical unit number (LUN) in SCSI.
NQN	NVMe Qualified Name (NQN) is used to identify the remote storage target (the storage array).
NVM	Non-Volatile Memory (NVM) is persistent memory used in many types of storage devices.
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.
NVMe-oF	Non-Volatile Memory Express over Fabrics (NVMe-oF) is a specification that enables NVMe commands and data to transfer over a network between a host and storage.
NVMe controller	An NVMe controller is created during the host connection process. It provides an access path between a host and the namespaces in the storage array.
NVMe queue	A queue is used for passing commands and messages over the NVMe interface.
NVMe subsystem	The storage array with an NVMe host connection.
RDMA	Remote direct memory access (RDMA) enables more direct data movement in and out of a server by implementing a transport protocol in the network interface card (NIC) hardware.
RoCE	RDMA over Converged Ethernet (RoCE) is a network protocol that allows remote direct memory access (RDMA) over an Ethernet network.
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.

How tos

Configure NVMe over InfiniBand ports

If your controller includes an NVMe over InfiniBand connection, you can configure the NVMe port settings from the System page.

Before you begin

- Your controller must include an NVMe over InfiniBand host port; otherwise, the NVMe over InfiniBand settings are not available in System Manager.
- You must know the IP address of the host connection.



The NVMe over InfiniBand settings and functions appear only if your storage array's controller includes an NVMe over InfiniBand port.

Steps

1. Select **Settings > System**.
2. Under **NVMe over InfiniBand settings**, select **Configure NVMe over InfiniBand ports**.
3. Select the controller with the NVMe over InfiniBand port you want to configure. Click **Next**.
4. Select the HIC port you want to configure from the drop-down list, and then enter the IP address.

If you are configuring an EF600 storage array with a 200Gb-capable HIC, this dialog box displays two IP Address fields, one for a physical port (external) and one for a virtual port (internal). You should assign a unique IP address for both ports. These settings allow the host to establish a path between each port, and for the HIC to achieve maximum performance. If you do not assign an IP address to the virtual port, the HIC will run at approximately half its capable speed.

5. Click **Finish**.
6. Reset the NVMe over InfiniBand port by clicking **Yes**.

Configure NVMe over RoCE ports

If your controller includes a connection for NVMe over RoCE (RDMA over Converged Ethernet), you can configure the NVMe port settings from the System page.

Before you begin


- Your controller must include an NVMe over RoCE host port; otherwise, the NVMe over RoCE settings are not available in System Manager.
- You must know the IP address of the host connection.

Steps

1. Select **Settings > System**.
2. Under **NVMe over ROCE settings**, select **Configure NVMe over ROCE ports**.
3. Select the controller with the NVMe over RoCE port you want to configure. Click **Next**.
4. Select the HIC port you want to configure from the drop-down list. Click **Next**.
5. Configure the port settings.

To see all port settings, click the **Show more port settings** link on the right of the dialog box.

Field Details

Port Setting	Description
Configured ethernet port speed	Select the speed that matches the speed capability of the SFP on the port.
Enable IPv4 / Enable IPv6	<div><div>Select one or both options to enable support for IPv4 and IPv6 networks.</div><div> If you want to disable port access, deselect both check boxes.</div></div>
MTU size (Available by clicking Show more port settings.)	<div>If necessary, enter a new size in bytes for the Maximum Transmission Unit (MTU).</div> <div>The default Maximum Transmission Unit (MTU) size is 1500 bytes per frame. You must enter a value between 1500 and 9000.</div>

If you selected **Enable IPv4**, a dialog box opens for selecting IPv4 settings after you click **Next**. If you selected **Enable IPv6**, a dialog box opens for selecting IPv6 settings after you click **Next**. If you selected both options, the dialog box for IPv4 settings opens first, and then after you click **Next**, the dialog box for IPv6 settings opens.

6. Configure the IPv4 and/or IPv6 settings, either automatically or manually.

Field Details

Port setting	Description
Automatically obtain configuration	Select this option to obtain the configuration automatically.
Manually specify static configuration	Select this option, and then enter a static address in the fields. (If desired, you can cut and paste addresses into the fields.) For IPv4, include the network subnet mask and gateway. For IPv6, include the routable IP address and router IP address. If you are configuring an EF600 storage array with a 200Gb-capable HIC, this dialog box displays two sets of fields for network parameters, one for a physical port (external) and one for a virtual port (internal). You should assign unique parameters for both ports. These settings allow the host to establish a path between each port, and for the HIC to achieve maximum performance. If you do not assign an IP address to the virtual port, the HIC will run at approximately half its capable speed.

7. Click **Finish**.

View NVMe over Fabrics statistics

You can view data about the NVMe over Fabrics connections to your storage array.

About this task

System Manager shows these types of NVMe over Fabrics statistics. All statistics are read-only and cannot be set.

- **NVMe Subsystem statistics** — Shows statistics for the NVMe controller and its queue. The NVMe controller provides an access path between a host and the namespaces in the storage array. You can review the NVMe subsystem statistics for such items as connection failures, resets, and shutdowns.
- **RDMA Interface statistics** — Provides statistics for all NVMe over Fabrics ports on the RDMA interface, which includes performance statistics and link error information associated with each switch port. This tab only appears when NVMe over Fabrics ports are available.

You can view each of these statistics as raw statistics or as baseline statistics. Raw statistics are all of the statistics that have been gathered since the controllers were started. Baseline statistics are point-in-time statistics that have been gathered since you set the baseline time.

Steps

1. Select **Settings > System**.
2. Select **View NVMe over Fabrics Statistics**.
3. **Optional:** To set the baseline, click **Set new baseline**.

Setting the baseline sets a new starting point for the collection of the statistics. The same baseline is used for all NVMe statistics.

FAQs

How do I interpret NVMe over Fabrics statistics?

The View NVMe over Fabrics Statistics dialog box displays statistics for the NVMe subsystem and the RDMA interface. All statistics are read-only, and cannot be set.

- **NVMe Subsystem statistics** — Shows statistics for the NVMe controller and its queue. The NVMe controller provides an access path between a host and the namespaces in the storage array. You can review the NVMe subsystem statistics for such items as connection failures, resets, and shutdowns. For more information about these statistics, click **View legend for table headings**.
- **RDMA Interface statistics** — Provides statistics for all NVMe over Fabrics ports on the RDMA interface, which includes performance statistics and link error information associated with each switch port. This tab only appears when NVMe over Fabrics ports are available. For more information about the statistics, click **View legend for table headings**.

You can view each of these statistics as raw statistics or as baseline statistics. Raw statistics are all of the statistics that have been gathered since the controllers were started. Baseline statistics are point-in-time statistics that have been gathered since you set the baseline time.

What else do I need to do to configure or diagnose NVMe over InfiniBand?

The following table lists the System Manager functions that you can use to configure and manage NVMe over InfiniBand sessions.



The NVMe over InfiniBand settings are available only if your storage array's controller includes an NVMe over InfiniBand port.

Configure and diagnose NVMe over InfiniBand

Action	Location
Configure NVMe over InfiniBand ports	<ol style="list-style-type: none">1. Select Hardware.2. Select Show back of shelf.3. Select a controller.4. Select Configure NVMe over InfiniBand ports. <p>or</p> <ol style="list-style-type: none">1. Select Settings > System.2. Scroll down to NVMe over InfiniBand settings, and then select Configure NVMe over InfiniBand Ports.
View NVMe over InfiniBand statistics	<ol style="list-style-type: none">1. Select Settings > System.2. Scroll down to NVMe over InfiniBand settings, and then select View NVMe over Fabrics Statistics.

What else do I need to do to configure or diagnose NVMe over RoCE?

You can configure and manage NVMe over RoCE from the Hardware and Settings pages.



The NVMe over RoCE settings are available only if your storage array's controller includes an NVMe over RoCE port.

Configure and diagnose NVMe over RoCE

Action	Location
Configure NVMe over RoCE ports	<ol style="list-style-type: none">1. Select Hardware.2. Select Show back of shelf.3. Select a controller.4. Select Configure NVMe over RoCE ports. <p>or</p> <ol style="list-style-type: none">1. Select Settings > System.2. Scroll down to NVMe over RoCE settings, and then select Configure NVMe over RoCE Ports.

Action	Location
View NVMe over Fabrics statistics	<ol style="list-style-type: none"> 1. Select Settings › System. 2. Scroll down to NVMe over RoCE settings, and then select View NVMe over Fabrics Statistics.

Why are there two IP addresses for one physical port?

The EF600 storage array can include two HICs — one external and one internal.

In this configuration, the external HIC is connected to an internal, auxiliary HIC. Each physical port that you can access from the external HIC has an associated virtual port from the internal HIC.

To achieve maximum 200Gb performance, you must assign a unique IP address for both the physical and virtual ports so the host can establish connections to each. If you do not assign an IP address to the virtual port, the HIC will run at approximately half its capable speed.

Why are there two sets of parameters for one physical port?

The EF600 storage array can include two HICs — one external and one internal.

In this configuration, the external HIC is connected to an internal, auxiliary HIC. Each physical port that you can access from the external HIC has an associated virtual port from the internal HIC.

To achieve maximum 200Gb performance, you must assign parameters for both the physical and virtual ports so the host can establish connections to each. If you do not assign parameters to the virtual port, the HIC will run at approximately half its capable speed.

System: Add-on features

Concepts

How add-on features work

Add-ons are features that are not included in the standard configuration of System Manager and might require a key to enable. An add-on feature can be either a single premium feature or a bundled feature pack.

The following steps provide an overview for enabling a premium feature or feature pack:

1. Obtain the following information:
 - Chassis serial number and the Feature Enable Identifier, which identify the storage array for the feature to be installed. These items are available in System Manager.
 - Feature Activation Code, which is available from the Support site when you purchase the feature.
2. Obtain the feature key by contacting your storage provider or by accessing the Premium Feature Activation site. Provide the chassis serial number, enable identifier, and feature code for activation.
3. Using System Manager, enable the premium feature or feature pack using the feature key file.

Add-on feature terminology

Learn how the add-on feature terms apply to your storage array.

Term	Description
Feature Enable Identifier	A Feature Enable Identifier is a unique string that identifies the specific storage array. This identifier ensures that when you obtain the premium feature, it is associated with only that particular storage array. This string is displayed under Add-Ons on the System page.
Feature key file	A feature key file is a file you receive for unlocking and enabling a premium feature or feature pack.
Feature pack	A feature pack is a bundle that changes storage array attributes (for example, changing the protocol from Fibre Channel to iSCSI). Feature packs require a special key to enable them.
Premium feature	A premium feature is an extra option that requires a key to enable it. It is not included in the standard configuration of System Manager.

How tos

Obtain a feature key file

To enable a premium feature or feature pack on your storage array, you must first obtain a feature key file. A key is associated with only one storage array.

About this task

This task describes how to gather required information for the feature, and then send a request for a feature key file. Required information includes:

- Chassis serial number
- Feature Enable Identifier
- Feature Activation Code

Steps

1. In System Manager, locate and record the chassis serial number. You can view this serial number by hovering your mouse over the Support Center tile.
2. In System Manager, locate the Feature Enable Identifier. Go to **Settings > System**, and then scroll down to **Add-ons**. Look for the **Feature Enable Identifier**. Record the number for the Feature Enable Identifier.
3. Locate and record the code for feature activation. For features packs, this code is provided in the appropriate instructions for performing the conversion.

NetApp instructions are available from [NetApp E-Series Systems Documentation Center](#).

For premium features, you can access the activation code from the Support site, as follows:

- a. Log in to [NetApp Support](#).

- b. Go to **Software Licenses** for your product.
 - c. Enter the serial number for the storage array chassis, and then click **Go**.
 - d. Look for the Feature Activation Codes in the **License Key** column.
 - e. Record the Feature Activation Code for the feature you want.
4. Request a feature key file by sending an email or a text document to your storage supplier with the following information: chassis serial number, the enable identifier, and the code for feature activation.

You can also go to [NetApp License Activation: Storage Array Premium Feature Activation](#) and enter the required information to obtain the feature or feature pack. (The instructions on this site are for premium features, not feature packs.)

After you finish

When you have a feature key file, you can enable the premium feature or feature pack.

Enable a premium feature

A premium feature is an extra option that requires a key to enable.

Before you begin

- You have obtained a feature key. If necessary, contact technical support for a key.
- You have loaded the key file on the management client (the system with a browser for accessing System Manager).

About this task

This task describes how to use System Manager to enable a premium feature.



If you want to disable a premium feature, you must use the Disable Storage Array Feature command (`disable storageArray (featurePack | feature=featureAttributeList)`) in the Command Line Interface (CLI).

Steps

1. Select **Settings > System**.
2. Under **Add-ons**, select **Enable Premium Feature**.

The Enable a Premium Feature dialog box opens.

3. Click **Browse**, and then select the key file.

The file name is displayed in the dialog box.

4. Click **Enable**.

Enable feature pack

A feature pack is a bundle that changes storage array attributes (for example, changing the protocol from Fibre Channel to iSCSI). Feature packs require a special key for enablement.

Before you begin

- You have followed the appropriate instructions for performing the conversion and for preparing your system for the new storage array attributes.



Conversion instructions are available from [NetApp E-Series Systems Documentation Center](#).

- The storage array is offline, so no hosts or applications are accessing it.
- All data is backed up.
- You have obtained a feature pack file.

The feature pack file is loaded on the management client (the system with a browser for accessing System Manager).



You must schedule a downtime maintenance window and stop all I/O operations between the host and controllers. In addition, be aware that you cannot access data on the storage array until you have successfully completed the conversion.

About this task

This task describes how to use System Manager to enable a feature pack. When you are done, you must restart the storage array.

Steps

1. Select **Settings** > **System**.
2. Under **Add-ons**, select **Change Feature Pack**.
3. Click **Browse**, and then select the key file.

The file name is displayed in the dialog box.

4. Type **CHANGE** in the field.
5. Click **Change**.

The feature pack migration begins and the controllers reboot. Unwritten cache data is deleted, which ensures no I/O activity. Both controllers automatically reboot for the new feature pack to take effect. The storage array returns to a responsive state after the reboot is complete.

Download the command line interface (CLI)

From System Manager you can download the command line interface (CLI) package. The CLI provides a text-based method for configuring and monitoring storage arrays. It communicates via https and uses the same syntax as the CLI available in the externally installed management software package. No key is required to download the CLI.

Before you begin

- A Java Runtime Environment (JRE), version 8 and above, must be available on the management system where you plan to run the CLI commands.

Steps

1. Select **Settings** > **System**.

2. Under **Add-ons**, select **Command Line Interface**.

The ZIP package downloads to the browser.

3. Save the ZIP file to the management system where you plan to run CLI commands for the storage array, and then extract the file.

You can now run CLI commands from an operating system prompt, such as the DOS C: prompt. A CLI command reference is available from the Help menu at the top right of the System Manager user interface.

System: Security key management

Concepts

How the Drive Security feature works

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.

How to implement Drive Security

To implement Drive Security, you perform the following steps.

1. Equip your storage array with secure-capable drives, either FDE drives or FIPS drives. (For volumes that require FIPS support, use only FIPS drives. Mixing FIPS and FDE drives in a volume group or pool will result in all drives being treated as FDE drives. Also, an FDE drive cannot be added to or used as a spare in an all-FIPS volume group or pool.)
2. Create a security key, which is a string of characters that is shared by the controller and drives for read/write access. You can create either an internal key from the controller's persistent memory or an external key from a key management server. For external key management, authentication must be established with the key management server.
3. Enable Drive Security for pools and volume groups:
 - Create a pool or volume group (look for **Yes** in the **Secure-capable** column in the Candidates table).
 - Select a pool or volume group when you create a new volume (look for **Yes** next to **Secure-capable** in the pool and volume group Candidates table).

How Drive Security works at the drive level

A secure-capable drive, either FDE or FIPS, encrypts data during writes and decrypts data during reads. This encryption and decryption does not affect the performance or user workflow. Each drive has its own unique encryption key, which can never be transferred from the drive.

The Drive Security feature provides an extra layer of protection with secure-capable drives. When volume groups or pools on these drives are selected for Drive Security, the drives look for a security key before allowing access to the data. You can enable Drive Security for pools and volume groups at any time, without affecting existing data on the drive. However, you cannot disable Drive Security without erasing all data on the drive.

How Drive Security works at the storage array level

With the Drive Security feature, you create a security key that is shared between the secure-enabled drives and controllers in a storage array. Whenever power to the drives is turned off and on, the secure-enabled drives change to a Security Locked state until the controller applies the security key.

If a secure-enabled drive is removed from the storage array and re-installed in a different storage array, the drive will be in a Security Locked state. The re-located drive looks for the security key before it makes the data accessible again. To unlock the data, you apply the security key from the source storage array. After a successful unlock process, the re-located drive will then use the security key already stored in the target storage array, and the imported security key file is no longer needed.



For internal key management, the actual security key is stored on the controller in a non-accessible location. It is not in human-readable format, nor is it user-accessible.

How Drive Security works at the volume level

When you create a pool or volume group from secure-capable drives, you can also enable Drive Security for those pools or volume groups. The Drive Security option makes the drives and associated volume groups and pools *secure-enabled*.

Keep the following guidelines in mind before creating secure-enabled volume groups and pools:

- Volume groups and pools must be comprised entirely of secure-capable drives. (For volumes that require FIPS support, use only FIPS drives. Mixing FIPS and FDE drives in a volume group or pool will result in all drives being treated as FDE drives. Also, an FDE drive cannot be added to or used as a spare in an all-FIPS volume group or pool.)
- Volume groups and pools must be in an optimal state.

How security key management works

When you implement the Drive Security feature, the secure-enabled drives (FIPS or FDE) require a security key for data access. A security key is a string of characters that is shared between these types of drives and the controllers in a storage array.

Whenever power to the drives is turned off and on, the secure-enabled drives change to a Security Locked state until the controller applies the security key. If a secure-enabled drive is removed from the storage array, the drive's data is locked. When the drive is re-installed in a different storage array, it looks for the security key before it makes the data accessible again. To unlock the data, you must apply the original security key.

You can create and manage security keys using one of the following methods:

- Internal key management on the controller's persistent memory.
- External key management on an external key management server.

Internal key management

Internal keys are maintained on the controller's persistent memory. To implement internal key management, you perform the following steps:

1. Install secure-capable drives in the storage array. These drives can be Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.

2. Make sure the Drive Security feature is enabled. If necessary, contact your storage vendor for instructions on enabling the Drive Security feature.
3. Create an internal security key, which involves defining an identifier and a pass phrase. The identifier is a string that is associated with the security key, and is stored on the controller and on all drives associated with the key. The pass phrase is used to encrypt the security key for backup purposes. To create an internal key, go to **Settings > System > Security key management > Create Internal Key**.

The security key is stored on the controller in a non-accessible location. You can then create secure-enabled volume groups or pools, or you can enable security on existing volume groups and pools.

External key management

External keys are maintained on a separate key management server, using a Key Management Interoperability Protocol (KMIP). To implement external key management, you perform the following steps:


1. Install secure-capable drives in the storage array. These drives can be Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.
2. Make sure the Drive Security feature is enabled. If necessary, contact your storage vendor for instructions on enabling the Drive Security feature.
3. Complete and download a client Certificate Signing Request (CSR) for authentication between the storage array and the key management server. Go to **Settings > Certificates > Key Management > Complete CSR**.
4. Create and download a client certificate from the key management server using the downloaded CSR file.
5. Ensure that the client certificate and a copy of the certificate for the key management server are available on your local host.
6. Create an external key, which involves defining the IP address of the key management server and the port number used for KMIP communications. During this process, you also load certificate files. To create an external key, go to **Settings > System > Security key management > Create External Key**.

The system connects to the key management server with the credentials you entered. You can then create secure-enabled volume groups or pools, or you can enable security on existing volume groups and pools.

Drive Security terminology

Learn how the Drive Security terms apply to your storage array.

Term	Description
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.
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.

Term	Description
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.
Management client	A local system (computer, tablet, etc.) that includes a browser for accessing System Manager.
Pass phrase	<p>The pass phrase is used to encrypt the security key for backup purposes. The same pass phrase used to encrypt the security key must be provided when the backed up security key is imported as the result of a drive migration or head swap. A pass phrase can have between 8 and 32 characters.</p> <div>  <p>The pass phrase for Drive Security is independent from the storage array's Administrator password.</p> </div>
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 <i>secure-capable</i> 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 <i>secure-enabled</i> .
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 <i>secure-capable</i> drives, the drives become <i>secure-enabled</i> . 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.
Security key	<p>A security key is a string of characters that is shared between the secure-enabled drives and controllers in a storage array. Whenever power to the drives is turned off and on, the secure-enabled drives change to a Security Locked state until the controller applies the security key. If a secure-enabled drive is removed from the storage array, the drive's data is locked. When the drive is re-installed in a different storage array, it looks for the security key before it makes the data accessible again. To unlock the data, you must apply the original security key. You can create and manage security keys using one of the following methods:</p> <ul style="list-style-type: none"> • Internal key management — Create and maintain security keys on the controller's persistent memory. • External key management — Create and maintain security keys on an external key management server.
Security key identifier	The security key identifier is a string that is associated with the security key during key creation. The identifier is stored on the controller and on all drives associated with the security key.

How to

Create internal security key

To use the Drive Security feature, you can create an internal security key that is shared by the controllers and secure-capable drives in the storage array. Internal keys are maintained on the controller's persistent memory.

Before you begin

- Secure-capable drives must be installed in the storage array. These drives can be Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.
- The Drive Security feature must be enabled. Otherwise, a Cannot Create Security Key dialog box opens during this task. If necessary, contact your storage vendor for instructions on enabling the Drive Security feature.



If both FDE and FIPS drives are installed in the storage array, they all share the same security key.

About this task

In this task, you define an identifier and a pass phrase to associate with the internal security key.



The pass phrase for Drive Security is independent from the storage array's Administrator password.

Steps

1. Select **Settings** > **System**.
2. Under **Security key management**, select **Create Internal Key**.

If you have not yet generated a security key, the **Create Security Key** dialog box opens.

3. Enter information in the following fields:

- **Define a security key identifier** — You can either accept the default value (storage array name and time stamp, which is generated by the controller firmware) or enter your own value. You can enter up to 189 alphanumeric characters without spaces, punctuation, or symbols.



Additional characters are generated automatically, appended to both ends of the string you enter. The generated characters ensure that the identifier is unique.

- **Define a pass phrase/Re-enter pass phrase** — Enter and confirm a pass phrase. The value can have between 8 and 32 characters, and must include each of the following:
 - An uppercase letter (one or more). Keep in mind that the pass phrase is case sensitive.
 - A number (one or more).
 - A non-alphanumeric character, such as !, *, @ (one or more).



Be sure to record your entries for later use. If you need to move a secure-enabled drive from the storage array, you must know the identifier and pass phrase to unlock drive data.

