



Hardware components

SANtricity 11.8

NetApp
April 05, 2024

Table of Contents

- Hardware components 1
 - Hardware component overview 1
 - Related information 1
 - Concepts 1
 - Manage shelf components 12
 - Manage controllers 17
 - Manage iSCSI ports 27
 - Manage NVMe ports 37
 - Manage drives 41
 - Manage hot spares 53
 - Shelf FAQs 56
 - Controller FAQs 59
 - iSCSI FAQs 62
 - NVMe FAQs 64
 - Drive FAQs 66

Hardware components

Hardware component overview

You can check component status on the Hardware page and perform some functions related to those components.

Which components can I manage?

You can check component status and perform some functions related to these components:

- **Shelves** — A *shelf* is a component that contains the hardware for the storage array (controllers, power/fan canisters, and drives). Shelves are available in three sizes for housing up to 12, 24, or 60 drives.
- **Controllers** — A *controller* is the combined hardware and firmware that implements storage array and management functions. It includes the cache memory, drive support, and the ports for host connections.
- **Drives** — A *drive* can be either a hard disk drive (HDD) or a solid state drive (SSD). Depending on the shelf size, up to 12, 24, or 60 drives can be installed in the shelf.

Learn more:

- [Hardware page](#)
- [Hardware terminology](#)

How do I view hardware components?

Go to the Hardware page, which provides a graphical depiction of the storage array's physical components. You can switch between the front and back views of the array shelves by selecting either **Show back of shelf** or **Show front of shelf** from the upper right of the shelf view.

Learn more:

- [View shelf component status and settings](#)
- [View controller settings](#)
- [View drive status and settings](#)

Related information

Learn more about concepts related to hardware:

- [Controller states](#)
- [Drive states](#)
- [Shelf loss protection and drawer loss protection](#)

Concepts

Hardware page and components

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 drives.

Hardware terminology

The following hardware terms apply to storage arrays.

General hardware terms:

Component	Description
Bay	A bay is a slot in the shelf where a drive or other component is installed.
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.
Drive	A drive is an electromagnetic mechanical device or solid state memory device that provides the physical storage media for data.
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.
IOM (ESM)	An IOM is an input/output module that includes SAS ports for connecting the drive shelf to the controller shelf. In previous controller models, the IOM was referred to as an environmental service module (ESM).
Power/fan canister	A power/fan canister is an assembly that slides into a shelf. It includes a power supply and an integrated fan.
SFP	An SFP is a Small Form-factor Pluggable (SFP) transceiver.
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.
Storage array	A storage array includes the shelves, controllers, drives, software, and firmware.

Controller terms:

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.
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.

Component	Description
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.).

Drive terms:

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 the EF300 or EF600 storage array to support resource-provisioned volumes.
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.
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.

Component	Description
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.

iSCSI terms:

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 CHAP <i>secret</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.
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.

Term	Description
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.

NVMe terms:

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.


Manage shelf components

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
Drives, controllers, and components views	To switch between front and back shelf views, select either Drives or Controllers & Components from the far right (the link that appears depends on the current view). The Drives view shows drives and any empty drive bays. The Controllers & Components 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 drives .
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 data-bbox="506 1234 561 1289"></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>
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.

Function	Description
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. Optionally, 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.


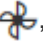



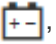


- **Controller shelf** — Contains a set of drives, one or two controller canisters, power/fan canisters, and other supporting components in a single shelf.

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> <p> You can also view this status by selecting the IOM icon , next to the Shelf drop-down list.</p>
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> <p> You can also view this status by selecting the Power Supply icon , next to the Shelf drop-down list.</p>

Tab	Description
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> <p> You can also view this status by selecting the Fan icon , next to the Shelf drop-down list.</p>
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**.

Manage controllers

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.
- Cache mirroring is disabled and all volumes are changed to write through cache mode.