4. Click **Create**.

The security key is stored on the controller in a non-accessible location. Along with the actual key, there is an encrypted key file that is downloaded from your browser.



The path for the downloaded file might depend on the default download location of your browser.

5. Record your key identifier, pass phrase, and the location of the downloaded key file, and then click **Close**.

Results

You can now create secure-enabled volume groups or pools, or you can enable security on existing volume groups and pools.



Whenever power to the drives is turned off and then on again, all the secure-enabled drives change to a Security Locked state. In this state, the data is inaccessible until the controller applies the correct security key during drive initialization. If someone physically removes a locked drive and installs it in another system, the Security Locked state prevents unauthorized access to its data.

After you finish

You should validate the security key to make sure the key file is not corrupted.

Create external security key

To use the Drive Security feature with a key management server, you must create an external key that is shared by the key management server and the secure-capable drives in the storage array.

Before you begin

- Secure-capable drives must be installed in the array. These drives can be Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.



If both FDE and FIPS drives are installed in the storage array, they all share the same security key.

- The Drive Security feature must be enabled. Otherwise, a **Cannot Create Security Key** dialog box opens during this task. If necessary, contact your storage vendor for instructions on enabling the Drive Security feature.
- The client and server certificates are available on your local host so the storage array and key management server can authenticate each other. The client certificate validates the controllers, while the server certificate validates the key management server.

About this task

In this task, you define the IP address of the key management server and the port number it uses, and then load certificates for external key management.

Steps

1. Select **Settings > System**.
2. Under **Security key management**, select **Create External Key**.



If internal key management is currently configured, a dialog box opens and asks you to confirm that you want to switch to external key management.

The **Create External Security Key** dialog box opens.

3. Under **Connect to Key Server**, enter information in the following fields:

- **Key management server address** — Enter the fully qualified domain name or the IP address (IPv4 or IPv6) of the server used for key management.
- **Key management port number** — Enter the port number used for the Key Management Interoperability Protocol (KMIP) communications. The most common port number used for key management server communications is 5696.
- **Select client certificate** — Click the first **Browse** button to select the certificate file for the storage array's controllers.
- **Select key management server's server certificate** — Click the second **Browse** button to select the certificate file for the key management server.

4. Click **Next**.

5. Under **Create/Backup Key**, enter information in the following field:

- **Define a pass phrase/Re-enter pass phrase** — Enter and confirm a pass phrase. The value can have between 8 and 32 characters, and must include each of the following:
 - An uppercase letter (one or more). Keep in mind that the pass phrase is case sensitive.
 - A number (one or more).
 - A non-alphanumeric character, such as !, *, @ (one or more).



Be sure to record your entries for later use. If you need to move a secure-enabled drive from the storage array, you must know the pass phrase to unlock drive data.

6. Click **Finish**.

The system connects to the key management server with the credentials you entered. A copy of the security key is then stored on your local system.



The path for the downloaded file might depend on the default download location of your browser.

7. Record your pass phrase and the location of the downloaded key file, and then click **Close**.

The page displays the following message with additional links for external key management:

Current key management method: External

8. Test the connection between the storage array and the key management server by selecting **Test Communication**.

Test results display in the dialog box.

Results

When external key management is enabled, you can create secure-enabled volume groups or pools, or you can enable security on existing volume groups and pools.



Whenever power to the drives is turned off and then on again, all the secure-enabled drives change to a Security Locked state. In this state, the data is inaccessible until the controller applies the correct security key during drive initialization. If someone physically removes a locked drive and installs it in another system, the Security Locked state prevents unauthorized access to its data.

After you finish

- You should validate the security key to make sure the key file is not corrupted.

Change security key

At any time, you can replace a security key with a new key. You might need to change a security key in cases where you have a potential security breach at your company and want to make sure unauthorized personnel cannot access the drives' data.

Before you begin

A security key already exists.

About this task

This task describes how to change a security key and replace it with a new one. After this process, the old key is invalidated.

Steps

1. Select **Settings > System**.
2. Under **Security key management**, select **Change Key**.

The Change Security Key dialog box opens.

3. Enter information in the following fields.

- **Define a security key identifier** -- (For internal security keys only.) Accept the default value (storage array name and time stamp, which is generated by the controller firmware) or enter your own value. You can enter up to 189 alphanumeric characters without spaces, punctuation, or symbols.



Additional characters are generated automatically and are appended to both ends of the string you enter. The generated characters help to ensure that the identifier is unique.

- **Define a pass phrase/Re-enter pass phrase** — In each of these fields, enter your pass phrase. The value can have between 8 and 32 characters, and must include each of the following:
 - An uppercase letter (one or more). Keep in mind that the pass phrase is case sensitive.
 - A number (one or more).
 - A non-alphanumeric character, such as !, *, @ (one or more).



Be sure to record your entries for later use — If you need to move a secure-enabled drive from the storage array, you must know the identifier and pass phrase to unlock drive data.

4. Click **Change**.

The new security key overwrites the previous key, which is no longer valid.



The path for the downloaded file might depend on the default download location of your browser.

5. Record your key identifier, pass phrase, and the location of the downloaded key file, and then click **Close**.

After you finish

You should validate the security key to make sure the key file is not corrupted.

Switch from external to internal key management

You can change the management method for Drive Security from an external key server to the internal method used by the storage array. The security key previously defined for external key management is then used for internal key management.

Before you begin

An external key was created.

About this task

In this task, you disable external key management and download a new backup copy to your local host. The existing key is still used for Drive Security, but will be managed internally in the storage array.

Steps

1. Select **Settings > System**.
2. Under **Security key management**, select **Disable External Key Management**.

The **Disable External Key Management** dialog box opens.

3. In **Define a pass phrase/Re-enter pass phrase**, enter and confirm a pass phrase for the backup of the key. The value can have between 8 and 32 characters, and must include each of the following:
 - An uppercase letter (one or more). Keep in mind that the pass phrase is case sensitive.
 - A number (one or more).
 - A non-alphanumeric character, such as !, *, @ (one or more).



Be sure to record your entries for later use. If you need to move a secure-enabled drive from the storage array, you must know the identifier and pass phrase to unlock drive data.

4. Click **Disable**.

The backup key is downloaded to your local host.

5. Record your key identifier, pass phrase, and the location of the downloaded key file, and then click **Close**.

Results

Drive Security is now managed internally through the storage array.

After you finish

- You should validate the security key to make sure the key file is not corrupted.

Edit key management server settings

If you configured external key management, you can view and edit the key management server settings at any time.

Before you begin

External key management must be configured.

Steps

1. Select **Settings > Systems**.
2. Under **Security key management**, select **View/Edit Key Management Server Settings**.
3. Edit information in the following fields:
 - **Key management server address** — Enter the fully qualified domain name or the IP address (IPv4 or IPv6) of the server used for key management.
 - **KMIP port number** — Enter the port number used for the Key Management Interoperability Protocol (KMIP) communications.
4. Click **Save**.

Back up security key

After creating or changing a security key, you can create a backup copy of the key file in case the original gets corrupted.

Before you begin

- A security key already exists.

About this task

This task describes how to back up a security key you previously created. During this procedure, you create a new pass phrase for the backup. This pass phrase does not need to match the pass phrase that was used when the original key was created or last changed. The pass phrase is applied only to the backup you are creating.

Steps

1. Select **Settings > System**.
2. Under **Security key management**, select **Back Up Key**.

The Back Up Security Key dialog box opens.

3. In the **Define a pass phrase/Re-enter pass phrase** fields, enter and confirm a pass phrase for this backup.

The value can have between 8 and 32 characters, and must include each of the following:

- An uppercase letter (one or more)
- A number (one or more)
- A non-alphanumeric character, such as **!**, *****, **@** (one or more)



Be sure to record your entry for later use. You need the pass phrase to access the backup of this security key.

4. Click **Back Up**.

A backup of the security key is downloaded to your local host, and then the **Confirm/Record Security Key Backup** dialog box opens.



The path for the downloaded security key file might depend on the default download location of your browser.

5. Record your pass phrase in a secure location, and then click **Close**.

After you finish

You should validate the backup security key.

Validate security key

You can validate the security key to make sure it has not been corrupted and to verify that you have a correct pass phrase.

Before you begin

A security key has been created.

About this task

This task describes how to validate the security key you previously created. This is an important step to make sure the key file is not corrupted and the pass phrase is correct, which ensures that you can later access drive data if you move a secure-enabled drive from one storage array to another.

Steps

1. Select **Settings > System**.
2. Under **Security key management**, select **Validate Key**.

The **Validate Security Key** dialog box opens.

3. Click **Browse**, and then select the key file (for example, `drivesecurity.slk`).
4. Enter the pass phrase associated with the key you selected.

When you select a valid key file and pass phrase, the **Validate** button becomes available.

5. Click **Validate**.

The results of the validation are displayed in the dialog box.

6. If the results show "The security key validated successfully," click **Close**. If an error message appears, follow the suggested instructions displayed in the dialog box.

Unlock drives using a security key

If you move secure-enabled drives from one storage array to another, you must import the appropriate security key to the new storage array. Importing the key unlocks the data on the drives.

Before you begin

- The target storage array (where you are moving the drives) must already have a security key configured. The migrated drives will be re-keyed to the target storage array.
- You must know the security key that is associated with the drives you want to unlock.
- The security key file is available on the management client (the system with a browser used for accessing System Manager). If you are moving the drives to a storage array that is managed by a different system, you need to move the security key file to that management client.

About this task

This task describes how to unlock data in secure-enabled drives that have been removed from a storage array and reinstalled in another. Once the array discovers the drives, a "Needs Attention" condition appears along with a status of "Security Key Needed" for these re-located drives. You can unlock drive data by importing their security key into the storage array. During this process, you select the security key file and enter the pass phrase for the key.



The pass phrase is not the same as the storage array's Administrator password.

If other secure-enabled drives are installed in the new storage array, they might use a different security key than the one you are importing. During the import process, the old security key is used only to unlock the data for the drives you are installing. When the unlock process is successful, the newly installed drives are re-keyed to the target storage array's security key.

Steps

1. Select **Settings > System**.
2. Under **Security key management**, select **Unlock Secure Drives**.

The Unlock Secure Drives dialog box opens. Any drives that require a security key are shown in the table.

3. **Optional:** Hover the mouse over a drive number to see the location of the drive (shelf number and bay number).
4. Click **Browse**, and then select the security key file that corresponds to the drive you want to unlock.

The key file you selected appears in the dialog box.

5. Enter the pass phrase associated with this key file.

The characters you enter are masked.

6. Click **Unlock**.

If the unlock operation is successful, the dialog box displays: "The associated secure drives have been unlocked."

Results

When all drives are locked and then unlocked, each controller in the storage array will reboot. However, if there are already some unlocked drives in the target storage array, then the controllers will not reboot.

FAQs

What do I need to know before creating a security key?

A security key is shared by controllers and secure-enabled drives within a storage array.

If a secure-enabled drive is removed from the storage array, the security key protects the data from unauthorized access.

You can create and manage security keys using one of the following methods:

- Internal key management on the controller's persistent memory.
- External key management on an external key management server.

Before creating an internal security key, you must do the following:

1. Install secure-capable drives in the storage array. These drives can be Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.
2. Make sure the Drive Security feature is enabled. If necessary, contact your storage vendor for instructions on enabling the Drive Security feature.

You can then create an internal security key, which involves defining an identifier and a pass phrase. The identifier is a string that is associated with the security key, and is stored on the controller and on all drives associated with the key. The pass phrase is used to encrypt the security key for backup purposes. When you are finished, the security key is stored on the controller in a non-accessible location. You can then create secure-enabled volume groups or pools, or you can enable security on existing volume groups and pools.

Before creating an external security key, you must do the following:

1. Install secure-capable drives in the storage array. These drives can be Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.
2. Make sure the Drive Security feature is enabled. If necessary, contact your storage vendor for instructions on enabling the Drive Security feature.
3. Complete and download a client Certificate Signing Request (CSR) for authentication between the storage array and the key management server. Go to **Settings > Certificates > Key Management > Complete CSR**.
4. Create and download a client certificate from the key management server using the downloaded CSR file.
5. Ensure that the client certificate and a copy of the certificate for the key management server are available on your local host.

You can then create an external key, which involves defining the IP address of the key management server and the port number used for KMIP communications. During this process, you also load certificate files. When you are finished, the system connects to the key management server with the credentials you entered. You can then create secure-enabled volume groups or pools, or you can enable security on existing volume groups and pools.

Why do I need to define a pass phrase?

The pass phrase is used to encrypt and decrypt the security key file stored on the local management client. Without the pass phrase, the security key cannot be decrypted and used to unlock data from a secure-enabled drive if it is re-installed in another storage array.

Why is it important to record security key information?

If you lose the security key information and do not have a backup, you could lose data when relocating secure-enabled drives or upgrading a controller. You need the security

key to unlock data on the drives.

Be sure to record the security key identifier, the associated pass phrase, and the location on the local host where the security key file was saved.

What do I need to know before backing up a security key?

If your original security key becomes corrupted and you do not have a backup, you will lose access to the data on drives if they are migrated from one storage array to another.

Before backing up a security key, keep these guidelines in mind:

- Make sure you know the security key identifier and pass phrase of the original key file.



Only internal keys use identifiers. When you created the identifier, additional characters were generated automatically and appended to both ends of the identifier string. The generated characters ensure that the identifier is unique.

- You create a new pass phrase for the backup. This pass phrase does not need to match the pass phrase that was used when the original key was created or last changed. The pass phrase is only applied to the backup you are creating.



The pass phrase for Drive Security should not be confused with the storage array's Administrator password. The pass phrase for Drive Security protects backups of a security key. The Administrator password protects the entire storage array from unauthorized access.

- The backup security key file is downloaded to your management client. The path for the downloaded file might depend on the default download location of your browser. Be sure to make a record of where your security key information is stored.

What do I need to know before unlocking secure drives?

To unlock the data from a secure-enabled drive that is migrated to a new storage array, you must import its security key.

Before unlocking secure-enabled drives, keep the following guidelines in mind:

- The target storage array (where you are moving the drives) must already have a security key. The migrated drives will be re-keyed to the target storage array.
- For the drives you are migrating, you know the security key identifier and the pass phrase that corresponds to the security key file.
- The security key file is available on the management client (the system with a browser used for accessing System Manager).
- If you are resetting a locked NVMe drive, you must enter the drive's security ID. To locate the security ID, you must physically remove the drive and find the PSID string (maximum of 32 characters) on the drive's label. Make sure the drive is reinstalled before you start the operation.

What is read/write accessibility?

The Drive Settings window includes information about the Drive Security attributes.

"Read/Write Accessible" is one of the attributes that displays if a drive's data has been locked.

To view Drive Security attributes, go to the Hardware page. Select a drive, click **View settings**, and then click **Show more settings**. At the bottom of the page, the Read/Write Accessible attribute value is **Yes** when the drive is unlocked. The Read/Write Accessible attribute value is **No, invalid security key** when the drive is locked. You can unlock a secure drive by importing a security key (go to **Settings > System > Unlock Secure Drives**).

What do I need to know about validating the security key?

After creating a security key, you should validate the key file to make sure it is not corrupt.

If the validation fails, do the following:

- If the security key identifier does not match the identifier on the controller, locate the correct security key file and then try the validation again.
- If the controller cannot decrypt the security key for validation, you might have incorrectly entered the pass phrase. Double-check the pass phrase, re-enter it if necessary, and then try the validation again. If the error message appears again, select a backup of the key file (if available) and re-try validation.
- If you still cannot validate the security key, the original file might be corrupted. Create a new backup of the key and validate that copy.

What is the difference between internal security key and external security key management?

When you implement the Drive Security feature, you can use an internal security key or an external security key to lock down data when a secure-enabled drive is removed from the storage array.

A security key is a string of characters, which is shared between the secure-enabled drives and controllers in a storage array. Internal keys are maintained on the controller's persistent memory. External keys are maintained on a separate key management server, using a Key Management Interoperability Protocol (KMIP).

Access Management

Concepts

How Access Management works

Access Management is a method of establishing user authentication in SANtricity System Manager.

Access Management configuration and user authentication works as follows:

1. An administrator logs in to System Manager with a user profile that includes Security Admin permissions.



For first-time login, the username `admin` is automatically displayed and cannot be changed. The `admin` user has full access to all functions in the system.

2. The administrator navigates to Access Management in the user interface. The storage array is pre-configured to use local user roles, which are an implementation of RBAC (role-based access control)

capabilities.

3. The administrator configures one or more of the following authentication methods:

- **Local user roles** — Authentication is managed through RBAC capabilities enforced in the storage array. Local user roles include pre-defined user profiles and roles with specific access permissions. Administrators can use these local user roles as the single method of authentication, or use them in combination with a directory service. No configuration is necessary, other than setting passwords for users.
- **Directory services** — Authentication is managed through an LDAP (Lightweight Directory Access Protocol) server and directory service, such as Microsoft's Active Directory. An administrator connects to the LDAP server, and then maps the LDAP users to the local user roles embedded in the storage array.
- **SAML** — Authentication is managed through an Identity Provider (IdP) using the Security Assertion Markup Language (SAML) 2.0. An administrator establishes communication between the IdP system and the storage array, and then maps IdP users to the local user roles embedded in the storage array.

4. The administrator provides users with login credentials for System Manager.

5. Users log in to the system by entering their credentials.



If authentication is managed with SAML and an SSO (single sign-on), the system might bypass the System Manager login dialog.

During login, the system performs the following background tasks:

- Authenticates the user name and password against the user account.
- Determines the user's permissions based on the assigned roles.
- Provides the user with access to tasks in the user interface.
- Displays the user name in the upper right of the interface.

Tasks available in System Manager

Access to tasks depends on a user's assigned roles, which include the following:

- **Storage admin** — Full read/write access to the storage objects (for example, volumes and disk pools), but no access to the security configuration.
- **Security admin** — Access to the security configuration in Access Management, certificate management, audit log management, and the ability to turn the legacy management interface (SYMBOL) on or off.
- **Support admin** — Access to all hardware resources on the storage array, failure data, MEL events, and controller firmware upgrades. No access to storage objects or the security configuration.
- **Monitor** — Read-only access to all storage objects, but no access to the security configuration.

An unavailable task is either grayed out or does not display in the user interface. For example, a user with the Monitor role can view all information about volumes, but cannot access functions for modifying that volume. The tabs for features such as **Copy Services** and **Add to Workload** will be grayed out; only **View/Edit Settings** is available.

Limitations in SANtricity Unified Manager and SANtricity Storage Manager

If SAML is configured for a storage array, users cannot discover or manage storage for that array from the SANtricity Unified Manager or the SANtricity Storage Manager interfaces.

When local user roles and directory services are configured, users must enter credentials before performing any of the following functions:

- Renaming the storage array
- Upgrading controller firmware
- Loading a storage array configuration
- Executing a script
- Attempting to perform an active operation when an unused session has timed out

Access Management terminology

Learn how the Access Management terms apply to your storage array.

Term	Description
Active Directory	Active Directory (AD) is a Microsoft directory service that uses LDAP for Windows domain networks.
Binding	Bind operations are used to authenticate clients to the directory server. Binding usually requires account and password credentials, but some servers allow for anonymous bind operations.
CA	A certificate authority (CA) is a trusted entity that issues electronic documents, called digital certificates, for Internet security. These certificates identify website owners, which allows for secure connections between clients and servers.
Certificate	A certificate identifies the owner of a site for security purposes, which prevents attackers from impersonating the site. The certificate contains information about the site owner and the identity of the trusted entity who certifies (signs) this information.
IdP	An Identity Provider (IdP) is an external system used to request credentials from a user and to determine if that user is successfully authenticated. The IdP can be configured to provide multi-factor authentication and to use any user database, such as Active Directory. Your security team is responsible for maintaining the IdP.
LDAP	Lightweight Directory Access Protocol (LDAP) is an application protocol for accessing and maintaining distributed directory information services. This protocol allows many different applications and services to connect to the LDAP server for validating users.
RBAC	Role-based access control (RBAC) is a method of regulating access to computer or network resources based on the roles of individual users. RBAC controls are enforced on the storage array and include predefined roles.

Term	Description
SAML	Security Assertion Markup Language (SAML) is an XML-based standard for authentication and authorization between two entities. SAML allows for multi-factor authentication, in which users must provide two or more items for proving their identity (for example, a password and fingerprint). The storage array's embedded SAML feature is SAML2.0 compliant for identity assertion, authentication, and authorization.
SP	A Service Provider (SP) is a system that controls user authentication and access. When Access Management is configured with SAML, the storage array acts as the Service Provider for requesting authentication from the Identity Provider.
SSO	Single sign-on (SSO) is an authentication service that allows for one set of login credentials to access multiple applications.

Permissions for mapped roles

The RBAC (role-based access control) capabilities enforced on the storage array include pre-defined user profiles with one or more roles mapped to them. Each role includes permissions for accessing tasks in SANtricity System Manager.

User profiles and mapped roles are accessible from **Settings > Access Management > Local User Roles** in the user interface of either System Manager.

The roles provide user access to tasks, as follows:

- **Storage admin** — Full read/write access to the storage objects (for example, volumes and disk pools), but no access to the security configuration.
- **Security admin** — Access to the security configuration in Access Management, certificate management, audit log management, and the ability to turn the legacy management interface (SYMbol) on or off.
- **Support admin** — Access to all hardware resources on the storage array, failure data, MEL events, and controller firmware upgrades. No access to storage objects or the security configuration.
- **Monitor** — Read-only access to all storage objects, but no access to the security configuration.

If a user does not have permissions for a certain task, that task is either grayed out or does not display in the user interface.

Access Management with local user roles

For Access Management, administrators can use RBAC (role-based access control) capabilities enforced in the storage array. These capabilities are referred to as "local user roles."

Configuration workflow

Local user roles are pre-configured for the storage array. To use local user roles for authentication, administrators can do the following:

1. An administrator logs in to SANtricity System Manager with a user profile that includes Security Admin

permissions.



The `admin` user has full access to all functions in the system.

2. An administrator reviews the user profiles, which are predefined and cannot be modified.
3. **Optional:** The administrator assigns new passwords for each user profile.
4. Users log in to the system with their assigned credentials.

Management

When using only local user roles for authentication, administrators can perform the following management tasks:

- Change passwords.
- Set a minimum length for passwords.
- Allow users to log in without passwords.

Access Management with directory services

For Access Management, administrators can use an LDAP (Lightweight Directory Access Protocol) server and a directory service, such as Microsoft's Active Directory.

Configuration workflow

If an LDAP server and directory service are used in the network, configuration works as follows:

1. An administrator logs in to SANtricity System Manager with a user profile that includes Security Admin permissions.



The `admin` user has full access to all functions in the system.

2. The administrator enters the configuration settings for the LDAP server. Settings include the domain name, URL, and Bind account information.
3. If the LDAP server uses a secure protocol (LDAPS), the administrator uploads a Certificate Authority (CA) certificate chain for authentication between the LDAP server and the storage array.
4. After the server connection is established, the administrator maps the user groups to the storage array's roles. These roles are predefined and cannot be modified.
5. The administrator tests the connection between the LDAP server and the storage array.
6. Users log in to the system with their assigned LDAP/Directory Services credentials.

Management

When using directory services for authentication, administrators can perform the following management tasks:

- Add a directory server.
- Edit directory server settings.
- Map LDAP users to local user roles.
- Remove a directory server.

Access Management with SAML

For Access Management, administrators can use the Security Assertion Markup Language (SAML) 2.0 capabilities embedded in the array.

Configuration workflow

SAML configuration works as follows:

1. An administrator logs in to System Manager with a user profile that includes Security Admin permissions.



The `admin` user has full access to all functions in System Manager.

2. The administrator goes to the **SAML** tab under Access Management.
3. An administrator configures communications with the Identity Provider (IdP). An IdP is an external system used to request credentials from a user and determine if the user is successfully authenticated. To configure communications with the storage array, the administrator downloads the IdP metadata file from the IdP system, and then uses System Manager to upload the file to the storage array.
4. An administrator establishes a trust relationship between the Service Provider and the IdP. A Service Provider controls user authorization; in this case, the controller in the storage array acts as the Service Provider. To configure communications, the administrator uses System Manager to export a Service Provider metadata file for each controller. From the IdP system, the administrator then imports those metadata files to the IdP.



Administrators should also make sure that the IdP supports the ability to return a Name ID on authentication.

5. The administrator maps the storage array's roles to user attributes defined in the IdP. To do this, the administrator uses System Manager to create the mappings.
6. The administrator tests the SSO login to the IdP URL. This test ensures the storage array and IdP can communicate.



Once SAML is enabled, you *cannot* disable it through the user interface, nor can you edit the IdP settings. If you need to disable or edit the SAML configuration, contact Technical Support for assistance.

7. From System Manager, the administrator enables SAML for the storage array.
8. Users log in to the system with their SSO credentials.

Management

When using SAML for authentication, administrators can perform the following management tasks:

- Modify or create new role mappings
- Export Service Provider files

Access restrictions

When SAML is enabled, users cannot discover or manage storage for that array from the SANtricity Unified Manager or the SANtricity Storage Manager interfaces.

In addition, the following clients cannot access storage array services and resources:

- Enterprise Management Window (EMW)
- Command-line interface (CLI)
- Software Developer Kits (SDK) clients
- In-band clients
- HTTP Basic Authentication REST API clients
- Login using standard REST API endpoint

How to

View local user roles

From the Local User Roles tab, you can view the mappings of the user profiles to the default roles. These mappings are part of the RBAC (role-based access controls) enforced in the storage array.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.

About this task

The user profiles and mappings cannot be changed. Only passwords can be modified.

Steps

1. Select **Settings > Access Management**.
2. Select the **Local User Roles** tab.

The user profiles are shown in the table:

- **Root admin** (admin) — Super administrator who has access to all functions in the system. This user profile includes all roles.
- **Storage admin** (storage) — The administrator responsible for all storage provisioning. This user profile includes the following roles: Storage Admin, Support Admin, and Monitor.
- **Security admin** (security) — The user responsible for security configuration, including access management, certificate management, and secure-enabled drive functions. This user profile includes the following roles: Security Admin and Monitor.
- **Support admin** (support) — The user responsible for hardware resources, failure data, and firmware upgrades. This user profile includes the following roles: Support Admin and Monitor.
- **Monitor** (monitor) — A user with read-only access to the system. This user profile includes only the Monitor role.

Change passwords

You can change the user passwords for each user profile in Access Management.

Before you begin

- You must be logged in as the local administrator, which includes Root admin permissions.

- You must know the local administrator password.

About this task

Keep these guidelines in mind when choosing a password:

- Any new local user passwords must meet or exceed the current setting for a minimum password (in View/Edit Settings).
- Passwords are case sensitive.
- Trailing spaces are not stripped from passwords when they are set. Be careful to include spaces if they were included in the password.
- For increased security, use at least 15 alphanumeric characters and change the password frequently.



Changing the password in System Manager also changes it in the command line interface (CLI). In addition, password changes cause the user's active session to terminate.

Steps

1. Select **Settings** > **Access Management**.
2. Select the **Local User Roles** tab.
3. Select a user from the table.

The Change Password button becomes available.

4. Select **Change Password**.

The Change Password dialog box opens.

5. If no minimum password length is set for local user passwords, you can check the box to require the selected user to enter a password to access the storage array, and then you can type the new password for the selected user.
6. Enter your local administrator password, and then click **Change**.

Results

If the user is currently logged in, the password change causes the user's active session to terminate.

Change local user password settings

You can set the minimum required length for all new or updated local user passwords on the storage array. You can also allow local users to access the storage array without entering a password.

Before you begin

- You must be logged in as the local administrator, which includes Root admin permissions.

About this task

Keep these guidelines in mind when setting the minimum length for local user passwords:

- Setting changes will not affect existing local user passwords.
- The minimum required length setting for local user passwords must be between 0 and 30 characters.
- Any new local user passwords must meet or exceed the current minimum length setting.

- Do not set a minimum length for the password if you want local users to access the storage array without entering a password.

Steps

1. Select **Settings > Access Management**.
2. Select the **Local User Roles** tab.
3. Select the **View/Edit Settings** button.

The **Local User Password Settings** dialog box opens.

4. Do one of the following:
 - To allow local users to access the storage array *without* entering a password, uncheck the "Require all local user passwords to be at least" checkbox.
 - To set a minimum password length for all local user passwords, check the "Require all local user passwords to be at least" checkbox and then use the spinner box to set the minimum required length for all local user passwords.

Any new local user passwords must meet or exceed the current setting.

5. Click **Save**.

Add directory server

To configure authentication for Access Management, you can establish communications between the storage array and an LDAP server, and then map the LDAP user groups to the array's predefined roles.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.
- User groups must be defined in your directory service.
- LDAP server credentials must be available, including the domain name, server URL, and optionally the bind account user name and password.
- For LDAPS servers using a secure protocol, the LDAP server's certificate chain must be installed on your local machine.

About this task

Adding a directory server is a two-step process. First you enter the domain name and URL. If your server uses a secure protocol, you must also upload a CA certificate for authentication if it is signed by a non-standard signing authority. If you have credentials for a bind account, you can also enter your user account name and password. Next, you map the LDAP server's user groups to the storage array's predefined roles.



During the procedure to add an LDAP server, the legacy management interface will be disabled. The legacy management interface (SYMBOL) is a method of communication between the storage array and the management client. When disabled, the storage array and management client use a more secure method of communication (REST API over https).

Steps



1. Select **Settings > Access Management**.

2. From the **Directory Services** tab, select **Add Directory Server**.

The Add Directory Server dialog box opens.

3. In the **Server Settings** tab, enter the credentials for the LDAP server.

Field Details

Setting	Description
Configuration settings	
Domain(s)	Enter the domain name of the LDAP server. For multiple domains, enter the domains in a comma separated list. The domain name is used in the login (<i>username@domain</i>) to specify which directory server to authenticate against.
Server URL	Enter the URL for accessing the LDAP server in the form of <code>ldap[s]://host:*port*</code> .
Upload certificate (optional)	<div>  <p>This field appears only if an LDAPS protocol is specified in the Server URL field above.</p> </div> <p>Click Browse and select a CA certificate to upload. This is the trusted certificate or certificate chain used for authenticating the LDAP server.</p>
Bind account (optional)	Enter a read-only user account for search queries against the LDAP server and for searching within the groups. Enter the account name in an LDAP-type format. For example, if the bind user is called "bindacct," then you might enter a value such as "CN=bindacct,CN=Users,DC=cpoc,DC=local."
Bind password (optional)	<div>  <p>This field appears when you enter a bind account above.</p> </div> <p>Enter the password for the bind account.</p>
Test server connection before adding	Select this checkbox if you want to make sure the storage array can communicate with the LDAP server configuration you entered. The test occurs after you click Add at the bottom of the dialog box. If this checkbox is selected and the test fails, the configuration is not added. You must resolve the error or de-select the checkbox to skip the testing and add the configuration.
Privilege settings	
Search base DN	Enter the LDAP context to search for users, typically in the form of <code>CN=Users, DC=cpoc, DC=local</code> .
Username attribute	Enter the attribute that is bound to the user ID for authentication. For example: <code>sAMAccountName</code> .
Group attribute(s)	Enter a list of group attributes on the user, which is used for group-to-role mapping. For example: <code>memberOf, managedObjects</code> .

4. Click the **Role Mapping** tab.
5. Assign LDAP groups to the predefined roles. A group can have multiple assigned roles.

Field Details

Setting	Description
Mappings	
Group DN	Specify the group distinguished name (DN) for the LDAP user group to be mapped.
Roles	<p>Click in the field and select one of the storage array's roles to be mapped to the Group DN. You must individually select each role you want to include for this group. The Monitor role is required in combination with the other roles to log in to SANtricity System Manager. The mapped roles include the following permissions:</p> <ul style="list-style-type: none">• Storage admin — Full read/write access to the storage objects (for example, volumes and disk pools), but no access to the security configuration.• Security admin — Access to the security configuration in Access Management, certificate management, audit log management, and the ability to turn the legacy management interface (SYMBOL) on or off.• Support admin — Access to all hardware resources on the storage array, failure data, MEL events, and controller firmware upgrades. No access to storage objects or the security configuration.• Monitor — Read-only access to all storage objects, but no access to the security configuration.



The Monitor role is required for all users, including the administrator. System Manager will not operate correctly for any user without the Monitor role present.

6. If desired, click **Add another mapping** to enter more group-to-role mappings.
7. When you are finished with the mappings, click **Add**.

The system performs a validation, making sure that the storage array and LDAP server can communicate. If an error message appears, check the credentials entered in the dialog box and re-enter the information if necessary.

Edit directory server settings and role mappings

If you previously configured a directory server in Access Management, you can change its settings at any time. Settings include the server connection information and the group-to-role mappings.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.
- A directory server must be defined.

Steps

1. Select **Settings** > **Access Management**.
2. Select the **Directory Services** tab.
3. If more than one server is defined, select the server you want to edit from the table.
4. Select **View/Edit Settings**.

The **Directory Server Settings** dialog box opens.

5. In the **Server Settings** tab, change the desired settings.

Setting	Description
Configuration settings	
Domain(s)	The domain name(s) of the LDAP server(s). For multiple domains, enter the domains in a comma separated list. The domain name is used in the login (<i>username@domain</i>) to specify which directory server to authenticate against.
Server URL	The URL for accessing the LDAP server in the form of <code>ldap[s]://host:*port*</code> .
Bind account (optional)	The read-only user account for search queries against the LDAP server and for searching within the groups.
Bind password (optional)	The password for the bind account. (This field appears when a bind account is entered.)
Test server connection before saving	Checks that the storage array can communicate with the LDAP server configuration. The test occurs after you click Save at the bottom of the dialog box. If this checkbox is selected and the test fails, the configuration is not changed. You must resolve the error or de-select the checkbox to skip the testing and re-edit the configuration.
Privilege settings	
Search base DN	The LDAP context to search for users, typically in the form of <code>CN=Users, DC=copc, DC=local</code> .
Username attribute	The attribute that is bound to the user ID for authentication. For example: <code>sAMAccountName</code> .
Group attribute(s)	A list of group attributes on the user, which is used for group-to-role mapping. For example: <code>memberOf, managedObjects</code> .

6. In the **Role Mapping** tab, change the desired mapping.

Setting	Description
Mappings	
Group DN	The domain name for the LDAP user group to be mapped.
Roles	<p>The storage array's roles to be mapped to the Group DN. You must individually select each role you want to include for this group. The Monitor role is required in combination with the other roles to log in to SANtricity System Manager. The storage array's roles include the following:</p> <ul style="list-style-type: none">• Storage admin — Full read/write access to the storage objects (for example, volumes and disk pools), but no access to the security configuration.• Security admin — Access to the security configuration in Access Management, certificate management, audit log management, and the ability to turn the legacy management interface (SYMBOL) on or off.• Support admin — Access to all hardware resources on the storage array, failure data, MEL events, and controller firmware upgrades. No access to storage objects or the security configuration.• Monitor — Read-only access to all storage objects, but no access to the security configuration.



The Monitor role is required for all users, including the administrator. System Manager will not operate correctly for any user without the Monitor role present.

7. If desired, click **Add another mapping** to enter more group-to-role mappings.

8. Click **Save**.

Results

After you complete this task, any active user sessions are terminated. Only your current user session is retained.

Remove directory server

To break the connection between a directory server and the storage array, you can remove the server information from the Access Management page. You might want to perform this task if you configured a new server, and then want to remove the old one.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.

About this task

After you complete this task, any active user sessions are terminated. Only your current user session is retained.

Steps

1. Select **Settings** > **Access Management**.
2. Select the **Directory Services** tab.
3. From the list, select the directory server you want to delete.
4. Click **Remove**.

The **Remove Directory Server** dialog box opens.

5. Type `remove` in the field, and then click **Remove**.

The directory server configuration settings, privilege settings, and role mappings are removed. Users can no longer log in with credentials from this server.

Configure SAML

To configure authentication for Access Management, you can use the Security Assertion Markup Language (SAML) capabilities embedded in the storage array. This configuration establishes a connection between an Identity Provider and the Storage Provider.

About this task

An Identity Provider (IdP) is an external system used to request credentials from a user and to determine if that user is successfully authenticated. The IdP can be configured to provide multi-factor authentication and to use any user database, such as Active Directory. Your security team is responsible for maintaining the IdP. A Service Provider (SP) is a system that controls user authentication and access. When Access Management is configured with SAML, the storage array acts as the Service Provider for requesting authentication from the Identity Provider. To establish a connection between the IdP and storage array, you share metadata files between these two entities. Next, you map the IdP user entities to the storage array roles. And finally, you test the connection and SSO logins before enabling SAML.



SAML and Directory Services. If you enable SAML when Directory Services is configured as the authentication method, SAML supersedes Directory Services in System Manager. If you disable SAML later, the Directory Services configuration returns to its previous configuration.



Editing and Disabling. Once SAML is enabled, you *cannot* disable it through the user interface, nor can you edit the IdP settings. If you need to disable or edit the SAML configuration, contact Technical Support for assistance.

Configuring SAML authentication is a multi-step procedure.

Step 1: Upload the IdP metadata file

To provide the storage array with IdP connection information, you import IdP metadata into System Manager.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.
- An IdP administrator has configured an IdP system.
- An IdP administrator has ensured that the IdP supports the ability to return a Name ID on authentication.
- An administrator has ensured that the IdP server and controller clocks are synchronized (either through an

NTP server or by adjusting the controller clock settings).

- An IdP metadata file is downloaded from the IdP system and is available on the local system used for accessing System Manager.

About this task

In this task, you upload a metadata file from the IdP into System Manager. The IdP system needs this metadata to redirect authentication requests to the correct URL and to validate responses received. You only need to upload one metadata file for the storage array, even if there are two controllers.

Steps

1. Select **Settings > Access Management**.
2. Select the **SAML** tab.

The page displays an overview of configuration steps.

3. Click the **Import Identity Provider (IdP) file** link.

The **Import Identity Provider File** dialog box opens.

4. Click **Browse** to select and upload the IdP metadata file you copied to your local system.

After you select the file, the IdP Entity ID is displayed.

5. Click **Import**.

Step 2: Export Service Provider files

To establish a trust relationship between the IdP and the storage array, you import the Service Provider metadata into the IdP.

Before you begin

- You know the IP address or domain name of each controller in the storage array.

About this task

In this task, you export metadata from the controllers (one file for each controller). The IdP needs this metadata to establish a trust relationship with the controllers and to process authorization requests. The file includes information such as the controller domain name or IP address, so that the IdP can communicate with the Service Providers.

Steps

1. Click the **Export Service Provider files** link.

The **Export Service Provider Files** dialog box opens.

2. Enter the controller IP address or DNS name in the **Controller A** field, and then click **Export** to save the metadata file to your local system. If the storage array includes two controllers, repeat this step for the second controller in the **Controller B** field.

After you click **Export**, the Service Provider metadata is downloaded to your local system. Make a note of where the file is stored.

3. From the local system, locate the Service Provider metadata file(s) you exported.

There is one XML-formatted file for each controller.

4. From the IdP server, import the Service Provider metadata file(s) to establish the trust relationship. You can either import the files directly or you can manually enter the controller information from the files.

Step 3: Map roles

To provide users with authorization and access to System Manager, you must map the IdP user attributes and group memberships to the storage array's predefined roles.

Before you begin

- An IdP administrator has configured user attributes and group membership in the IdP system.
- The IdP metadata file is imported into System Manager.
- A Service Provider metadata file for each controller is imported into the IdP system for the trust relationship.

About this task

In this task, you use System Manager to map IdP groups to local user roles.

Steps

1. Click the link for mapping System Manager roles.

The Role Mapping dialog box opens.

2. Assign IdP user attributes and groups to the predefined roles. A group can have multiple assigned roles.

Field Details

Setting	Description
Mappings	
User Attribute	Specify the attribute (for example, "member of") for the SAML group to be mapped.
Attribute Value	Specify the attribute value for the group to be mapped.
Roles	<p>Click in the field and select one of the storage array's roles to be mapped to the Attribute. You must individually select each role you want to include. The Monitor role is required in combination with the other roles to log in to System Manager. The Security Admin role is also required for at least one group. The mapped roles include the following permissions:</p> <ul style="list-style-type: none">• Storage admin — Full read/write access to the storage objects (for example, volumes and disk pools), but no access to the security configuration.• Security admin — Access to the security configuration in Access Management, certificate management, audit log management, and the ability to turn the legacy management interface (SYMBOL) on or off.• Support admin — Access to all hardware resources on the storage array, failure data, MEL events, and controller firmware upgrades. No access to storage objects or the security configuration.• Monitor — Read-only access to all storage objects, but no access to the security configuration.



The Monitor role is required for all users, including the administrator. System Manager will not operate correctly for any user without the Monitor role present.

3. If desired, click **Add another mapping** to enter more group-to-role mappings.



Role mappings can be modified after SAML is enabled.

4. When you are finished with the mappings, click **Save**.

Step 4: Test SSO login

To ensure that the IdP system and storage array can communicate, you can optionally test an SSO login. This test is also performed during the final step for enabling SAML.

Before you begin

- The IdP metadata file is imported into System Manager.
- A Service Provider metadata file for each controller is imported into the IdP system for the trust relationship.

Steps

1. Select the **Test SSO Login** link.

A dialog box opens for entering SSO credentials.

2. Enter login credentials for a user with both Security Admin permissions and Monitor permissions.

A dialog box opens while the system tests the login.

3. Look for a Test Successful message. If the test completes successfully, go to the next step for enabling SAML.

If the test does not complete successfully, an error message appears with further information. Make sure that:

- The user belongs to a group with permissions for Security Admin and Monitor.
- The metadata you uploaded for the IdP server is correct.
- The controller addresses in the SP metadata files are correct.

Step 5: Enable SAML

Your final step is to enable SAML user authentication.

Before you begin

- The IdP metadata file is imported into System Manager.
- A Service Provider metadata file for each controller is imported into the IdP system for the trust relationship.
- At least one Monitor and one Security Admin role mapping is configured.

About this task

This task describes how to finish the SAML configuration for user authentication. During this process, the system also prompts you to test an SSO login. The SSO Login test process is described in the previous step.



Editing and Disabling. Once SAML is enabled, you *cannot* disable it through the user interface, nor can you edit the IdP settings. If you need to disable or edit the SAML configuration, contact Technical Support for assistance.

Steps

1. From the **SAML** tab, select the **Enable SAML** link.

The **Confirm Enable SAML** dialog box opens.

2. Type `enable`, and then click **Enable**.
3. Enter user credentials for an SSO login test.

Results

After the system enables SAML, it terminates all active sessions and begins authenticating users through SAML.

Change SAML role mappings

If you previously configured SAML for Access Management, you can change the role mappings between the IdP groups and the storage array's predefined roles.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.
- An IdP administrator has configured user attributes and group membership in the IdP system.
- SAML is configured and enabled.

Steps

1. Select **Settings** > **Access Management**.
2. Select the **SAML** tab.
3. Select **Role Mapping**.

The **Role Mapping** dialog box opens.

4. Assign IdP user attributes and groups to the predefined roles. A group can have multiple assigned roles.



Be careful that you do not remove your permissions while SAML is enabled, or you will lose access to System Manager.

Field Details

Setting	Description
Mappings	
User Attribute	Specify the attribute (for example, "member of") for the SAML group to be mapped.
Attribute Value	Specify the attribute value for the group to be mapped.
Roles	<p>Click in the field and select one of the storage array's roles to be mapped to the attribute. You must individually select each role you want to include for this group. The Monitor role is required in combination with the other roles to log in to System Manager. A Security Admin role must be assigned to at least one group. The mapped roles include the following permissions:</p> <ul style="list-style-type: none">• Storage admin — Full read/write access to the storage objects (for example, volumes and disk pools), but no access to the security configuration.• Security admin — Access to the security configuration in Access Management, certificate management, audit log management, and the ability to turn the legacy management interface (SYMBOL) on or off.• Support admin — Access to all hardware resources on the storage array, failure data, MEL events, and controller firmware upgrades. No access to storage objects or the security configuration.• Monitor — Read-only access to all storage objects, but no access to the security configuration.



The Monitor role is required for all users, including the administrator. System Manager will not operate correctly for any user without the Monitor role present.

5. **Optional:** Click **Add another mapping** to enter more group-to-role mappings.

6. Click **Save**.

Results

After you complete this task, any active user sessions are terminated. Only your current user session is retained.

Export SAML Service Provider files

If necessary, you can export Service Provider metadata for the storage array and re-import the file(s) into the Identity Provider (IdP) system.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.

- SAML is configured and enabled.

About this task

In this task, you export metadata from the controllers (one file for each controller). The IdP needs this metadata to establish a trust relationship with the controllers and to process authentication requests. The file includes information such as the controller domain name or IP address that the IdP can use for sending requests.

Steps

1. Select **Settings** > **Access Management**.
2. Select the **SAML** tab.
3. Select **Export**.

The **Export Service Provider Files** dialog box opens.

4. For each controller, click **Export** to save the metadata file to your local system.



The domain name fields for each controller are read-only.

Make a note of where the file is stored.

5. From the local system, locate the Service Provider metadata file(s) you exported.

There is one XML-formatted file for each controller.

6. From the IdP server, import the Service Provider metadata file(s). You can either import the files directly or you can manually enter the controller information from them.
7. Click **Close**.