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 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 3 minutes. The default IP addresses are as follows:

```
Controller 1 (port 1): IP Address: 192.168.128.101
```

```
Controller 1 (port 2): IP Address: 192.168.129.101
```

```
Controller 2 (port 1): IP Address: 192.168.128.102
```

```
Controller 2 (port 2): IP Address: 192.168.129.102
```

When assigning IP addresses:

- Reserve Port 2 on the controllers for Customer 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. 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.

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	<p>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.</p> <p> 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.</p>

8. Click **Finish**.

Results

The management port configuration is displayed in the controller settings, Management Ports tab.

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 (NVS RAM) 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	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. <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;">  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. </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	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. Network Time Protocol (NTP) is a networking protocol for clock synchronization between computer systems in data networks.

5. Click **Close**.

Configure remote login (SSH)

By enabling remote login, you allow users from outside the local area network to start an SSH session and access settings on the controller.

For SANtricity versions 11.74 and later, you can also configure multifactor authorization (MFA) by requiring users to enter an SSH key and/or SSH password. For SANtricity versions 11.73 and earlier, this feature does *not* include an option for multifactor authorization with SSH keys and passwords.



Security risk — For security reasons, only technical support personnel should use the Remote Login feature.

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 configure remote login.

The controller's context menu appears.

4. Select **Configure remote login (SSH)**. (For SANtricity versions 11.73 and earlier, this menu item is **Change remote login**.)

The dialog box opens for enabling remote login.

5. Select the **Enable remote login** checkbox.

This setting provides remote login with three options for authorization:

- **Password only**. For this option, you are done and can click **Save**. If you have a duplex system, you can enable remote login on the second controller by following the previous steps.
 - **Either SSH key or password**. For this option, proceed to the next step.
 - **Both password and SSH key**. For this option, select the **Require authorized public key and password for remote login** checkbox and proceed to the next step.
6. Populate the **Authorized public key** field. This field contains a list of authorized public keys, in the format of the OpenSSH **authorized_keys** file.

When populating the **Authorized public key** field, be aware of the following guidelines:

- The **Authorized public key** field applies to both controllers and only needs to be configured on the first controller.
- The **authorized_keys** file should contain only one key per line. Lines starting with # and empty lines are ignored. For more information about the file format, see [Configuring Authorized Keys for OpenSSH](#).
- An **authorized_keys** file should look similar to the following example:

```
ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQBAQDj1G20rYTk4ok+xFjkPHYp/R0LfJqEYDLXA5AJ4
9w3DvAWLrUg+1CpNq76WSqmQBmoG9jgbcAB5ABGdswdeMQZHilJcu29iJ3OKKv6S1CulA
j1tHymwtbdhPuipd2wIDAQAB
```

7. When you're done, click **Save**.
8. For duplex systems, you can enable remote login on the second controller by following the steps above. If you are configuring the option for both a password and SSH key, be sure to select the **Require authorized public key and password for remote login** checkbox again.
9. After technical support is finished troubleshooting, you can disable remote login by returning to the Configure Remote Login dialog box and de-selecting the **Enable remote login** checkbox. If remote login is enabled on a second controller, a confirmation dialog opens and allows you to disable remote login on the second one as well.

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.



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.

Manage iSCSI ports

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.

- In the drop-down list, select the port you want to configure, and then click **Next**.
- 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
Configured ethernet port speed (Appears only for certain types of Host Interface Cards)	Select the speed that matches the speed capability of the SFP on the port.
Forward Error Correction (FEC) mode (Appears only for certain types of Host Interface Cards)	<p>If desired, select one of the FEC modes for the specified host port.</p> <p> The Reed Solomon mode does not support the 25 Gbps port speed.</p>
Enable IPv4 / Enable IPv6	<p>Select one or both options to enable support for IPv4 and IPv6 networks.</p> <p> If you want to disable port access, deselect both check boxes.</p>
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	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.

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.

- 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 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 **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 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. **Optional:** 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.
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 **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**.

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.

Manage NVMe ports

NVMe overview

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.