View audit log activity

By viewing audit logs, users with Security Admin permissions can monitor user actions, authentication failures, invalid login attempts, and the user session lifespan.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.

Steps

1. Select **Settings** > **Access Management**.
2. Select the **Audit Log** tab.




Audit log activity appears in tabular format, which includes the following columns of information:

- **Date/Time** — Timestamp of when the storage array detected the event (in GMT).
- **Username** — The user name associated with the event. For any non-authenticated actions on the storage array, "N/A" appears as the user name. Non-authenticated actions might be triggered by the internal proxy or some other mechanism.
- **Status Code** — HTTP status code of the operation (200, 400, etc.) and descriptive text associated with the event.

- **URL Accessed** — Full URL (including host) and query string.
- **Client IP Address** — IP address of the client associated with the event.
- **Source** — Logging source associated with the event, which can be System Manager, CLI, Web Services, or Support Shell.

3. Use the selections on the Audit Log page to view and manage events.

Selection Details

Selection	Description
Show events from the...	Limit events shown by date range (last 24 hours, last 7 days, last 30 days, or a custom date range).
Filter	Limit events shown by the characters entered in the field. Use quotes (") for an exact word match, enter OR to return one or more words, or enter a dash (--) to omit words.
Refresh	Select Refresh to update the page to the most current events.
View/Edit Settings	Select View/Edit Settings to open a dialog box that allows you to specify a full log policy and level of actions to be logged.
Delete events	Select Delete to open a dialog box that allows you to remove old events from the page.
Show/hide columns	<p>Click the Show/Hide column icon  to select additional columns for display in the table. Additional columns include:</p> <ul style="list-style-type: none"> • Method — The HTTP method (for example, POST, GET, DELETE, etc.). • CLI Command Executed — The CLI command (grammar) executed for Secure CLI requests. • CLI Return Status — A CLI status code or a request for input files from the client. • SYMBOL Procedure — The SYMBOL procedure executed. • SSH Event Type — Secure Shell (SSH) events type, such as login, logout, and login_fail. • SSH Session PID — Process ID number of the SSH session. • SSH Session Duration(s) — The number of seconds the user was logged in.
Toggle column filters	Click the Toggle icon  to open filtering fields for each column. Enter characters within a column field to limit events shown by those characters. Click the icon again to close the filtering fields.
Undo changes	Click the Undo icon  to return the table to the default configuration.
Export	Click Export to save the table data to a comma separated value (CSV) file.

Define audit log policies

You can change the overwrite policy and the types of events recorded in the audit log.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.

About this task

This task describes how to change the audit log settings, which include the policy for overwriting old events and the policy for recording event types.



Steps

1. Select **Settings > Access Management**.
2. Select the **Audit Log** tab.
3. Select **View/Edit Settings**.

The **Audit Log Settings** dialog box opens.

4. Change the overwrite policy or types of events recorded.

Field Details

Setting	Description
Overwrite policy	<div>Determines the policy for overwriting old events when the maximum capacity is reached:</div> <div><ul style="list-style-type: none">• Allow the oldest events in the audit log to be overwritten when the audit log is full — Overwrites the old events when the audit log reaches 50,000 records.• Require audit log events to be manually deleted — Specifies that events will not be automatically deleted; instead, a threshold warning appears at the set percentage. Events must be deleted manually.</div> <div><div></div><div>If the overwrite policy is disabled and the audit log entries reach the maximum limit, access to System Manager is denied to users without Security Admin permissions. To restore system access to users without Security Admin permissions, a user assigned to the Security Admin role must delete the old event records.</div></div> <div><div></div><div>Overwrite policies do not apply if a syslog server is configured for archiving audit logs.</div></div>
Level of actions to be logged	<div>Determines types of events to be logged:</div> <div><ul style="list-style-type: none">• Record modification events only — Shows only the events where a user action involves making a change in the system.• Record all modification and read-only events — Shows all events, including a user action that involves reading or downloading information.</div>

5. Click **Save**.

Delete events from the audit log

You can clear the audit log of old events, which makes searching through events more manageable. You have the option of saving old events to a CSV (comma-separated values) file upon deletion.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.

About this task

This task describes how to remove old events from the audit log.

Steps

1. Select **Settings** > **Access Management**.
2. Select the **Audit Log** tab.
3. Select **Delete**.

The Delete Audit Log dialog box opens.

4. Select or enter the number of oldest events that you want to delete.
5. If you want to export the deleted events to a CSV file (recommended), keep the checkbox selected. You will be prompted to enter a file name and location when you click **Delete** in the next step. Otherwise, if you do not want to save events to a CSV file, click the checkbox to deselect it.
6. Click **Delete**.

A confirmation dialog box opens.

7. Type `delete` in the field, and then click **Delete**.

The oldest events are removed from the Audit Log page.

Configure syslog server for audit logs

If you want to archive audit logs onto an external syslog server, you can configure communications between that server and the storage array. After the connection is established, audit logs are automatically saved to the syslog server.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.
- The syslog server address, protocol, and port number must be available. The server address can be a fully qualified domain name, an IPv4 address, or an IPv6 address.
- If your server uses a secure protocol (for example, TLS), a Certificate Authority (CA) certificate must be available on your local system. CA certificates identify website owners for secure connections between servers and clients.

Steps

1. Select **Settings** › **Access Management**.
2. From the **Audit Log** tab, select **Configure Syslog Servers**.

The **Configure Syslog Servers** dialog box opens.

3. Click **Add**.

The **Add Syslog Server** dialog box opens.

4. Enter information for the server, and then click **Add**.
 - **Server address** — Enter a fully qualified domain name, an IPv4 address, or an IPv6 address.
 - **Protocol** — Select a protocol from the drop-down list (for example, TLS, UDP, or TCP).
 - **Upload certificate (optional)** — If you selected the TLS protocol and have not yet uploaded a signed CA certificate, click **Browse** to upload a certificate file. Audit logs are not archived to a syslog server without a trusted certificate.



If the certificate becomes invalid later, the TLS handshake will fail. As a result, an error message is posted to the audit log and messages are no longer sent to the syslog server. To resolve this issue, you must fix the certificate on the syslog server and then go to **Settings** › **Audit Log** › **Configure Syslog Servers** › **Test All**.

- **Port** — Enter the port number for the syslog receiver.

After you click **Add**, the **Configure Syslog Servers** dialog box opens and displays your configured syslog server on the page.

5. To test the server connection with the storage array, select **Test All**.

Results

After configuration, all new audit logs are sent to the syslog server. Previous logs are not transferred.

Edit syslog server settings for audit log records

You can change the settings for the syslog server used for archiving audit logs, and also upload a new Certificate Authority (CA) certificate for the server.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.
- The syslog server address, protocol, and port number must be available. The server address can be a fully qualified domain name, an IPv4 address, or an IPv6 address.
- If you are uploading a new CA certificate, the certificate must be available on your local system.

Steps

1. Select **Settings** › **Access Management**.
2. From the **Audit Log** tab, select **Configure Syslog Servers**.

Configured syslog servers are displayed on the page.

3. To edit the server information, select the **Edit** (pencil) icon to the right of the server name, and then make desired changes in the following fields:
 - **Server Address** — Enter a fully qualified domain name, an IPv4 address, or an IPv6 address.
 - **Protocol** — Select a protocol from the drop-down list (for example, TLS, UDP, or TCP).
 - **Port** — Enter the port number for the syslog receiver.
4. If you changed the protocol to the secure TLS protocol (from either UDP or TCP), click **Import Trusted Certificate** to upload a CA certificate.
5. To test the new connection with the storage array, select **Test All**.

Results

After configuration, all new audit logs are sent to the syslog server. Previous logs are not transferred.

FAQs

Why can't I log in?

If you receive an error when attempting to log in to System Manager, review these possible causes.

Login errors to System Manager might occur for one of these reasons:

- You entered an incorrect username or password.
- You have insufficient privileges.
- The directory server (if configured) might be unavailable. If this is the case, try logging in with a local user role.
- You attempted to log in unsuccessfully multiple times, which triggered the lockout mode. Wait 10 minutes to re-login.
- A lockout condition was triggered and your audit log might be full. Go to Access Management and delete old events from the audit log.
- SAML authentication is enabled. Refresh your browser to log in.

Login errors to a remote storage array for mirroring tasks might occur for one of these reasons:

- You have entered an incorrect password.
- You attempted to log in unsuccessfully multiple times, which triggered the lockout mode. Wait 10 minutes to log in again.
- The maximum number of client connections used on the controller has been reached. Check for multiple users or clients.

What do I need to know before adding a directory server?

Before adding a directory server in Access Management, make sure you meet the following requirements.

- User groups must be defined in your directory service.
- LDAP server credentials must be available, including the domain name, server URL, and optionally the bind account user name and password.

- For LDAPS servers using a secure protocol, the LDAP server's certificate chain must be installed on your local machine.

What do I need to know about mapping to storage array roles?

Before mapping groups to roles, review the following guidelines.

The storage array's embedded RBAC (role-based access control) capabilities include the following roles:

- **Storage admin** — Full read/write access to the storage objects (for example, volumes and disk pools), but no access to the security configuration.
- **Security admin** — Access to the security configuration in Access Management, certificate management, audit log management, and the ability to turn the legacy management interface (SYMBOL) on or off.
- **Support admin** — Access to all hardware resources on the storage array, failure data, MEL events, and controller firmware upgrades. No access to storage objects or the security configuration.
- **Monitor** — Read-only access to all storage objects, but no access to the security configuration.

Directory Services

If you are using an LDAP (Lightweight Directory Access Protocol) server and Directory Services, make sure that:

- An administrator has defined user groups in the directory service.
- You know the group domain names for the LDAP user groups.
- The Monitor role is required for all users, including the administrator. System Manager will not operate correctly for any user without the Monitor role present.

SAML

If you are using the Security Assertion Markup Language (SAML) capabilities embedded in the storage array, make sure that:

- An Identity Provider (IdP) administrator has configured user attributes and group membership in the IdP system.
- You know the group membership names.
- The Monitor role is required for all users, including the administrator. System Manager will not operate correctly for any user without the Monitor role present.

Which external management tools may be affected by this change?

When you make certain changes in System Manager, such as switching the management interface or using SAML for an authentication method, some external tools and features might be restricted from use.

Management interface

Tools that communicate directly with the legacy management interface (SYMBOL), such as the SANtricity SMI-S Provider or OnCommand Insight (OCI), do not work unless the Legacy Management Interface setting is enabled. In addition, you cannot use legacy CLI commands or perform mirroring operations if this setting is disabled.

Contact technical support for more information.

SAML authentication

When SAML is enabled, the following clients cannot access storage array services and resources:

- Enterprise Management Window (EMW)
- Command-line interface (CLI)
- Software Developer Kits (SDK) clients
- In-band clients
- HTTP Basic Authentication REST API clients
- Login using standard REST API endpoint

Contact technical support for more information.

What do I need to know before configuring and enabling SAML?

Before configuring and enabling the Security Assertion Markup Language (SAML) capabilities for authentication, make sure you meet the following requirements and understand SAML restrictions.

Requirements

Before you begin, make sure that:

- An Identity Provider (IdP) is configured in your network. An IdP is an external system used to request credentials from a user and determine if the user is successfully authenticated. Your security team is responsible for maintaining the IdP.
- An IdP administrator has configured user attributes and groups in the IdP system.
- An IdP administrator has ensured that the IdP supports the ability to return a Name ID on authentication.
- An administrator has ensured that the IdP server and controller clocks are synchronized (either through an NTP server or by adjusting the controller clock settings).
- An IdP metadata file is downloaded from the IdP system and available on the local system used for accessing System Manager.
- You know the IP address or domain name of each controller in the storage array.

Restrictions

In addition to the requirements above, make sure you understand the following restrictions:

- Once SAML is enabled, you *cannot* disable it through the user interface, nor can you edit the IdP settings. If you need to disable or edit the SAML configuration, contact Technical Support for assistance. We recommend that you test the SSO logins before you enable SAML in the final configuration step. (The system also performs an SSO login test before enabling SAML.)
- If you disable SAML in the future, the system automatically restores the previous configuration (Local User Roles and/or Directory Services).
- If Directory Services are currently configured for user authentication, SAML overrides that configuration.
- When SAML is configured, the following clients cannot access storage array resources:

- Enterprise Management Window (EMW)
- Command-line interface (CLI)
- Software Developer Kits (SDK) clients
- In-band clients
- HTTP Basic Authentication REST API clients
- Login using standard REST API endpoint

What types of events are recorded in the audit log?

The audit log can record modification events, or both modification and read-only events.

Depending on the policy settings, the following types of events are shown:

- **Modification events** — User actions from within System Manager that involve changes to the system, such as provisioning storage.
- **Modification and read-only events** — User actions that involve changes to the system, as well as events that involve viewing or downloading information, such as viewing volume assignments.

What do I need to know before configuring a syslog server?

You can archive audit logs onto an external syslog server.

Before configuring a syslog server, keep the following guidelines in mind.

- Make sure you know the server address, protocol, and port number. The server address can be a fully qualified domain name, an IPv4 address, or an IPv6 address.
- If your server uses a secure protocol (for example, TLS), a Certificate Authority (CA) certificate must be available on your local system. CA certificates identify website owners for secure connections between servers and clients.
- After configuration, all new audit logs are sent to the syslog server. Previous logs are not transferred.
- The Overwrite Policy settings (available from **View/Edit Settings**) do not affect how logs are managed with a syslog server configuration.
- Audit logs follow the RFC 5424 messaging format.

The syslog server is no longer receiving audit logs. What do I do?

If you configured a syslog server with a TLS protocol, the server cannot receive messages if the certificate becomes invalid for any reason. An error message about the invalid certificate is posted to the audit log.

To resolve this issue, you must first fix the certificate for the syslog server. Once a valid certificate chain is in place, go to **Settings > Audit Log > Configure Syslog Servers > Test All**.

Certificates

Concepts

How certificates work

Certificates are digital files that identify online entities, such as websites and servers, for secure communications on the internet.

Certificates ensure that web communications are transmitted in encrypted form, privately and unaltered, only between the specified server and client. Using System Manager, you can manage certificates between the browser on a host management system (acting as the client) and the controllers in a storage system (acting as the servers).

A certificate can be signed by a trusted authority, or it can be self-signed. "Signing" simply means that someone validated the owner's identity and determined that their devices can be trusted. Storage arrays ship with an automatically generated self-signed certificate on each controller. You can continue to use the self-signed certificates, or you can obtain CA-signed certificates for a more secure connection between the controllers and the host systems.



Although CA-signed certificates provide better security protection (for example, preventing man-in-the-middle attacks), they also require fees that can be expensive if you have a large network. In contrast, self-signed certificates are less secure, but they are free. Therefore, self-signed certificates are most often used for internal testing environments, not in production environments.

Signed certificates

A signed certificate is validated by a certificate authority (CA), which is a trusted third-party organization. Signed certificates include details about the owner of the entity (typically, a server or website), date of certificate issue and expiration, valid domains for the entity, and a digital signature composed of letters and numbers.

When you open a browser and enter a web address, your system performs a certificate-checking process in the background to determine if you are connecting to a website that includes a valid, CA-signed certificate. Generally, a site that is secured with a signed certificate includes a padlock icon and an https designation in the address. If you attempt to connect to a website that does not contain a CA-signed certificate, your browser displays a warning that the site is not secure.

The CA takes steps to verify your identity during the application process. They might send an email to your registered business, verify your business address, and perform an HTTP or DNS verification. When the application process is complete, the CA sends you digital files to load on a host management system. Typically, these files include a chain of trust, as follows:

- **Root** — At the top of the hierarchy is the root certificate, which contains a private key used to sign other certificates. The root identifies a particular CA organization. If you use the same CA for all your network devices, you need only one root certificate.
- **Intermediate** — Branching off from the root are the intermediate certificates. The CA issues one or more intermediate certificates to act as middlemen between a protected root and server certificates.
- **Server** — At the bottom of the chain is the server certificate, which identifies your specific entity, such as a website or other device. Each controller in an storage array requires a separate server certificate.

Self-signed certificates

Each controller in the storage array includes a pre-installed, self-signed certificate. A self-signed certificate is similar to a CA-signed certificate, except that it is validated by the owner of the entity instead of a third party. Like a CA-signed certificate, a self-signed certificate contains its own private key, and also ensures that data is

encrypted and sent over an HTTPS connection between a server and client. However, a self-signed certificate does not use the same chain of trust as a CA-signed certificate.

Self-signed certificates are not “trusted” by browsers. Each time you attempt to connect to a website that contains only a self-signed certificate, the browser displays a warning message. You must click a link in the warning message that allows you to proceed to the website; by doing so, you are essentially accepting the self-signed certificate.

Certificates used for key management server

If you are using an external key management server with the Drive Security feature, you can also manage certificates for authentication between that server and the controllers.

Certificate terminology

The following terms apply to certificate management.

Term	Description
CA	A certificate authority (CA) is a trusted entity that issues electronic documents, called digital certificates, for Internet security. These certificates identify website owners, which allows for secure connections between clients and servers.
CSR	A Certificate Signing Request (CSR) is a message that is sent from an applicant to a certificate authority (CA). The CSR validates the information the CA requires to issue a certificate.
Certificate	A certificate identifies the owner of a site for security purposes, which prevents attackers from impersonating the site. The certificate contains information about the site owner and the identity of the trusted entity who certifies (signs) this information.
Certificate chain	A hierarchy of files that adds a layer of security to the certificates. Typically, the chain includes one root certificate at the top of the hierarchy, one or more intermediate certificates, and the server certificates that identify the entities.
Client certificate	For security key management, a client certificate validates the storage array's controllers, so the key management server can trust their IP addresses.
Intermediate certificate	One or more intermediate certificates branch off from the root in the certificate chain. The CA issues one or more intermediate certificates to act as middlemen between a protected root and server certificates.
Key management server certificate	For security key management, a key management server certificate validates the server, so the storage array can trust its IP address.
Keystore	A keystore is a repository on your host management system that contains private keys, along with their corresponding public keys and certificates. These keys and certificates identify your own entities, such as the controllers.

Term	Description
OCSP server	The Online Certificate Status Protocol (OCSP) server determines if the certificate authority (CA) has revoked any certificates before their scheduled expiration date, and then blocks the user from accessing a server if the certificate is revoked.
Root certificate	The root certificate is at the top of the hierarchy in the certificate chain, and contains a private key used to sign other certificates. The root identifies a particular CA organization. If you use the same CA for all your network devices, you need only one root certificate.
Signed certificate	A certificate that is validated by a certificate authority (CA). This data file contains a private key and ensures that data is sent in encrypted form between a server and a client over an HTTPS connection. In addition, a signed certificate includes details about the owner of the entity (typically, a server or website) and a digital signature composed of letters and numbers. A signed certificate uses a chain of trust, and therefore is most often used in production environments. Also referred to as a "CA-signed certificate" or a "management certificate."
Self-signed certificate	A self-signed certificate is validated by the owner of the entity. This data file contains a private key and ensures that data is sent in encrypted form between a server and a client over an HTTPS connection. It also includes a digital signature composed of letters and numbers. A self-signed certificate does not use the same chain of trust as a CA-signed certificate, and therefore is most often used in test environments. Also referred to as a "preinstalled" certificate.
Server certificate	The server certificate is at the bottom of the certificate chain. It identifies your specific entity, such as a website or other device. Each controller in a storage system requires a separate server certificate.

How to

Use CA-signed certificates for controllers

You can obtain CA-signed certificates for secure communications between the controllers and the browser used for accessing System Manager.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.

About this task

Using CA-signed certificates is a three-step procedure.

Step 1: Complete and submit a CSR for the controllers

You must first generate a certificate signing request (CSR) file for each controller in the storage array, and then submit the file(s) to a certificate authority (CA).

Before you begin

- You must know the IP address or DNS name of each controller.

About this task

The CSR provides information about your organization, the IP address or DNS name of the controller, and a key pair that identifies the web server in the controller. During this task, one CSR file is generated if there is only one controller in the storage array and two CSR files if there are two controllers.



Do not generate a new CSR after submission to the CA. When you generate a CSR, the system creates a private and public key pair. The public key is part of the CSR, while the private key is kept in the keystore. When you receive the signed certificates and import them into the keystore, the system ensures that both the private and public keys are the original pair. Therefore, you must not generate a new CSR after submitting one to the CA. If you do, the controllers generate new keys, and the certificates you receive from the CA will not work.

Steps

1. Select **Settings** > **Certificates**.
2. From the **Array Management** tab, select **Complete CSR**.



If you see a dialog box prompting you to accept a self-signed certificate for the second controller, click **Accept Self-Signed Certificate** to proceed.

3. Enter the following information, and then click **Next**:

- **Organization** — The full, legal name of your company or organization. Include suffixes, such as Inc. or Corp.
- **Organizational unit (optional)** — The division of your organization that is handling the certificate.
- **City/Locality** — The city where your storage array or business is located.
- **State/Region (optional)** — The state or region where your storage array or business is located.
- **Country ISO code** — Your country's two-digit ISO (International Organization for Standardization) code, such as US.



Some fields might be pre-populated with the appropriate information, such as the IP address of the controller. Do not change prepopulated values unless you are certain they are incorrect. For example, if you have not yet completed a CSR, the controller IP address is set to "localhost." In this case, you must change "localhost" to the DNS name or IP address of the controller.

4. Verify or enter the following information about controller A in your storage array:

- **Controller A common name** — The IP address or DNS name of controller A is displayed by default. Make sure this address is correct; it must match exactly what you enter to access System Manager in the browser.
- **Controller A alternate IP addresses** — If the common name is an IP address, you can optionally enter any additional IP addresses or aliases for controller A. For multiple entries, use a comma-delimited format.
- **Controller A alternate DNS names** — If the common name is a DNS name, enter any additional DNS names for controller A. For multiple entries, use a comma-delimited format. If there are no alternate DNS names, but you entered a DNS name in the first field, copy that name here. If the storage array has only one controller, the **Finish** button is available. If the storage array has two controllers, the **Next** button is available.



Do not click the **Skip this step** link when you are initially creating a CSR request. This link is provided in error-recovery situations. In rare cases, a CSR request might fail on one controller, but not on the other. This link allows you to skip the step for creating a CSR request on controller A if it is already defined, and continue to the next step for re-creating a CSR request on controller B.

5. If there is only one controller, click **Finish**. If there are two controllers, click **Next** to enter information for controller B (same as above), and then click **Finish**.

For a single controller, one CSR file is downloaded to your local system. For dual controllers, two CSR files are downloaded. The folder location of the download depends on your browser.

6. Locate the downloaded CSR file(s). The folder location depends on your browser.
7. Submit the CSR file(s) to a CA and request signed certificates in PEM format.
8. Wait for the CA to return the certificates, and then go to [Step 2: Import signed certificates for controllers](#).

Step 2: Import signed certificates for controllers

After you receive signed certificates, you import the files for the controllers.

Before you begin

- The CA returned signed certificate files.
- The files are available on your local system.
- If the CA provided a chained certificate (for example, a .p7b file), you must unpack the chained file into individual files: the root certificate, one or more intermediate certificates, and the server certificates that identify the controllers. You can use the Windows `certmgr` utility to unpack the files (right-click and select **All Tasks > Export**). When the exports are complete, a CER file is shown for each certificate file in the chain.

About this task

This task describes how to upload the certificate files.

Steps

1. Select **Settings > Certificates**.
2. From the **Array Management** tab, select **Import**.

A dialog box opens for importing the certificate file(s).

3. Click the **Browse** buttons to first select the root and intermediate files, and then select each server certificate for the controllers. The root and intermediate files are the same for both controllers. Only the server certificates are unique for each controller.

The file names are displayed in the dialog box.

4. Click **Import**.

The file(s) are uploaded and validated.

Results

The session is automatically terminated. You must log in again for the certificate(s) to take effect. When you log

in again, the new CA-signed certificate is used for your session.

Reset management certificates

You can revert the certificates on the controllers from using CA-signed certificates back to the factory-set, self-signed certificates.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.
- CA-signed certificates must be previously imported.

About this task

The Reset function deletes the current CA-signed certificate files from each controller. The controllers will then revert to using self-signed certificates.

Steps

1. Select **Settings > Certificates**.
2. From the **Array Management** tab, select **Reset**.

A Confirm **Reset Management Certificates** dialog box opens.

3. Type `reset` in the field, and then click **Reset**.

After your browser refreshes, the browser might block access to the destination site and report that the site is using HTTP Strict Transport Security. This condition arises when you switch back to self-signed certificates. To clear the condition that is blocking access to the destination, you must clear the browsing data from the browser.

Results

The controllers revert to using self-signed certificates. As a result, the system prompts users to manually accept the self-signed certificate for their sessions.

View imported certificate information

From the Certificates page, you can view the certificate type, issuing authority, and the valid date range of certificates for the storage array.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.

Steps

1. Select **Settings > Certificates**.
2. Select one of the tabs to view information about the certificates.

Tab	Description
Array Management	View information about the CA-signed certificates imported for each controller, including the root file, intermediate file(s), and the server file(s).

Tab	Description
Trusted	<p>View information about all other types of certificates imported for the controllers. Use the filter field under Show certificates that are... to view either user-installed or pre-installed certificates.</p> <ul style="list-style-type: none"> • User-installed. Certificates that a user uploaded to the storage array, which can include trusted certificates when the controller acts as a client (instead of a server), LDAPS certificates, and Identity Federation certificates. • Pre-installed. Self-signed certificates included with the storage array.
Key Management	View information about the CA-signed certificates imported for an external key management server.

Import certificates for controllers when acting as clients

If the controller rejects a connection because it cannot validate the chain of trust for a network server, you can import a certificate from the Trusted tab that allows the controller (acting as a client) to accept communications from that server.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.
- The certificate files are installed on your local system.

About this task

Importing certificates from the Trusted tab might be necessary if you want to allow another server to contact the controllers (for example, an LDAP server or a syslog server that uses TLS).

Steps

1. Select **Settings** > **Certificates**.
2. From the **Trusted** tab, select **Import**.

A dialog box opens for importing the trusted certificate files.

3. Click **Browse** to select the certificate files for the controllers.

The file names display in the dialog box.

4. Click **Import**.

Results

The files are uploaded and validated.

Enable certificate revocation checking

You can enable automatic checks for revoked certificates, so that an Online Certificate Status Protocol (OCSP) server blocks users from making non-secure connections.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.
- A DNS server is configured on both controllers, which enables use of a fully qualified domain name for the OCSP server. This task is available from the Hardware page.
- If you want to specify your own OCSP server, you must know the URL of that server.

About this task

Automatic revocation checking is helpful in cases where the CA improperly issued a certificate, or a private key is compromised.

During this task, you can configure an OCSP server or use the server specified in the certificate file. The OCSP server determines if the CA has revoked any certificates before their scheduled expiration date, and then blocks the user from accessing a site if the certificate is revoked.

Steps

1. Select **Settings** > **Certificates**.
2. Select the **Trusted** tab.



You can also enable revocation checking from the **Key Management** tab.

3. Click **Uncommon Tasks**, and then select **Enable Revocation Checking** from the drop-down menu.
4. Select **I want to enable revocation checking**, so that a checkmark appears in the checkbox and additional fields appear in the dialog box.
5. In the **OCSP responder address** field, you can optionally enter a URL for an OCSP responder server. If you do not enter an address, the system uses the OCSP server's URL from the certificate file.
6. Click **Test Address** to make certain the system can open a connection to the specified URL.
7. Click **Save**.

Results

If the storage array attempts to connect to a server with a revoked certificate, the connection is denied and an event is logged.

Delete trusted certificates

You can delete the user-installed certificates previously imported from the Trusted tab.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.
- If you are updating a trusted certificate with a new version, the updated certificate must be imported before you delete the old certificate.



You might lose access to a system if you delete a certificate used to authenticate the controllers and another server, such as an LDAP server, before you import a replacement certificate.

About this task

This task describes how to delete user-installed certificates. The pre-installed, self-signed certificates cannot

be deleted.

Steps

1. Select **Settings > Certificates**.
2. Select the **Trusted** tab.

The table shows the storage array's trusted certificates.

3. From the table, select the certificate you want to remove.
4. Click **Uncommon Tasks > Delete**

A Confirm Delete Trusted Certificate dialog box opens.

5. Type `delete` in the field, and then click **Delete**.

Use CA-signed certificates for authentication with a key management server

For secure communications between a key management server and the storage array controllers, you must configure the appropriate sets of certificates.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.

About this task

Authenticating between the controllers and a key management server is a two-step procedure.

Step 1: Complete and submit CSR for authentication with a key management server

You must first generate a certificate signing request (CSR) file, and then use the CSR to request a signed client certificate from a certificate authority (CA) that is trusted by the key management server. You can also create and download a client certificate from the key management server using the downloaded CSR file.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.

About this task

This task describes how to generate the CSR file, which you will then use to request a signed client certificate from a CA that is trusted by the key management server. A client certificate validates the storage array's controllers, so the key management server can trust their Key Management Interoperability Protocol (KMIP) requests. During this task, you must provide information about your organization.

Steps

1. Select **Settings > Certificates**.
2. From the **Key Management** tab, select **Complete CSR**.
3. Enter the following information:
 - **Common name** — A name that identifies this CSR, such as the storage array name, which will be displayed in the certificate files.
 - **Organization** — The full, legal name of your company or organization. Include suffixes, such as Inc. or

Corp.

- **Organizational unit (optional)** — The division of your organization that is handling the certificate.
- **City/Locality** — The city or locality where your organization is located.
- **State/Region (optional)** — The state or region where your organization is located.
- **Country ISO code** — The two-digit ISO (International Organization for Standardization) code, such as US, where your organization is located.

4. Click **Download**.

A CSR file is saved to your local system.

5. Request a signed client certificate from a CA that is trusted by the key management server.
6. When you have a client certificate, go to [Step 2: Import certificates for the key management server](#).

Step 2: Import certificates for the key management server

As the next step, you import certificates for authentication between the storage array and the key management server. There are two types of certificates: the client certificate validates the storage array's controllers, while the key management server certificate validates the server.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.
- You have a signed client certificate file (see [Step 1: Complete and submit CSR for authentication with a key management server](#)), and you have copied that file to the host where you are accessing System Manager. A client certificate validates the storage array's controllers, so the key management server can trust their Key Management Interoperability Protocol (KMIP) requests.
- You must retrieve the server certificate file from the key management server, and then copy that file to the host where you are accessing System Manager. A key management server certificate validates the key management server, so the storage array can trust its IP address.



For more information about the server certificate, consult the documentation for your key management server.

About this task

This task describes how to upload certificate files for authentication between the storage array controllers and the key management server. You must load both the client certificate file for the controllers and the server certificate file for the key management server.

Steps

1. Select **Settings > Certificates**.
2. From the **Key Management** tab, select **Import**.

A dialog box opens for importing the certificate files.

3. Next to **Select client certificate**, click the **Browse** button to select the client certificate file for the storage array's controllers.

The file name displays in the dialog box.

4. Next to **Select key management server's server certificate**, click the **Browse** button to select the server certificate file for your key management server.

The file name displays in the dialog box.

5. Click **Import**.

The files are uploaded and validated.

Export key management server certificates

You can save a certificate for a key management server to your local machine.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.
- Certificates must be previously imported.

Steps

1. Select **Settings > Certificates**.
2. Select the **Key Management** tab.
3. From the table, select the certificate you want to export, and then click **Export**.

A Save dialog box opens.

4. Enter a filename and click **Save**.

FAQs

Why does the Cannot Access Other Controller dialog box appear?

When you perform certain operations related to CA certificates (for example, importing a certificate), you might see a dialog box prompting you to accept a self-signed certificate for the second controller.

In storage arrays with two controllers (duplex configurations), this dialog box sometimes appears if SANtricity System Manager cannot communicate with the second controller or if your browser cannot accept the certificate during a certain point in an operation.

If this dialog box opens, click **Accept Self-Signed Certificate** to proceed. If another dialog box prompts you for a password, enter your Administrator password used for accessing System Manager.

If this dialog box appears again and you cannot complete a certificate task, try one of the following procedures:

- Use a different browser type to access this controller, accept the certificate, and continue.
- Access the second controller with System Manager, accept the self-signed certificate, and then return to the first controller and continue.

How do I know what certificates need to be uploaded to System Manager for external key management?

For external key management, you import two types of certificates for authentication

between the storage array and the key management server so the two entities can trust each other.

A client certificate validates the storage array's controllers, so the key management server can trust their Key Management Interoperability Protocol (KMIP) requests. To obtain a client certificate, you use System Manager to complete a CSR for the storage array. You can then upload the CSR to a key management server and generate a client certificate from there. Once you have a client certificate, copy that file to the host where you are accessing System Manager.

A key management server certificate validates the key management server, so the storage array can trust its IP address. Retrieve the server certificate file from the key management server, and then copy that file to the host where you are accessing System Manager.

What do I need to know about certificate revocation checking?

System Manager allows you to check for revoked certificates by using an Online Certificate Status Protocol (OCSP) server, instead of uploading Certificate Revocation Lists (CRLs).

Revoked certificates should no longer be trusted. A certificate might be revoked for several reasons; for example, if the Certificate Authority (CA) improperly issued the certificate, a private key was compromised, or the identified entity did not adhere to policy requirements.

After you establish a connection to an OCSP server in System Manager, the storage array performs revocation checking whenever it connects to an AutoSupport server, External Key Management Server (EKMS), Lightweight Directory Access Protocol over SSL (LDAPS) server, or a Syslog server. The storage array attempts to validate these servers' certificates to ensure that they have not been revoked. The server then returns a value of "good," "revoked," or "unknown" for that certificate. If the certificate is revoked or the array cannot contact the OCSP server, the connection is refused.



Specifying an OCSP responder address in System Manager or in the command line interface (CLI) overrides the OCSP address found in the certificate file.

What types of servers will revocation checking be enabled for?

The storage array performs revocation checking whenever it connects to an AutoSupport server, External Key Management Server (EKMS), Lightweight Directory Access Protocol over SSL (LDAPS) server, or a Syslog server.

Support

Support center

Concepts

AutoSupport feature overview

The AutoSupport feature monitors the health of a storage array and sends automatic dispatches to technical support.

Technical support uses the AutoSupport data reactively to speed the diagnosis and resolution of customer

issues and proactively to detect and avoid potential issues.

AutoSupport data includes information about a storage array's configuration, status, performance, and system events. The AutoSupport data does not contain any user data. Dispatches can be sent immediately or on a schedule (daily and weekly).

Key benefits

Some key benefits of the AutoSupport feature include:

- Expedited case resolution times
- Sophisticated monitoring for faster incident management
- Automated reporting according to a schedule, as well as automated reporting about critical events
- Automated hardware replacement requests for selected components such as drives
- Nonintrusive alerting to notify you of a problem and provide information for technical support to take corrective action
- AutoSupport analysis tools that monitor dispatches for known configuration issues

Individual AutoSupport features

The AutoSupport feature is made up of three individual features that you enable separately.

- **Basic AutoSupport** — Allows your storage array to automatically collect and send data to technical support.
- **AutoSupport OnDemand** — Allows technical support to request retransmission of a previous AutoSupport dispatch when needed for troubleshooting an issue. All transmissions are initiated from the storage array, not from the AutoSupport server. The storage array checks in periodically with the AutoSupport server to determine if there are any pending retransmission requests and responds accordingly.
- **Remote Diagnostics** — Allows technical support to request a new, up-to-date AutoSupport dispatch when needed for troubleshooting an issue. All transmissions are initiated from the storage array, not from the AutoSupport server. The storage array checks in periodically with the AutoSupport server to determine if there are any pending new requests and responds accordingly.

Difference between AutoSupport and Collect Support Data


Two methods of collecting support data exist in the storage array:

- The AutoSupport feature — Data is automatically collected.
- The Collect Support Data option — Data must be collected and sent manually.

The AutoSupport feature is easier to use because data is collected and sent automatically. AutoSupport data can be used proactively to prevent problems before they occur. The AutoSupport feature speeds troubleshooting because technical support already has access to the data. For these reasons, the AutoSupport feature is the preferred data collection method to use.

Workflow for the AutoSupport feature

In SANtricity System Manager, you configure the AutoSupport feature by following these steps.



How to

View storage array information

View storage array profile

The storage array profile provides a description of all of the components and properties of the storage array.

About this task

You can use the storage array profile as an aid during recovery or as an overview of the current configuration of the storage array. You might want to save a copy of the storage array profile on the management client and keep a hard copy of the storage array profile with the storage array. Create a new copy of the storage array profile if your configuration changes.

Steps

1. Select **Support > Support Center > Support Resources** tab.
2. Scroll down to **Launch detailed storage array information**, and then select **Storage Array Profile**.

The report appears on your screen.

Field details

Section	Description
Storage Array	<p>Shows all of the options that you can configure and the system static options for your storage array. These options include the number of controllers, drive shelves, drives, disk pools, volume groups, volumes, and hot spare drives; the maximum number of drive shelves, drives, Solid State Disks (SSDs), and volumes allowed; the number of snapshot groups, snapshot images, snapshot volumes and consistency groups; information about features; information about firmware versions; information about the chassis serial number; AutoSupport status and AutoSupport schedule information; the settings for automatic support data collection and scheduled support data collection; the storage array World-Wide Identifier (WWID); and the media scan and cache settings.</p>
Storage	<p>Shows a list of all of the storage devices in the storage array. Depending on your storage array configuration, the Storage section might show these sub-sections.</p> <ul style="list-style-type: none">• Disk Pools — Shows a list of all of the disk pools in the storage array.• Volume Groups — Shows a list of all of the volume groups in the storage array. Volumes and free capacity are listed in the order in which they were created.• Volumes — Shows a list of all of the volumes in the storage array. The information listed includes the volume name, the volume status, the capacity, the RAID level, the volume group or disk pool, the drive type, and additional details.• Missing Volumes — Shows a list of all of the volumes in the storage array that currently have a missing status. The information listed includes the World Wide Identifier (WWID) for each missing volume.

Section	Description
Copy Services	<p>Shows a list of all the copy services that are used for the storage array. Depending on your storage array configuration, the Copy Services section might show these sub-sections:</p> <ul style="list-style-type: none"> • Volume Copies — Shows a list of all copy pairs in the storage array. The information listed includes the number of copies, the copy pair names, the status, the start timestamp, and additional details. • Snapshot Groups — Shows a list of all snapshot groups in the storage array. • Snapshot Images — Shows a list of all snapshots in the storage array. • Snapshot Volumes — Shows a list of all snapshot volumes in the storage array. • Consistency Groups — Shows a list of all consistency groups in the storage array. • Member Volumes — Shows a list of all consistency group member volumes in the storage array. • Mirror Groups — Shows a list of all mirrored volumes. • Reserved Capacity — Shows a list of all reserved capacity volumes in the storage array.
Host Assignments	<p>Shows a list of host assignments in the storage array. The information listed includes the volume name, logical unit number (LUN), controller ID, host name or host cluster name, and volume status. Additional information listed includes topology definitions and host type definitions.</p>

Section	Description
Hardware	<p>Shows a list of all of the hardware in the storage array. Depending on your storage array configuration, the Hardware section might show these sub-sections.</p> <ul style="list-style-type: none"> • Controllers — Shows a list of all of the controllers in the storage array and includes the controller location, status, and configuration. In addition, it includes drive channel information, host channel information, and Ethernet port information. • Drives — Shows a list of all of the drives in the storage array. The drives are listed in shelf ID, drawer ID, slot ID order. The information listed includes the shelf ID, the drawer ID, the slot ID, the status, the raw capacity, the media type, the interface type, the current data rate, the product ID, and the firmware version for each drive. The Drive section also includes drive channel information, hot spare coverage information, and wear life information (only for SSD drives). The wear life information includes the percent endurance used, which is the amount of data written to the SSD drives to date, divided by the total theoretical write limit for the drives. • Drive Channels — Shows information for all of the drive channels in the storage array. The information listed includes the channel status, the link status (if applicable), drive counts, and cumulative error counts. • Shelves — Shows information for all of the shelves in the storage array. The information listed includes drive types, and status information for each component of the shelf. Shelf components might include battery packs, Small Form-factor Pluggable (SFP) transceivers, power-fan canisters, or input/output module (IOM) canisters. <p>The Hardware section also shows the security key identifier if a security key is used by the storage array.</p>
Features	<p>Shows a list of the feature packs installed and maximum allowed number of snapshot groups, snapshots (legacy), and volumes per host or host cluster. The information in the Features section also includes Drive Security; that is, whether the storage array is security enabled or security disabled.</p>

3. To search the storage array profile, type a search term in the **Find** text box, and then click **Find**.

All matching terms are highlighted. To scroll through all the results one at a time, continue to click **Find**.

4. To save the storage array profile, click **Save**.

The file is saved in the Downloads folder for your browser with the name `storage-array-profile.txt`.

View software and firmware inventory

The software and firmware inventory lists the firmware versions for each component in

your storage array.

About this task

A storage array is made up of many components, which might include controllers, drives, drawers, and input/output modules (IOMs). Each of these components contains firmware. Some versions of firmware depend on other versions of firmware. To capture information about all of the firmware versions in your storage array, view the software and firmware inventory. Technical support can analyze the software and firmware inventory to detect any firmware mismatches.

Steps

1. Select **Support > Support Center > Support Resources** tab.
2. Scroll down to **Launch detailed storage array information**, and then select **Software and Firmware Inventory**.

The software and firmware inventory report appears on the screen.

3. To save the software and firmware inventory, click **Save**.

The file is saved in the Downloads folder for your browser with the filename `firmware-inventory.txt`.

4. Follow the instructions provided by technical support to send the file to them.

Collect diagnostic data

Collect support data manually

You can gather various types of inventory, status, and performance data about your storage array in a single file. Technical support can use the file for troubleshooting and further analysis.

About this task



If the AutoSupport feature is enabled, you can also collect this data by going to the **AutoSupport** tab and selecting **Send AutoSupport Dispatch**.

You can run only one collection operation at a time. If you try to start another operation, you receive an error message.



Perform this operation only when instructed to do so by technical support.

Steps

1. Select **Support > Support Center > Diagnostics** tab.
2. Select **Collect Support Data**.
3. Click **Collect**.

The file is saved in the Downloads folder for your browser with the name `support-data.7z`. If your shelf contains drawers, the diagnostic data for that shelf is archived in a separate zipped file named `tray-component-state-capture.7z`.

4. Follow the instructions provided by technical support to send the file to them.

Retrieve recovery support files

Technical support can use recovery support files to troubleshoot issues. System Manager automatically saves these files.

Before you begin

Technical support has requested that you send them additional files for troubleshooting.

About this task

Recovery support files include these types of files:

- Support data files
- AutoSupport history
- AutoSupport log
- SAS/RLS diagnostics files
- Recovery profile data
- Database capture files

Steps

1. Select **Support** > **Support Center** > **Diagnostics** tab.
2. Select **Retrieve Recovery Support Files**.

A dialog box lists all the recovery support files that your storage array has collected. To find particular files, you can sort any of the columns or type characters in the **Filter** box.

3. Select a file, and then click **Download**.

The file is saved in the Downloads folder for your browser.

4. If you need to save additional files, repeat the previous step.
5. Click **Close**.
6. Follow the instructions provided by technical support to send the file to them.

Retrieve trace buffers

You can retrieve the trace buffers from the controllers and send the file to technical support for analysis.

About this task

The firmware uses the trace buffers to record processing, especially exception conditions, that might be useful for debugging. You can retrieve trace buffers without interrupting the operation of the storage array and with minimal effect on performance.



Perform this operation only when instructed to do so by technical support.

Steps

1. Select **Support > Support Center > Diagnostics** tab.
2. Select **Retrieve Trace Buffers**.
3. Select the check box next to each controller for which you want to retrieve trace buffers.

You can select one or both controllers. If the controller status message to the right of a check box is Failed or Disabled, the check box is disabled.

4. Click **Yes**.

The file is saved in the Downloads folder for your browser with the filename `trace-buffers.7z`.

5. Follow the instructions provided by technical support to send the file to them.

Collect drive data

You can collect log data from all drives in your storage array and send the file to technical support for analysis.

About this task

Log data consists of statistical information that is maintained by each of the drives in your storage array. Technical support can use this information to analyze the performance of your drives and for troubleshooting problems that might exist.



Perform this operation only when instructed to do so by technical support.

Steps

1. Select **Support > Support Center > Diagnostics** tab.
2. Select **Collect Drive Data**.

A dialog box appears listing all drives in your storage array.

3. In the first column of the table, you can either select individual drives for which you want to collect data (click the check box next to each drive) or select all drives (select the check box in the table header).

To find particular drives, you can sort any of the columns or type characters in the **Filter** box.

4. Click **Collect**.

The file is saved in the Downloads folder for your browser with the name `drive-data.7z`.

5. Follow the instructions provided by technical support to send the file to them.

Collect I/O path statistics

You can save the I/O path statistics file and send it to technical support for analysis.

About this task

Technical support uses the I/O path statistics to help diagnose performance issues. Application performance issues can be caused by memory utilization, CPU utilization, network latency, I/O latency, or other issues. The I/O path statistics are collected automatically during support data collection or you can collect them manually. In addition, if you have AutoSupport turned on, the I/O path statistics are automatically collected and sent to

technical support.

The counters for the I/O path statistics are reset after you confirm that you want to collect the I/O path statistics. The counters are reset even if you subsequently cancel the operation. The counters are also reset when the controller resets (reboots).



Perform this operation only when instructed to do so by technical support.

Steps

1. Select **Support** › **Support Center** › **Diagnostics** tab.
2. Select **Collect I/O Path Statistics**.
3. Confirm that you want to perform the operation by typing `collect`, and then click **Collect**.

The file is saved in the Downloads folder for your browser with the filename `io-path-statistics.7z`.

4. Follow the instructions provided by technical support to send the file to them.

Retrieve health image

You can review a health image for the controller. A health image is a raw data dump of the controller's processor memory that technical support can use to diagnose a problem with a controller.

About this task

The firmware automatically generates a health image when it detects certain errors. After a health image is generated, the controller that had the error reboots and an event is logged in the event log.

If you have AutoSupport turned on, the health image is automatically sent to technical support. If you do not have AutoSupport turned on, you need to contact technical support for instructions on retrieving the health image and sending it to them for analysis.



Perform this operation only when instructed to do so by technical support.

Steps

1. Select **Support** › **Support Center** › **Diagnostics** tab.
2. Select **Retrieve Health Image**.

You can look at the details section to see the size of the health image before downloading the file.

3. Click **Collect**.

The file is saved in the Downloads folder for your browser with the name `health-image.7z`.

4. Follow the instructions provided by technical support to send the file to them.

Take recovery action on storage array status

View unreadable sectors log

You can save the unreadable sectors log and send the file to technical support for

analysis.

About this task

The unreadable sectors log contains detailed records of unreadable sectors caused by drives reporting unrecoverable media errors. Unreadable sectors are detected during normal I/O and during modification operations, such as reconstructions. When unreadable sectors are detected on a storage array, a Needs Attention alert appears for the storage array. The Recovery Guru distinguishes which unreadable sector condition needs attention. Any data contained in an unreadable sector cannot be recovered and should be considered lost.

The unreadable sectors log can store up to 1,000 unreadable sectors. When the unreadable sectors log reaches 1,000 entries, the following conditions apply:

- If new unreadable sectors are detected during reconstruction, the reconstruction fails, and no entry is logged.
- For new unreadable sectors detected during I/O, the I/O fails, and no entry is logged.



These actions include RAID 5 writes and RAID 6 writes that would have succeeded before the overflow.



Possible loss of data — Recovery from unreadable sectors is a complicated procedure that can involve several different methods. Perform this operation only when instructed to do so by technical support.

Steps

1. Select **Support > Support Center > Diagnostics** tab.
2. Select **View/Clear Unreadable Sectors**.
3. To save the unreadable sectors log:
 - a. In the first column of the table, you can either select individual volumes for which you want to save the unreadable sectors log (click the check box next to each volume) or select all volumes (select the check box in the table header).

To find particular volumes, you can sort any of the columns or type characters in the **Filter** box.

- b. Click **Save**.

The file is saved in the Downloads folder for your browser with the name `unreadable-sectors.txt`.

4. If technical support instructs you to clear the unreadable sectors log, perform the following steps:
 - a. In the first column of the table, you can either select individual volumes for which you want to clear the unreadable sectors log (click the check box next to each volume) or select all volumes (select the check box in the table header).
 - b. Click **Clear**, and confirm that you want to perform the operation.

Re-enable drive ports

You can indicate to the controller that corrective action has been taken to recover from a mis-wire condition.

Steps

1. Select **Support > Support Center > Diagnostics** tab.
2. Select **Re-enable Drive Ports**, and confirm that you want to perform the operation.

This option appears only when the storage array has disabled drive ports.

The controller re-enables any SAS ports that were disabled when a mis-wire was detected.

Clear recovery mode

After restoring a storage array configuration, use the Clear Recovery Mode operation to resume I/O on the storage array and return it to normal operations.

Before you begin

- If you want to return the storage array to a previous configuration, you must restore the configuration from the backup before clearing recovery mode.
- You must perform validation checks or check with technical support to make sure that the restore was successful. After determining that the restore was successful, recovery mode can be cleared.

About this task

The storage array contains a configuration database that includes a record of its logical configuration (pools, volume groups, volumes, and so on). If you intentionally clear the storage array configuration or if the configuration database gets corrupted, the storage array enters recovery mode. Recovery mode stops I/O and freezes the configuration database, which gives you time to do one of the following:

- Restore the configuration from the automatic backup that is stored in the controller's flash devices. You must contact technical support to do this.
- Restore the configuration from a previous Save Configuration Database operation. Save Configuration Database operations are performed through the command line interface (CLI).
- Reconfigure the storage array from scratch.

After the storage array configuration has been restored or redefined and you have verified that all is well, you must manually clear recovery mode.



You cannot cancel the Clear Recovery Mode operation after it starts. Clearing recovery mode can take a long time. Perform this operation only when instructed to do so by technical support.

Steps

1. Select **Support > Support Center > Diagnostics** tab.
2. Select **Clear Recovery Mode**, and confirm that you want to perform this operation.

This option appears only if the storage array is in recovery mode.

Manage iSCSI connections

View iSCSI statistics packages

You can view data about the iSCSI connections to your storage array.

About this task

System Manager shows these types of iSCSI statistics. All statistics are read-only and cannot be set.

- **Ethernet MAC statistics** — Provides statistics for the media access control (MAC). MAC also provides an addressing mechanism called the physical address or the MAC address. The MAC address is a unique address that is assigned to each network adapter. The MAC address helps deliver data packets to a destination within the subnetwork.
- **Ethernet TCP/IP statistics** — Provides statistics for the TCP/IP, which is the Transmission Control Protocol (TCP) and Internet Protocol (IP) for the iSCSI device. With TCP, applications on networked hosts can create connections to one another, over which they can exchange data in packets. The IP is a data-oriented protocol that communicates data across a packet-switched inter-network. The IPv4 statistics and the IPv6 statistics are shown separately.
- **Local Target/Initiator (Protocol) statistics** — Shows statistics for the iSCSI target, which provides block level access to its storage media, and shows the iSCSI statistics for the storage array when used as an initiator in asynchronous mirroring operations.
- **DCBX Operational States statistics** — Displays the operational states of the various Data Center Bridging Exchange (DCBX) features.
- **LLDP TLV statistics** — Displays the Link Layer Discovery Protocol (LLDP) Type Length Value (TLV) statistics.
- **DCBX TLV statistics** — Displays the information that identifies the storage array host ports in a Data Center Bridging (DCB) environment. This information is shared with network peers for identification and capability purposes.

You can view each of these statistics as raw statistics or as baseline statistics. Raw statistics are all of the statistics that have been gathered since the controllers were started. Baseline statistics are point-in-time statistics that have been gathered since you set the baseline time.

Steps

1. Select **Support > Support Center > Diagnostics** tab.
2. Select **View iSCSI Statistics Packages**.
3. Click a tab to view the different sets of statistics.
4. To set the baseline, click **Set new baseline**.

Setting the baseline sets a new starting point for the collection of the statistics. The same baseline is used for all iSCSI statistics.

View the different types of iSCSI statistics

You can review different sets of statistics as either raw or baseline statistics: Ethernet MAC statistics, Ethernet TCP/IP statistics, Target (protocol) statistics, Local initiator (protocol) statistics, DCBX operational state statistics, LLDP TLV statistics, and DCBX TLV statistics.

MAC transmit and MAC receive statistics

When you select Ethernet MAC statistics, these MAC transmit statistics appear. You can view each of these statistics as raw statistics or as baseline statistics.

Statistic	Definition
F	Frame count
B	Byte count
MF	Multicast frame count
BF	Broadcast frame count
PF	Pause frame count
CF	Control frame count
FDF	Frame deferral count
FED	Frame excess deferral count
FLC	Frame late collisions count
FA	Frame abort count
FSC	Frame single collision count
FMC	Frame multiple collisions count
FC	Frame collision count
FDR	Frame dropped count
JF	Jumbo frame count

When you select Ethernet MAC statistics, these MAC receive statistics appear.

Statistic	Definition
F	Frame count
B	Byte count
MF	Multicast frame count
BF	Broadcast frame count
PF	Pause frame count

Statistic	Definition
CF	Control frame count
FLE	Frame length error count
FD	Frame dropped count
FCRCE	Frame CRC error count
FEE	Frame encoding error count
LFE	Large frame error count
SFE	Small frame error count
J	Jabber count
UCC	Unknown control frame count
CSE	Carrier sense error count

Ethernet TCP/IP statistics

When you select Ethernet TCP/IP statistics, the TCP statistics in this table appear. You can view each of these statistics as raw statistics or as baseline statistics.

Statistic	Definition
TxS	Transmitted segment count
TxB	Transmitted byte count
RTxTE	Retransmit timer expired count
TxDACK	Transmit delayed ACK count
TxACK	Transmit ACK count
RxS	Received segment count
RxB	Received byte count
RxDACK	Received duplicate ACK count
RxACK	Received ACK count

Statistic	Definition
RxSEC	Received segment error count
RxSOOC	Received segment out-of-order count
RxWP	Received window probe count
RxWU	Received window update count

When you select Ethernet TCP/IP statistics, the IP statistics in this table appear.

Statistic	Definition
TxP	Transmitted packet count
TxB	Transmitted byte count
TxF	Transmitted fragment count
RxP	Packets received count. Select Show IPv4 to show the IPv4 packets received count. Select Show IPv6 to show the IPv6 packets received count.
RxB	Received byte count
RxF	Received fragment count
RxPE	Received packet error count
DR	Datagram reassembly count
DRE-OLFC	Datagram reassembly error, overlapped fragment count
DRE-OOFC	Datagram reassembly error, out-of-order fragment count
DRE-TOC	Datagram reassembly error, time-out count

iSCSI target statistics and local initiator statistics

When you select Target (protocol) statistics or Local initiator (protocol) statistics, the following statistics are shown. You can view each of these statistics as raw statistics or as baseline statistics.

Statistic	Definition
SL	Successful iSCSI login count
UL	Unsuccessful iSCSI login count
SA	Successful iSCSI authentication count (when authentication is enabled)
UA	Unsuccessful iSCSI authentication count (when authentication is enabled)
PDU	Correct iSCSI PDUs processed count
HDE	iSCSI PDUs with header digest errors count
DDE	iSCSI PDUs with data digest errors count
PE	PDUs with iSCSI protocol errors count
UST	Unexpected iSCSI session terminations count
UCT	Unexpected iSCSI connection termination count

DCBX operational state statistics

When you select Data Center Bridging Exchange (DCBX) Operational State Statistics, these statistics appear.

Statistic	Definition
iSCSI Host Port	Indicates the location of the detected host port in Controller #, Port # format.
Priority Group	Indicates the operational state of the Priority Group (PG) application. The state is either Enabled or Disabled.
Priority-based Flow Control	Indicates the operational state of the Priority-based Flow Control (PFC) feature. The state is either Enabled or Disabled.
iSCSI Feature	Indicates the operational state of the Internet Small Computer System Interface (iSCSI) application. The state is either Enabled or Disabled.
FCoE Bandwidth	Indicates the state of the Fibre Channel over Ethernet (FCoE) Bandwidth. The state is either True or False.

Statistic	Definition
No FCoE/FIP Map Mismatch	Indicates whether a map mismatch exists between FCoE and FCoE Initialization Protocol (FIP). The value is either True or False.

You can find additional DCBX operational state statistics in the state capture file.

LLDP TLV statistics

When you select Link Layer Discovery Protocol (LLDP) Type Length Value (TLV) Statistics, these statistics appear. Two sets of statistics appear: one for the local device and one for the remote device. The local device refers to the controller. The remote device refers to the peer device that the controller is attached to, typically a switch.

Statistic	Definition
iSCSI Host Port	Indicates the location of the detected host port in Controller #, Port # format.
Chassis ID	Indicates the chassis ID.
Chassis ID Subtype	Indicates the chassis ID subtype.
Port ID	Indicates the port ID.
Port ID Subtype	Indicates the port ID subtype.
Time to Live	Indicates the number of seconds that the recipient LLDP agent considers the information to be valid.

You can find additional LLDP TLV statistics in the state capture file.

DCBX TLV statistics

When you select Data Center Bridging Exchange (DCBX) Type Length Value (TLV) Statistics, these statistics appear:

- **Local statistics** — The DCBX parameters configured on the controller at the factory.
- **Operational statistics** — The DCBX parameters that result from DCBX negotiations.
- **Remote statistics** — The DCBX parameters from the peer device that the controller is connected to, typically a switch.

Statistic	Definition
iSCSI Host Port	Indicates the location of the detected host port in Controller #, Port # format.

Statistic	Definition
Flow Control Mode	The Flow Control Mode of the entire port. Valid values are Disabled, Standard, Per Priority, or Indeterminate.
Protocol	The communication protocol. Valid values are FCoE, FIP, iSCSI, or UNKNOWN.
Priority	Integer value indicating the priority number of the communication.
Priority Group	Integer value representing the priority group to which the protocol has been assigned.
Priority Group % Bandwidth	Percentage value indicating the amount of bandwidth allocated to the priority group.
DCBX PFC Status	Priority-based Flow Control (PFC) status of the specific port. The value is either enabled or disabled.

You can find additional DCBX TLV statistics in the state capture file.

View iSCSI sessions

You can view detailed information about the iSCSI connections to your storage array. iSCSI sessions can occur with hosts or remote storage arrays in an asynchronous mirror relationship.

Steps

1. Select **Support > Support Center > Diagnostics** tab.
2. Select **View/End iSCSI Sessions**.

A list of the current iSCSI sessions appears.

3. To see additional information about a specific iSCSI session, select a session, and then click **View Details**.

Field Details

Item	Description
Session Identifier (SSID)	A hexadecimal string that identifies a session between an iSCSI initiator and an iSCSI target. The SSID is composed of the ISID and the TPGT.
Initiator Session ID (ISID)	The initiator part of the session identifier. The initiator specifies the ISID during login.
Target Portal Group	The iSCSI target.
Target Portal Group Tag (TPGT)	The target part of the session identifier. A 16-bit numerical identifier for an iSCSI target portal group.
Initiator iSCSI name	The worldwide unique name of the initiator.
Initiator iSCSI label	The user label set in System Manager.
Initiator iSCSI alias	A name that also can be associated with an iSCSI node. The alias allows an organization to associate a user-friendly string with the iSCSI name. However, the alias is not a substitute for the iSCSI name. The initiator iSCSI alias only can be set at the host, not in System Manager
Host	A server that sends input and output to the storage array.
Connection ID (CID)	A unique name for a connection within the session between the initiator and the target. The initiator generates this ID and presents it to the target during login requests. The connection ID is also presented during logouts that close connections.
Ethernet port identifier	The controller port associated with the connection.
Initiator IP address	The IP address of the initiator.
Negotiated login parameters	The parameters that are transacted during the login of the iSCSI session.
Authentication method	The technique to authenticate users who want access to the iSCSI network. Valid values are CHAP and None .
Header digest method	The technique to show possible header values for the iSCSI session. HeaderDigest and DataDigest can be either None or CRC32C . The default value for both is None .
Data digest method	The technique to show possible data values for the iSCSI session. HeaderDigest and DataDigest can be either None or CRC32C . The default value for both is None .

Item	Description
Maximum connections	The greatest number of connections allowed for the iSCSI session. The maximum number of connections can be 1 through 4. The default value is 1 .
Target alias	The label associated with the target.
Initiator alias	The label associated with the initiator.
Target IP address	The IP address of the target for the iSCSI session. DNS names are not supported.
Initial R2T	The initial ready to transfer status. The status can be either Yes or No .
Maximum burst length	The maximum SCSI payload in bytes for this iSCSI session. The maximum burst length can be from 512 to 262,144 (256 KB). The default value is 262,144 (256 KB) .
First burst length	The SCSI payload in bytes for unsolicited data for this iSCSI session. The first burst length can be from 512 to 131,072 (128 KB). The default value is 65,536 (64 KB) .
Default time to wait	The minimum number of seconds to wait before you attempt to make a connection after a connection termination or a connection reset. The default time to wait value can be from 0 to 3600. The default is 2 .
Default time to retain	The maximum number of seconds that connection is still possible following a connection termination or a connection reset. The default time to retain can be from 0 to 3600. The default value is 20 .
Maximum outstanding R2T	The maximum number of "ready to transfers" outstanding for this iSCSI session. The maximum outstanding ready to transfer value can be from 1 to 16. The default is 1 .
Error recovery level	The level of error recovery for this iSCSI session. The error recovery level value is always set to 0 .
Maximum receive data segment length	The maximum amount of data that either the initiator or the target can receive in any iSCSI payload data unit (PDU).
Target name	The official name of the target (not the alias). The target name with the <i>iqn</i> format.
Initiator name	The official name of the initiator (not the alias). The initiator name that uses either the <i>iqn</i> or <i>eui</i> format.

4. To save the report to a file, click **Save**.

The file is saved in the Downloads folder for your browser with the filename `iscsi-session-connections.txt`.

End iSCSI session

You can end an iSCSI session that is no longer needed. iSCSI sessions can occur with hosts or remote storage arrays in an asynchronous mirror relationship.

About this task

You might want to end an iSCSI session for these reasons:

- **Unauthorized access** — If an iSCSI initiator is logged on and should not have access, you can end the iSCSI session to force the iSCSI initiator off the storage array. The iSCSI initiator could have logged on because the None authentication method was available.
- **System downtime** — If you need to take down a storage array and you see that iSCSI initiators are still logged on, you can end the iSCSI sessions to get the iSCSI initiators off the storage array.

Steps

1. Select **Support > Support Center > Diagnostics** tab.
2. Select **View/End iSCSI Sessions**.

A list of the current iSCSI sessions appears.

3. Select the session that you want to end.
4. Click **End Session**, and confirm that you want to perform the operation.

View iSER over InfiniBand statistics

If your storage array's controller includes an iSER over InfiniBand port, you can view data about the host connections.

About this task

System Manager shows the following types of iSER over InfiniBand statistics. All statistics are read-only and cannot be set.

- **Local Target (Protocol) statistics** — Provides statistics for the iSER over InfiniBand target, which shows block-level access to its storage media.
- **iSER over InfiniBand Interface statistics** — Provides statistics for all iSER ports on the InfiniBand interface, which includes performance statistics and link error information associated with each switch port.

You can view each of these statistics as raw statistics or as baseline statistics. Raw statistics are all of the statistics that have been gathered since the controllers were started. Baseline statistics are point-in-time statistics that have been gathered since you set the baseline time.

Steps

1. Select **Support > Support Center > Diagnostics** tab.
2. Select **View iSER over InfiniBand Statistics**.

3. Click a tab to view the different sets of statistics.
4. To set the baseline, click **Set new baseline**.

Setting the baseline sets a new starting point for the collection of the statistics. The same baseline is used for all iSER over InfiniBand statistics.

Manage NVMe connections

View NVMe over Fabrics statistics

You can view data about the NVMe over Fabrics connections to your storage array.

About this task

System Manager shows these types of NVMe over Fabrics statistics. All statistics are read-only and cannot be set.

- **NVMe Subsystem statistics** — Shows statistics for the NVMe controller and its queue. The NVMe controller provides an access path between a host and the namespaces in the storage array. You can review the NVMe subsystem statistics for such items as connection failures, resets, and shutdowns.
- **RDMA Interface statistics** — Provides statistics for all NVMe over Fabrics ports on the RDMA interface, which includes performance statistics and link error information associated with each switch port. This tab only appears when NVMe over Fabrics ports are available.

You can view each of these statistics as raw statistics or as baseline statistics. Raw statistics are all of the statistics that have been gathered since the controllers were started. Baseline statistics are point-in-time statistics that have been gathered since you set the baseline time.

Steps

1. Select **Support > Support Center > Diagnostics** tab.
2. Select **View NVMe over Fabrics Statistics**.
3. To set the baseline, click **Set new baseline**.

Setting the baseline sets a new starting point for the collection of the statistics. The same baseline is used for all NVMe statistics.

Manage AutoSupport

Enable or disable AutoSupport features

You enable the AutoSupport feature and the individual AutoSupport features during initial setup or you can enable them later.

Before you begin

If you want to enable either AutoSupport OnDemand or Remote Diagnostics, the AutoSupport delivery method must be set to HTTPS.

About this task

You can disable the AutoSupport feature at any time, but you are strongly advised to leave it enabled. Enabling the AutoSupport feature can significantly speed problem determination and resolution should a problem occur on your storage array.

The AutoSupport feature is made up of three individual features that you enable separately.

- **Basic AutoSupport** — Allows your storage array to automatically collect and send data to technical support.
- **AutoSupport OnDemand** — Allows technical support to request retransmission of a previous AutoSupport dispatch when needed for troubleshooting an issue. All transmissions are initiated from the storage array, not from the AutoSupport server. The storage array checks in periodically with the AutoSupport server to determine if there are any pending retransmission requests and responds accordingly.
- **Remote Diagnostics** — Allows technical support to request a new, up-to-date AutoSupport dispatch when needed for troubleshooting an issue. All transmissions are initiated from the storage array, not from the AutoSupport server. The storage array checks in periodically with the AutoSupport server to determine if there are any pending new requests and responds accordingly.

Steps

1. Select **Support > Support Center > AutoSupport** tab.
2. Select **Enable/Disable AutoSupport Features**.
3. Select the check boxes next to the AutoSupport features that you want to enable.

The features depend on each other as indicated by the indentation of the items in the dialog box. For example, you must enable AutoSupport OnDemand before you can enable Remote Diagnostics.

4. Click **Save**.

Enable AutoSupport maintenance window

Enable the AutoSupport maintenance window to suppress automatic ticket creation on error events. Under normal operation mode, the storage array uses AutoSupport to open a case with Support if there is an issue.

Steps

1. Select **Support > Support Center > AutoSupport** tab.
2. Select **Enable AutoSupport Maintenance window**.
3. Enter the email address to receive a confirmation that the maintenance window request has been processed.

Depending on your configuration, you may be able to enter up to five email addresses. If you want to add more than one address, select **Add another email** to open another field.

4. Specify the duration (in hours) to enable the maintenance window.

The maximum supported duration is 72 hours.

5. Click **Yes**.

AutoSupport automatic ticket creation on error events is temporarily suppressed for the specified duration window.

After you finish

The maintenance window does not begin until the storage array's request is processed by the AutoSupport servers. Wait until you have received a confirmation email before performing any maintenance activities on

your storage array.

Disable AutoSupport maintenance window

Disable the AutoSupport maintenance window to allow automatic ticket creation on error events. When AutoSupport maintenance window is disabled, the storage array will use AutoSupport to open a case with Support if there is an issue.

Steps

1. Select **Support > Support Center > AutoSupport** tab.
2. Select **Disable AutoSupport Maintenance window**.
3. Enter the email address to receive a confirmation that the disable maintenance window request has been processed.

Depending on your configuration, you may be able to enter up to five email addresses. If you want to add more than one address, select **Add another email** to open another field.

4. Click **Yes**.

AutoSupport automatic ticket creation on error events is enabled.

After you finish

The maintenance window will not end until the storage array's request has been processed by the AutoSupport servers. Wait until you have received a confirmation email before proceeding.

Configure AutoSupport delivery method

The AutoSupport feature supports the HTTPS, HTTP, and SMTP protocols for delivering dispatches to technical support.

Before you begin

- The AutoSupport feature must be enabled. You can see whether it is enabled on the AutoSupport page.
- A DNS server must be installed and configured in your network. The DNS server address must be configured in System Manager (this task is available from the Hardware page).

About this task

Review the different protocols:

- **HTTPS** — Allows you to connect directly to the destination technical support server using HTTPS. If you want to enable either AutoSupport OnDemand or Remote Diagnostics, the AutoSupport delivery method must be set to HTTPS.
- **HTTP** — Allows you to connect directly to the destination technical support server using HTTP.
- **Email** — Allows you to use an email server as the delivery method for sending AutoSupport dispatches.



Differences between the HTTPS/HTTP and Email methods. The Email delivery method, which uses SMTP, has some important differences from the HTTPS and HTTP delivery methods. First, the size of the dispatches for the Email method are limited to 5MB, which means that some ASUP data collections will not be dispatched. Second, the AutoSupport OnDemand feature is available only on HTTP and HTTPS methods.

Steps

1. Select **Support > Support Center > AutoSupport** tab.
2. Select **Configure AutoSupport Delivery Method**.

A dialog box appears, which lists the dispatch delivery methods.

3. Select the desired delivery method, and then select the parameters for that delivery method. Do one of the following:
 - If you selected HTTPS or HTTP, select one of the following delivery parameters:
 - **Directly** — This delivery parameter is the default selection. Choosing this option allows you to connect directly to the destination technical support system using the HTTPS or HTTP protocol.
 - **Via Proxy server** — Choosing this option allows you to specify the HTTP proxy server details required for establishing connection with the destination technical support system. You must specify the host address and port number. However, you only need to enter the host authentication details (user name and password) if required.
 - **Via Proxy auto-configuration script (PAC)** — Specify the location of a Proxy Auto-Configuration (PAC) Script file. A PAC file allows the system to automatically choose the appropriate proxy server for establishing a connection with the destination technical support system.
 - If you selected Email, enter the following information:
 - The mail server address as a fully qualified domain name, IPv4 address, or IPv6 address.
 - The email address that appears in the From field of the AutoSupport dispatch email.
 - **Optional:** If you want to perform a configuration test.) The email address where a confirmation is sent when the AutoSupport system receives the test dispatch.
4. Click **Test Configuration** to test the connection to the technical support server using the specified delivery parameters. If you enabled the AutoSupport On-Demand feature, the system will also test the connection for AutoSupport OnDemand dispatch delivery.

If the configuration test fails, check your configuration settings and run the test again. If the test continues to fail, contact technical support.

5. Click **Save**.

Schedule AutoSupport dispatches

System Manager automatically creates a default schedule for AutoSupport dispatches. If you prefer, you can specify your own schedule.

Before you begin

The AutoSupport feature must be enabled. You can see whether it is enabled on the AutoSupport page.

About this task

- **Daily time** — Daily dispatches are collected and sent every day during the time range that you specify. System Manager selects a random time during the range. All times are in Coordinated Universal Time (UTC), which might be different from the storage array's local time. You must convert your storage array's local time into UTC.
- **Weekly day** — Weekly dispatches are collected and sent once a week. System Manager selects a random day from the days that you specify. Deselect any days on which you do not want to allow a weekly dispatch to occur. System Manager selects a random day from the days that you allow.

- **Weekly time** — Weekly dispatches are collected and sent once a week during the time range that you specify. System Manager selects a random time during the range. All times are in Coordinated Universal Time (UTC), which might be different from the storage array's local time. You must convert your storage array's local time into UTC.

Steps

1. Select **Support** › **Support Center** › **AutoSupport** tab.
2. Select **Schedule AutoSupport Dispatches**.

The Schedule AutoSupport Dispatches wizard appears.

3. Follow the steps in the wizard.

Send AutoSupport dispatches

System Manager allows you to send AutoSupport dispatches to technical support, without waiting for a scheduled dispatch.

Before you begin

The AutoSupport feature must be enabled. You can see whether it is enabled on the AutoSupport page.

About this task

This operation collects support data and automatically sends it to technical support, so they can troubleshoot issues.

Steps

1. Select **Support** › **Support Center** › **AutoSupport** tab.
2. Select **Send AutoSupport Dispatch**.

The **Send AutoSupport Dispatch** dialog box appears.

3. Confirm the operation by selecting **Send**.

View AutoSupport status

The AutoSupport page shows you whether the AutoSupport feature and the individual AutoSupport features are currently enabled.

Steps

1. Select **Support** › **Support Center** › **AutoSupport** tab.
2. Look at the right side of the page just below the tabs to see whether the basic AutoSupport feature is enabled.
3. Hover your cursor over the question mark to see whether individual AutoSupport features are enabled.

View AutoSupport log

The AutoSupport log provides information about status, dispatch history, and errors encountered during the delivery of AutoSupport dispatches.

About this task

Multiple log files can exist. When the current log file reaches 200 KB, it is archived and a new log file is created. The archived log file name is `ASUPMessages.n`, where `n` is an integer from 1 to 9. If multiple log files exist, you can choose to view the most current log or a previous log.

- **Current log** — Shows a list of the latest captured events.
- **Archived log** — Shows a list of earlier events.

Steps

1. Select **Support > Support Center > AutoSupport** tab.
2. Select **View AutoSupport Log**.

A dialog box appears, which lists the current AutoSupport log.

3. If you want to see previous AutoSupport logs, select the **Archived** radio button, and then select a log from the **Select AutoSupport log** drop-down list.

The Archived option appears only if archived logs exist on the storage array.

The selected AutoSupport log appears in the dialog box.

4. **Optional:** To search the AutoSupport log, type a term in the **Find** box, and click **Find**.

Click **Find** again to search for additional occurrences of the term.

FAQs

What data am I collecting?

The AutoSupport feature and the manual Support Data Collection feature provide ways to collect data in a customer support bundle for remote troubleshooting and problem analysis by technical support.

The customer support bundle gathers all types of information about the storage array into a single compressed file. The information collected includes the physical configuration, logical configuration, version information, events, log files, and performance data. The information is used only by technical support to solve problems with the storage array.

What does unreadable sectors data show me?

You can display detailed data about unreadable sectors detected on the drives in your storage array.

The unreadable sectors log shows the most recent unreadable sector first. The log contains the following information about the volumes that contain the unreadable sectors. The fields are sortable.

Field	Description
Affected Volume	Shows the label of the volume. If a missing volume contains unreadable sectors, the World Wide Identifier appears for the missing volume.

Field	Description
Logical Unit Number (LUN)	Shows the LUN for the volume. If the volume does not have a LUN, the dialog box shows NA.
Assigned To	Shows the hosts or host clusters that have access to the volume. If the volume is not accessible by a host, host cluster, or even a Default Cluster, the dialog box shows NA.

To see additional information about the unreadable sectors, click the plus (+) sign next to a volume.

Field	Description
Date/Time	Shows the date and the time that the unreadable sector was detected.
Volume Logical Block Address	Shows the logical block address (LBA) of the volume.
Drive Location	Shows the drive shelf, drawer (if your drive shelf has drawers), and the bay location.
Drive Logical Block Address	Shows the LBA of the drive.
Failure Type	Shows one of the following failure types: <ul style="list-style-type: none"> • Physical — A physical media error. • Logical — A read error elsewhere in the stripe causing unreadable data. For example, an unreadable sector due to media errors elsewhere in the volume. • Inconsistent — Inconsistent redundancy data. • Data Assurance — A Data Assurance error.

What is a health image?

A health image is a raw data dump of the controller's processor memory that technical support can use to diagnose a problem with a controller.

The firmware automatically generates a health image when it detects certain errors. Under certain troubleshooting scenarios, technical support might request that you retrieve the health image file and send it to them.

What else do I need to do to configure or diagnose iSCSI?

iSCSI sessions can occur with hosts or remote storage arrays in an asynchronous mirror relationship. The following tables list the System Manager functions that you can use to configure and manage these iSCSI sessions.



The iSCSI settings are only available if your storage array supports iSCSI.

Configure iSCSI

Action	Location
Manage iSCSI settings	<ol style="list-style-type: none">1. Select Settings > System.2. Scroll down to iSCSI settings to view all the management functions.
Configure iSCSI ports	<ol style="list-style-type: none">1. Select Hardware.2. Select Show back of shelf.3. Select a controller.4. Select Configure iSCSI ports.
Set the host CHAP secret	<ol style="list-style-type: none">1. Select Settings > System.2. Scroll down to iSCSI settings, and then select Configure Authentication. <p>or</p> <ol style="list-style-type: none">1. Select Storage > Hosts.2. Select a host member.3. Click View/Edit Settings > Host Ports tab.

Diagnose iSCSI

Action	Location
View or end iSCSI sessions	<ol style="list-style-type: none">1. Select Settings > System.2. Scroll down to iSCSI settings, and then select View/End iSCSI Sessions. <p>or</p> <ol style="list-style-type: none">1. Select Support > Support Center > Diagnostics tab.2. Select View/End iSCSI Sessions.

Action	Location
View iSCSI statistics	<ol style="list-style-type: none"> 1. Select Settings › System. 2. Scroll down to iSCSI settings, and then select View iSCSI Statistics Packages. <p>or</p> <ol style="list-style-type: none"> 1. Select Support › Support Center › Diagnostics tab. 2. Select View iSCSI Statistics Packages.

What do the AutoSupport features do?

The AutoSupport feature is made up of three individual features that you enable separately.

- **Basic AutoSupport** — Allows your storage array to automatically collect and send data to technical support.
- **AutoSupport OnDemand** — Allows technical support to request retransmission of a previous AutoSupport dispatch when needed for troubleshooting an issue. All transmissions are initiated from the storage array, not from the AutoSupport server. The storage array checks in periodically with the AutoSupport server to determine if there are any pending retransmission requests and responds accordingly.
- **Remote Diagnostics** — Allows technical support to request a new, up-to-date AutoSupport dispatch when needed for troubleshooting an issue. All transmissions are initiated from the storage array, not from the AutoSupport server. The storage array checks in periodically with the AutoSupport server to determine if there are any pending new requests and responds accordingly.

What type of data is collected through the AutoSupport feature?

The AutoSupport feature contains three standard dispatch types: event dispatches, scheduled dispatches, and on-demand and remote diagnostics dispatches.

The AutoSupport data does not contain any user data.

• Event dispatches

When events occur on the system that warrant proactive notification to technical support, the AutoSupport feature automatically sends an event-triggered dispatch.

- Sent when a support event on the managed storage array occurs.
- Includes a comprehensive snapshot of what was going on with the storage array at the time the event occurred.

• Scheduled dispatches

The AutoSupport feature automatically sends several dispatches on a regular schedule.

- **Daily dispatches** — Sent once every day during a user-configurable time interval. Includes the current system event logs and performance data.
- **Weekly dispatches** — Sent once every week during a user-configurable time interval and day.

Includes configuration and system state information.

- **AutoSupport OnDemand and Remote Diagnostics dispatches**

- **AutoSupport OnDemand** — Allows technical support to request retransmission of a previous AutoSupport dispatch when needed for troubleshooting an issue. All transmissions are initiated from the storage array, not from the AutoSupport server. The storage array checks in periodically with the AutoSupport server to determine if there are any pending retransmission requests and responds accordingly.
- **Remote Diagnostics** — Allows technical support to request a new, up-to-date AutoSupport dispatch when needed for troubleshooting an issue. All transmissions are initiated from the storage array, not from the AutoSupport server. The storage array checks in periodically with the AutoSupport server to determine if there are any pending new requests and responds accordingly.

How do I configure the delivery method for the AutoSupport feature?

The AutoSupport feature supports the protocols HTTPS, HTTP, and SMTP for delivering AutoSupport dispatches to technical support.

Before you begin

- The AutoSupport feature must be enabled. You can see whether it is enabled on the AutoSupport page.
- A DNS server must be installed and configured in your network. The DNS server address must be configured in System Manager (this task is available from the Hardware page).

About this task

Review the different protocols:

- **HTTPS** — Allows you to connect directly to the destination technical support server using HTTPS. If you want to enable either AutoSupport OnDemand or Remote Diagnostics, the AutoSupport delivery method must be set to HTTPS.
- **HTTP** — Allows you to connect directly to the destination technical support server using HTTP.
- **Email** — Allows you to use an email server as the delivery method for sending AutoSupport dispatches.



Differences between the HTTPS/HTTP and Email methods. The Email delivery method, which uses SMTP, has some important differences from the HTTPs and HTTP delivery methods. First, the size of the dispatches for the Email method are limited to 5MB, which means that some ASUP data collections will not be dispatched. Second, the AutoSupport OnDemand feature is available only on HTTP and HTTPS methods.

Steps

1. Select **Support > Support Center > AutoSupport** tab.
2. Select **Configure AutoSupport Delivery Method**.

A dialog box appears, which lists the dispatch delivery methods.

3. Select the desired delivery method, and then select the parameters for that delivery method. Do one of the following:
 - If you selected HTTPS or HTTP, select one of the following delivery parameters:
 - **Directly** — This delivery parameter is the default selection. Choosing this option allows you to connect directly to the destination technical support system using the HTTPS or HTTP protocol.

- **Via Proxy server** — Choosing this option allows you to specify the HTTP proxy server details required for establishing connection with the destination technical support system. You must specify the host address and port number. However, you only need to enter the host authentication details (user name and password) if required.
- **Via Proxy auto-configuration script (PAC)** — Specify the location of a Proxy Auto-Configuration (PAC) Script file. A PAC file allows the system to automatically choose the appropriate proxy server for establishing a connection with the destination technical support system.

◦ If you selected Email, enter the following information:

- The mail server address as a fully qualified domain name, IPv4 address, or IPv6 address.
- The email address that appears in the From field of the AutoSupport dispatch email.
- **Optional:** If you want to perform a configuration test.) The email address where a confirmation is sent when the AutoSupport system receives the test dispatch.

4. Click **Test Configuration** to test the connection to the technical support server using the specified delivery parameters. If you enabled the AutoSupport On-Demand feature, the system will also test the connection for AutoSupport OnDemand dispatch delivery.

If the configuration test fails, check your configuration settings and run the test again. If the test continues to fail, contact technical support.

5. Click **Save**.

Event log

Concepts

Critical events

Critical events indicate a problem with the storage array. If you resolve the critical event immediately, you might prevent loss of data access.

When a critical event occurs, it is logged in the event log. All critical events are sent to the SNMP management console or to the email recipient that you have configured to receive alert notifications. If the shelf ID is not known at the time of the event, the shelf ID is listed as "Shelf unknown."

When you receive a critical event, refer to the Recovery Guru procedure for a detailed description of the critical event. Complete the Recovery Guru procedure to correct the critical event. To correct certain critical events, you might need to contact technical support.

Event log

The event log provides a historical record of events that have occurred on the storage array and aids technical support in troubleshooting events leading up to failures.

The event log is a detailed record of events that occur in the storage array. It records configuration events and storage array component failures. You can use the event log as a supplementary diagnostic tool to the Recovery Guru for tracing storage array events. Always refer to the Recovery Guru first when you attempt to recover from component failures in the storage array.

The events in the event log are categorized with different statuses. Events that you need to take action on have the following statuses:

- Critical
- Warning

Events that are informational and do not require any immediate action are the following:

- Informational

How to

View events using the event log


You can view the event log, which provides a historical record of events that have occurred on the storage array.

Steps

1. Select **Support > Event Log**.

The Event Log page appears.

Page details

Item	Description
View All field	Toggles between all events, and only the critical and warning events.
Filter field	Filters the events. Useful for displaying only events related to a specific component, a specific event, etc.
Select columns icon.	Allows you to select other columns to view. Other columns give you additional information about the event.
Check boxes	Allows you to select the events to save. The check box in the table header selects all events.
Date/Time column	<p>The date and time stamp of the event, according to the controller clock.</p> <div>  <p>The event log initially sorts events based on sequence number. Usually, this sequence corresponds to the date and time. However, the two controller clocks in the storage array could be unsynchronized. In this case, some perceived inconsistencies could appear in the event log relative to events and the date and time shown.</p> </div>

Item	Description
Priority column	<p>These priority values exist:</p> <ul style="list-style-type: none"> • Critical — A problem exists with the storage array. However, if you take immediate action, you might prevent losing access to data. Critical events are used for alert notifications. All critical events are sent to any network management client (through SNMP traps) or to the email recipient that you configured. • Warning — An error has occurred that has degraded the performance and the ability of the storage array to recover from another error. • Informational — Non-critical information related to the storage array.
Component Type column	The component that is affected by the event. The component could be hardware, such as a drive or a controller, or it could be software, such as controller firmware.
Component Location column	The physical location of the component in the storage array.
Description column	<p>A description of the event.</p> <p>Example — Drive write failure - retries exhausted</p>
Sequence Number column	A 64-bit number that uniquely identifies a specific log entry for a storage array. This number increments by one with every new event log entry. To display this information, click the Select columns icon.
Event Type column	A 4-digit number that identifies each type of logged event. To display this information, click the Select columns icon.
Event Specific Codes column	This information is used by technical support. To display this information, click the Select columns icon.

Item	Description
Event Category column	<ul style="list-style-type: none"> • Failure – A component in the storage array has failed; for example, drive failure or battery failure. • State Change – An element of the storage array that has changed state; for example, a volume transitioned to Optimal status, or a controller transitioned to Offline status. • Internal – Internal controller operations that do not require user action; for example, the controller has completed start-of-day. • Command – A command that has been issued to the storage array; for example, a hot spare has been assigned. • Error – An error condition has been detected on the storage array; for example, a controller is unable to synchronize and purge cache, or a redundancy error is detected on the storage array. • General – Any event that does not fit well into any other category. To display this information, click the Select columns icon.
Logged By column	The name of the controller that logged the event. To display this information, click the Select columns icon.

2. To retrieve new events from the storage array, click **Refresh**. It can take several minutes for an event to be logged and become visible in the Event Log page.
3. To save the event log to a file:

- a. Select the check box next to each event that you want to save.
- b. Click **Save**.

The file is saved in the Downloads folder for your browser with the name `major-event-log-timestamp.log`.

4. To clear events from the event log:

The event log stores approximately 8,000 events before it replaces an event with a new event. If you want to keep the events, you can save them, and clear them from the event log.

- a. First, save the event log.
- b. Click **Clear All**, and confirm that you want to perform the operation.

Upgrade center

Concepts

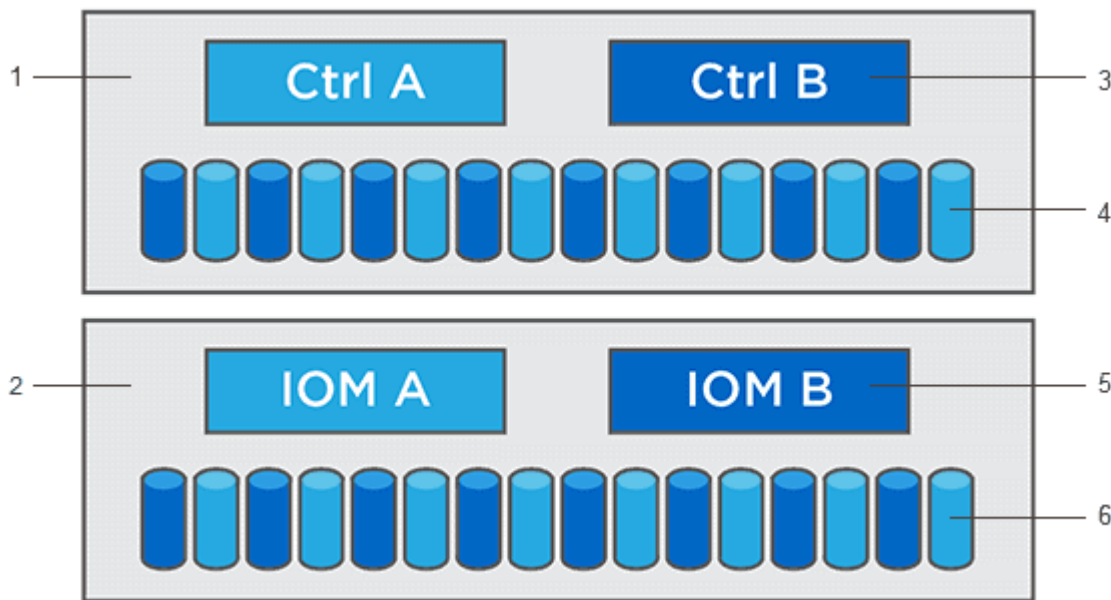
Controller software and firmware upgrades

You can upgrade your storage array's software and firmware for all the latest features and bug fixes.

Components included in the SANtricity OS controller software upgrade

Several storage array components contain software or hardware that you might want to upgrade occasionally.

- **Management software** — System Manager is the software that manages the storage array.
- **Controller firmware** — Controller firmware manages the I/O between hosts and volumes.
- **Controller NVSRAM** — Controller NVSRAM is a controller file that specifies the default settings for the controllers.
- **IOM firmware** — The I/O module (IOM) firmware manages the connection between a controller and a drive shelf. It also monitors the status of the components.
- **Supervisor software** — Supervisor software is the virtual machine on a controller in which the software runs.



¹ Controller shelf; ² Drive shelf; ³ Software, controller firmware, controller NVSRAM, supervisor software; ⁴ Drive firmware; ⁵ IOM firmware; ⁶ Drive firmware

You can view your current software and firmware versions in the **Software and Firmware Inventory** dialog box. Go to **Support > Upgrade Center**, and then click the link for **Software and Firmware Inventory**.

As part of the upgrade process, the host's multipath/failover driver and/or HBA driver might also need to be upgraded so the host can interact with the controllers correctly. To determine if this is the case, see the [Netapp Interoperability Matrix Tool](#).

When to stop I/O

If your storage array contains two controllers and you have a multipath driver installed, the storage array can

continue processing I/O while the upgrade occurs. During the upgrade, controller A fails over all of its volumes to controller B, upgrades, takes back its volumes and all of controller B's volumes, and then upgrades controller B.

Pre-upgrade health check


A pre-upgrade health check runs as part of the upgrade process. The pre-upgrade health check assesses all storage array components to make sure the upgrade can proceed. The following conditions might prevent the upgrade:

- Failed assigned drives
- Hot spares in use
- Incomplete volume groups
- Exclusive operations running
- Missing volumes
- Controller in Non-optimal status
- Excess number of event log events
- Configuration database validation failure
- Drives with old versions of DACstore

You also can run the pre-upgrade health check separately without doing an upgrade.

Workflow for controller software and firmware upgrade

In SANtricity System Manager, you can upgrade the controller software and firmware by following these steps.



Drive firmware upgrades

Drive firmware controls the low-level operating characteristics of a drive. Periodically, the drive manufacturers release updates to drive firmware to add new features, improve performance, and fix defects.

Online and offline drive firmware upgrades

There are two types of drive firmware upgrade methods: online and offline.

Online

During an online upgrade, drives are upgraded sequentially, one at a time. The storage array continues processing I/O while the upgrade occurs. You do not have to stop I/O. If a drive can do an online upgrade, the online method is used automatically.

Drives that can do an online upgrade include the following:

- Drives in an Optimal pool
- Drives in an Optimal redundant volume group (RAID 1, RAID 5, and RAID 6)
- Unassigned drives
- Standby hot spare drives

Doing an online drive firmware upgrade can take several hours exposing the storage array to potential volume failures. Volume failure could occur in these cases:

- In a RAID 1 or RAID 5 volume group, one drive fails while a different drive in the volume group is being upgraded.
- In a RAID 6 pool or volume group, two drives fail while a different drive in the pool or volume group is being upgraded.

Offline (parallel)

During an offline upgrade, all drives of the same drive type are upgraded at the same time. This method requires stopping I/O activity to the volumes associated with the selected drives. Because multiple drives can be upgraded concurrently (in parallel), the overall downtime is significantly reduced. If a drive can do only an offline upgrade, the offline method is used automatically.

The following drives **MUST** use the offline method:

- Drives in a non-redundant volume group (RAID 0)
- Drives in a non-optimal pool or volume group
- Drives in SSD cache

Compatibility

Each drive firmware file contains information about the drive type on which the firmware runs. You can download the specified firmware file only to a compatible drive. System Manager automatically checks compatibility during the upgrade process.

How to

Upgrade software and firmware

You can upgrade your storage array's software and, optionally, the IOM firmware and the nonvolatile static random access memory (NVS RAM) to make sure you have all the latest features and bug fixes.

Before you begin

- You know whether you want to upgrade your IOM firmware.

Normally, you should upgrade all components at the same time. However, you might decide not to upgrade the IOM firmware if you do not want to upgrade it as part of the SANtricity OS software upgrade or if technical support has instructed you to downgrade your IOM firmware (you can only downgrade firmware by using the command line interface).

- You know whether you want to upgrade the controller NVSRAM file.

Normally, you should upgrade all components at the same time. However, you might decide not to upgrade the controller NVSRAM file if your file has either been patched or is a custom version and you do not want to overwrite it.

- You know whether you want to activate your OS upgrade now or later.

Reasons for activating later might include:

- **Time of day** — Activating the software and firmware can take a long time, so you might want to wait until I/O loads are lighter. The controllers fail over during activation so performance might be lower than usual until the upgrade completes.
- **Type of package** — You might want to test the new software and firmware on one storage array before upgrading the files on other storage arrays.
- You know whether you want to switch from unsecured drives or internally secured drives to use an external key management server (KMS) for drive security.
- You know whether you want to use role-based access control in your storage array.

About this task

You can choose to upgrade only the OS software file or only the Controller NVSRAM file or you can choose to upgrade both files.

Perform this operation only when instructed to do so by technical support.



Risk of data loss or risk of damage to the storage array — Do not make changes to the storage array while the upgrade is occurring. Maintain power to the storage array.

Steps

1. If your storage array contains only one controller or you do not have a multipath driver installed, stop I/O activity to the storage array to prevent application errors. If your storage array has two controllers and you have a multipath driver installed, you do not need to stop I/O activity.
2. Select **Support > Upgrade Center**.
3. Download the new file from the Support site to your management client.
 - a. In the area labeled SANtricity OS Controller Software upgrade, click **NetApp Support**.
 - b. On the Support web site, click the **Downloads** tab, and then select **Software**.
 - c. Select **SANtricity OS Controller Software**.
 - d. Follow the remaining instructions.



Digitally signed firmware is required in version 8.42 and above. If you attempt to download unsigned firmware, an error is displayed and the download is aborted.

4. If you do NOT want to upgrade the IOM firmware at this time, click **Suspend IOM Auto-Synchronization**.

If you have a storage array with a single controller, the IOM firmware is not upgraded.

5. Under SANtricity OS Software upgrade, click **Begin Upgrade**.

The **Upgrade SANtricity OS Software** dialog box appears.

6. Select one or more files to begin the upgrade process:

- a. Select the SANtricity OS Software file by clicking **Browse** and navigating to the OS software file you downloaded from the Support web site.
- b. Select the Controller NVSRAM file by clicking **Browse** and navigating to the NVSRAM file that you downloaded from the Support site. Controller NVSRAM files have a filename similar to N2800-830000-000.dlp.

These actions occur:

- By default, only the files that are compatible with the current storage array configuration appear.
- When you select a file for upgrade, the file's name and size appear.

7. **Optional:** If you selected a SANtricity OS Software file to upgrade, you can transfer the files to the controller without activating them by selecting the **Transfer files now, but do not upgrade (activate upgrade later)** check box.
8. Click **Start**, and confirm that you want to perform the operation.

You can cancel the operation during the pre-upgrade health check, but not during transferring or activating.

9. **Optional:** To see a list of what was upgraded, click **Save Log**.

The file is saved in the Downloads folder for your browser with the namedrive_upgrade_log-timestamp.txt.

After you finish

- Verify that all components appear on the Hardware page.
- Verify the new software and firmware versions by checking the **Software and Firmware Inventory** dialog box (go to **Support > Upgrade Center**, and then click the link for **Software and Firmware Inventory**).
- If you upgraded controller NVSRAM, any custom settings that you have applied to the existing NVSRAM are lost during the process of activation. You need to apply the custom settings to the NVSRAM again after the process of activation is complete.

Activate controller software and firmware

You can choose to activate the upgrade files immediately or wait until a more convenient time.

About this task

You can download and transfer the files without activating them. You might choose to activate later for these reasons:

- **Time of day** — Activating the software and firmware can take a long time, so you might want to wait until I/O loads are lighter. The controllers fail over during activation so performance might be lower than usual

until the upgrade completes.

- **Type of package** — You might want to test the new software and firmware on one storage array before upgrading the files on other storage arrays.

When you have software or firmware that has been transferred but not activated, you see a notification in the Notifications area of the System Manager Home page and also on the Upgrade Center page.



You cannot stop the activation process after it starts.

Steps

1. Select **Support › Upgrade Center**.
2. In the area labeled SANtricity OS Controller Software upgrade, click **Activate**, and confirm that you want to perform the operation.

You can cancel the operation during the pre-upgrade health check, but not during activating.

The pre-upgrade health check begins. If the pre-upgrade health check passes, the upgrade process proceeds to activating the files. If the pre-upgrade health check fails, use the Recovery Guru or contact technical support to resolve the problem.

On successful completion of the pre-upgrade health check, activation occurs. The time it takes to activate depends on your storage array configuration and the components that you are activating.

3. **Optional:** To see a list of what was upgraded, click **Save Log**.

The file is saved in the Downloads folder for your browser with the `namedrive_upgrade_log-timestamp.txt`.

After you finish

- Verify that all components appear on the Hardware page.
- Verify the new software and firmware versions by checking the **Software and Firmware Inventory** dialog box (go to **Support › Upgrade Center**, and then click the link for **Software and Firmware Inventory**).
- If you upgraded controller NVSRAM, any custom settings that you have applied to the existing NVSRAM are lost during the process of activation. You need to apply the custom settings to the NVSRAM again after the process of activation is complete.

Upgrade drive firmware

You can upgrade your drives' firmware to make sure you have all the latest features and bug fixes.

Before you begin

- You have backed up your data using disk-to-disk backup, volume copy (to a volume group not affected by the planned firmware upgrade), or a remote mirror.
- The storage array has an Optimal status.
- All drives have an Optimal status.
- No configuration changes are running on the storage array.
- If the drives are capable of only an offline upgrade, I/O activity to all volumes associated with the drives is stopped.

Steps

1. Select **Support > Upgrade Center**.
2. Download the new files from the Support site to your management client.
 - a. Under Drive Firmware upgrade, click **NetApp Support**.
 - b. On the NetApp Support web site, click the **Downloads** tab, and then select **Firmware**.
 - c. Select **Disk Drive & Firmware Matrix**.
 - d. Follow the remaining instructions.
3. Under Drive Firmware upgrade, click **Begin Upgrade**.

A dialog box appears, which lists the drive firmware files currently in use.

4. Extract (unzip) the files you downloaded from the Support site.
5. Click **Browse**, and select the new drive firmware files that you downloaded from the Support site.

Drive firmware files have a filename similar to D_HUC101212CSS600_30602291_MS01_2800_0002 with the extension of .dlp.

You can select up to four drive firmware files, one at a time. If more than one drive firmware file is compatible with the same drive, you get a file conflict error. Decide which drive firmware file you want to use for the upgrade and remove the other one.

6. Click **Next**.

The **Select Drives** dialog box appears, which lists the drives that you can upgrade with the selected files.

Only drives that are compatible appear.

The selected firmware for the drive appears in the **Proposed Firmware** information area. If you must change the firmware, click **Back** to return to the previous dialog.

7. Select the type of upgrade you want to perform:
 - **Online (default)**- Shows the drives that can support a firmware download *while the storage array is processing I/O*. You do not have to stop I/O to the associated volumes using these drives when you select this upgrade method. These drives are upgraded one at a time while the storage array is processing I/O to those drives.
 - **Offline (parallel)** - Shows the drives that can support a firmware download *only while all I/O activity is stopped* on any volumes that use the drives. You must stop all I/O activity on any volumes that use the drives you are upgrading when you select this upgrade method. Drives that do not have redundancy must be processed as an offline operation. This requirement includes any drive associated with SSD cache, a RAID 0 volume group, or any pool or volume group that is degraded. The offline (parallel) upgrade is typically faster than the online (default) method.
8. In the first column of the table, select the drive or drives you want to upgrade.
9. Click **Start**, and confirm that you want to perform the operation.

If you need to stop the upgrade, click **Stop**. Any firmware downloads currently in progress complete. Any firmware downloads that have not started are canceled.



Stopping the drive firmware upgrade might result in data loss or unavailable drives.

10. **Optional:** To see a list of what was upgraded, click **Save Log**.

The file is saved in the Downloads folder for your browser with the namedrive_upgrade_log-timestamp.txt.

11. If any of the following errors occur during the upgrade procedure, take the appropriate recommended action.

If you encounter this firmware download error...	Then do the following...
<ul style="list-style-type: none">Failed assigned drives	<p>One reason for the failure might be that the drive does not have the appropriate signature. Make sure that the affected drive is an authorized drive. Contact technical support for more information.</p> <p>When replacing a drive, make sure that the replacement drive has a capacity equal to or greater than the failed drive you are replacing.</p> <p>You can replace the failed drive while the storage array is receiving I/O.</p>
<ul style="list-style-type: none">Check storage array	<ul style="list-style-type: none">Make sure that an IP address has been assigned to each controller.Make sure that all cables connected to the controller are not damaged.Make sure that all cables are tightly connected.
<ul style="list-style-type: none">Integrated hot spare drives	<p>This error condition must be corrected before you can upgrade the firmware. Launch System Manager and use the Recovery Guru to resolve the problem.</p>
<ul style="list-style-type: none">Incomplete volume groups	<p>If one or more volume groups or disk pools are incomplete, you must correct this error condition before you can upgrade the firmware.</p> <p>Launch System Manager and use the Recovery Guru to resolve the problem.</p>
<ul style="list-style-type: none">Exclusive operations \ (other than background media/parity scan) currently running on any volume groups	<p>If one or more exclusive operations are in progress, the operations must complete before the firmware can be upgraded. Use System Manager to monitor the progress of the operations.</p>
<ul style="list-style-type: none">Missing volumes	<p>You must correct the missing volume condition before the firmware can be upgraded. Launch System Manager and use the Recovery Guru to resolve the problem.</p>
<ul style="list-style-type: none">Either controller in a state other than Optimal	<p>One of the storage array controllers needs attention. This condition must be corrected before the firmware can be upgraded. Launch System Manager and use the Recovery Guru to resolve the problem.</p>

If you encounter this firmware download error...	Then do the following...
<ul style="list-style-type: none"> • Mismatched Storage Partition information between Controller Object Graphs 	An error occurred while validating the data on the controllers. Contact technical support to resolve this issue.
<ul style="list-style-type: none"> • SPM Verify Database Controller check fails 	A storage partitions mapping database error occurred on a controller. Contact technical support to resolve this issue.
<ul style="list-style-type: none"> • Configuration Database Validation (If supported by the storage array's controller version) 	A configuration database error occurred on a controller. Contact technical support to resolve this issue.
<ul style="list-style-type: none"> • MEL Related Checks 	Contact technical support to resolve this issue.
<ul style="list-style-type: none"> • More than 10 DDE Informational or Critical MEL events were reported in the last 7 days 	Contact technical support to resolve this issue.
<ul style="list-style-type: none"> • More than 2 Page 2C Critical MEL Events were reported in the last 7 days 	Contact technical support to resolve this issue.
<ul style="list-style-type: none"> • More than 2 Degraded Drive Channel Critical MEL events were reported in the last 7 days 	Contact technical support to resolve this issue.
<ul style="list-style-type: none"> • More than 4 critical MEL entries in the last 7 days 	Contact technical support to resolve this issue.

After you finish

Your drive firmware upgrade is complete. You can resume normal operations.

Review the possible software and firmware upgrade errors

Errors can occur during either the controller software upgrade or the drive firmware upgrade.

Firmware download error	Description	Recommended action
Failed assigned drives	Failed to upgrade an assigned drive in the storage array.	<p>One reason for the failure might be that the drive does not have the appropriate signature. Make sure that the affected drive is an authorized drive. Contact technical support for more information.</p> <p>When replacing a drive, make sure that the replacement drive has a capacity equal to or greater than the failed drive you are replacing.</p> <p>You can replace the failed drive while the storage array is receiving I/O.</p>
Integrated hot spare drives	If the drive is marked as a hot spare and is in use for a volume group, the firmware upgrade process fails.	This error condition must be corrected before you can upgrade the firmware. Launch SANtricity System Manager and use the Recovery Guru to resolve the problem.
Incomplete volume groups	If any drive that is part of a volume group is bypassed, removed or unresponsive, it is considered an incomplete volume group. An incomplete volume group prevents firmware upgrades.	If one or more volume groups or disk pools are incomplete, you must correct this error condition before you can upgrade the firmware. Launch SANtricity System Manager and use the Recovery Guru to resolve the problem.
Exclusive operations (other than background media/parity scan) currently running on any volume groups	Cannot upgrade the firmware if any exclusive operations are in progress on a volume.	If one or more exclusive operations are in progress, the operations must complete before the firmware can be upgraded. Use SANtricity System Manager to monitor the progress of the operations.
Missing volumes	Cannot upgrade the firmware if any volume is missing.	You must correct the missing volume condition before the firmware can be upgraded. Launch SANtricity System Manager and use the Recovery Guru to resolve the problem.

Firmware download error	Description	Recommended action
Either controller in a state other than Optimal	Cannot upgrade the firmware if either controller is in a state other than optimal.	One of the storage array controllers needs attention. This condition must be corrected before the firmware can be upgraded. Launch SANtricity System Manager and use the Recovery Guru to resolve the problem.
SPM Verify Database Controller check fails	Cannot upgrade the firmware because the storage partitions mappings database is corrupted.	A storage partitions mapping database error occurred on a controller. Contact technical support to resolve this issue.
Configuration Database Validation (If supported by the storage array's controller version)	Cannot upgrade the firmware because the configuration database is corrupted.	A configuration database error occurred on a controller. Contact technical support to resolve this issue.
MEL Related Checks	Cannot upgrade the firmware because the event log contains errors.	Contact technical support to resolve this issue.
More than 10 DDE Informational or Critical MEL events were reported in the last 7 days	Cannot upgrade the firmware because there are more than 10 DDE informational or critical MEL events reported in the last seven days.	Contact technical support to resolve this issue.
More than 2 Page 2C Critical MEL Events were reported in the last 7 days	Cannot upgrade the firmware because there are more than two page 2C critical MEL Events reported in the last seven days.	Contact technical support to resolve this issue.
More than 2 Degraded Drive Channel Critical MEL events were reported in the last 7 days	Cannot upgrade the firmware because there are more than two degraded drive channel critical MEL events reported in the last seven days.	Contact technical support to resolve this issue.
More than 4 critical MEL entries in the last 7 days	Cannot upgrade the firmware because there are more than four critical event log entries reported in the last seven days.	Contact technical support to resolve this issue.
A valid management IP address is required.	A valid controller IP address is required to perform this operation.	Contact technical support to resolve this issue.

Firmware download error	Description	Recommended action
The command requires an active management IP address for each controllers to be provided.	A controller IP address for each controller associated with the storage array is required for this operation.	Contact technical support to resolve this issue.
Unhandled download file type returned.	The specified download file is not supported.	Contact technical support to resolve this issue.
An error occurred during the firmware download upload procedure.	The firmware download failed because the controller cannot process the request. Verify the storage array is optimal and retry the operation.	If this error occurs again after verifying the storage array is optimal, contact technical support to resolve this issue.
An error occurred during the firmware activation procedure.	The firmware activation failed because the controller cannot process the request. Verify the storage array is optimal and retry the operation.	If this error occurs again after verifying the storage array is optimal, contact technical support to resolve this issue.
Timeout has reached while waiting for controller {0} to reboot.	The management software is unable to reconnect to controller {0} following a reboot. Validate there is an operational connection path to the storage array and retry the operation if it did not complete successfully.	If this error occurs again after verifying the storage array is optimal, contact technical support to resolve this issue.

You can correct some of these conditions by using the Recovery Guru in SANtricity System Manager. However, for some of the conditions, you might need to contact technical support. The information about the latest controller firmware download is available from the storage array. This information helps technical support to understand the error conditions that prevented the firmware upgrade and download.

FAQs

What do I need to know before upgrading the SANtricity OS Software?

Before you upgrade your controller's software and firmware, be aware of these items.

- You have read the document and the `readme.txt` file and have determined that you want to do the upgrade.
- You know whether you want to upgrade your IOM firmware.

Normally, you should upgrade all components at the same time. However, you might decide not to upgrade the IOM firmware if you do not want to upgrade it as part of the SANtricity OS controller software upgrade or if technical support has instructed you to downgrade your IOM firmware (you can only downgrade firmware by using the command line interface).

- You know whether you want to upgrade the controller NVSRAM file.

Normally, you should upgrade all components at the same time. However, you might decide not to upgrade the controller NVSRAM file if your file has either been patched or is a custom version and you do not want to overwrite it.

- You know whether you want to activate now or later.

Reasons for activating later might include:

- **Time of day** — Activating the software and firmware can take a long time, so you might want to wait until I/O loads are lighter. The controllers fail over during activation so performance might be lower than usual until the upgrade completes.
- **Type of package** — You might want to test the new software and firmware on one storage array before upgrading the files on other storage arrays.

These components are part of the SANtricity OS controller software upgrade:

- **Management software** — System Manager is the software that manages the storage array.
- **Controller firmware** — Controller firmware manages the I/O between hosts and volumes.
- **Controller NVSRAM** — Controller NVSRAM is a controller file that specifies the default settings for the controllers.
- **IOM firmware** — The I/O module (IOM) firmware manages the connection between a controller and a drive shelf. It also monitors the status of the components.
- **Supervisor software** — Supervisor software is the virtual machine on a controller in which the software runs.

As part of the upgrade process, the host's multipath/failover driver and/or HBA driver might also need to be upgraded so the host can interact with the controllers correctly.



To determine if this is the case, see the [NetApp Interoperability Matrix Tool](#).

If your storage array contains only one controller or you do not have a multipath driver installed, stop I/O activity to the storage array to prevent application errors. If your storage array has two controllers and you have a multipath driver installed, you do not need to stop I/O activity.

Do not make changes to the storage array while the upgrade occurs.

What do I need to know before suspending IOM auto-synchronization?

Suspending IOM auto-synchronization prevents the IOM firmware from being upgraded the next time a SANtricity OS controller software upgrade occurs.

Normally, controller software and IOM firmware are upgraded as a bundle. You can suspend IOM auto-synchronization if you have a special build of IOM firmware that you want to preserve on your enclosure. Otherwise, you will revert to the IOM firmware bundled with the controller software the next time you do a controller software upgrade.

Why is my firmware upgrade progressing so slowly?

The firmware upgrade progress depends on the overall load of the system.

During an online upgrade of drive firmware, if a volume transfer takes place during the rapid reconstruction process, the system initiates a full reconstruction on the volume that was transferred. This operation might take

a considerable amount of time. Actual full reconstruction time depends on several factors, including the amount of I/O activity occurring during the reconstruction operation, the number of drives in the volume group, the rebuild priority setting, and the drive performance.

What do I need to know before upgrading drive firmware?

Before upgrading your drive firmware, be aware of these items.

- As a precaution, back up your data using disk-to-disk backup, volume copy (to a volume group not affected by the planned firmware upgrade), or a remote mirror.
- You might want to upgrade only a few drives to test the behavior of the new firmware to ensure it is functioning correctly. If the new firmware is functioning correctly, upgrade the remaining drives.
- If you have any failed drives, fix them before you start the firmware upgrade.
- If the drives can do an offline upgrade, stop I/O activity to all volumes associated with the drives. When I/O activity is stopped, no configuration operations associated with those volumes can occur.
- Do not remove any drives while upgrading drive firmware.
- Do not make any configuration changes to the storage array while upgrading drive firmware.

How do I choose which type of upgrade to perform?

You choose the type of upgrade to perform on the drive depending on the state of the pool or volume group.

- **Online**

If the pool or volume group supports redundancy and is Optimal, you can use the Online method to upgrade the drive firmware. The Online method downloads firmware *while the storage array is processing I/O* to the associated volumes using these drives. You do not have to stop I/O to the associated volumes using these drives. These drives are upgraded one at a time to the volumes associated with the drives. If the drive is not assigned to a pool or volume group, its firmware can be updated by the Online or the Offline method. System performance may be impacted when you use the Online method to upgrade drive firmware.

- **Offline**

If the pool or volume group does not support redundancy (RAID 0), or is degraded, you must use the Offline method to upgrade the drive firmware. The Offline method will upgrade firmware *only while all I/O activity is stopped* to the associated volumes using these drives. You must stop all I/O to any associated volumes using these drives. If the drive is not assigned to a pool or volume group, its firmware may be updated by the Online or the Offline method.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.