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	Select one or both options to enable support for IPv4 and IPv6 networks.  If you want to disable port access, deselect both check boxes.
MTU size (Available by clicking Show more port settings.)	If necessary, enter a new size in bytes for the Maximum Transmission Unit (MTU). The default Maximum Transmission Unit (MTU) size is 1500 bytes per frame. You must enter a value between 1500 and 9000.

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	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.

8. 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.

Manage drives

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.

State	Definition
Present	The controller can communicate with the drive on both ports.
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).

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
 - Drive capacity (highest to lowest)
 - Drive speed (highest to lowest)
- Security attributes might include:

- Secure-capable
- Secure-enabled
- DA (Data Assurance) capable
- FIPS compliant
- FIPS compliant (FIPS 140-2)
- FIPS compliant (FIPS 140-3)

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. Optionally, 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.  Not available for NVMe drives.
Secure-capable	Shows whether the drive is secure-capable (Yes) or not (No). Secure-capable drives can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives (level 140-2 or 140-3), 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.
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.
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.

Settings	Description
DULBE capable	Indicates whether the option for Deallocated or Unwritten Logical Block Error (DULBE) is enabled (Yes) or not (No). DULBE is an option on NVMe drives that allows the EF300 or EF600 storage array to support resource-provisioned volumes.

6. Click **Close**.

Replace drive logically

If a drive fails or you want to replace it for any other reason, you can logically replace the failed drive with an unassigned drive or a fully integrated hot spare.

About this task

When you logically replace a drive, it 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)

Before you begin

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.



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.

Erase drives

You can use the Erase option to prepare an unassigned drive for removal from the system. This procedure permanently removes data, ensuring that the data cannot be read again.

Before you begin

The drive must be in an Unassigned state.

About this task

Use the Erase option only if you want to permanently remove all data on a drive. If the drive is secure-enabled, the Erase option performs a cryptographic erase and resets the drive's security attributes back to secure-capable.



The Erase feature does not support some older drive models. If you attempt to erase one of these older models, an error message appears.

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. Optionally, you can use the filter fields to view all the unassigned drives in the shelf. From the **Show drives that are...** drop-down list, select **Unassigned**.

The shelf view shows only the unassigned drives; all others are grayed out.

4. To open the drive's context menu, click on a drive that you want to erase. (If you want to select multiple drives, you can do so in the Erase Drives dialog box.)



Possible loss of data — The Erase operation cannot be undone. Make sure you select the correct drives during the procedure.

5. From the context menu, select **Erase**.

The Erase Drives dialog box opens, showing all the eligible drives for an erase operation.

6. If desired, select additional drives from the table. You cannot select *all* drives; be sure one drive remains deselected.
7. Confirm the operation by typing `erase`, and then click **Erase**.



Be sure you want to continue with this operation. Once you click Yes in the next dialog, the operation cannot be aborted.

8. In the Estimated Completion Time dialog box, click **Yes** to continue with the erase operation.

Results

The Erase operation might take several minutes or several hours. You can view the status in **Home > View Operations in Progress**. When the Erase operation completes, the drives are available for use in another volume group or disk pool, or in another storage array.

After you finish

If you want to use the drive again, you must initialize it first. To do this, select **Initialize** from the drive's context menu.

Unlock or reset locked NVMe or FIPS drives

If you insert one or more locked NVMe or FIPS 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 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 or FIPS 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.

You can access the Unlock feature from either the Hardware page or from **Settings > System > Security key management**. The task below provides instructions from 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 the NVMe or FIPS 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 or FIPS 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 or FIPS drives. For a Reset operation, you must individually select each drive you want to reset.

Results

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

Manage hot spares

Hot spare drive overview

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.

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).

Shelf 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	<p>Pool candidates must include drives from all drawers, and there must be an equal number of drives in each drawer.</p> <p>The pool must include drives from at least five drawers and there must be an equal number of drives in each drawer.</p> <p>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.</p>	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.

Controller 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, subnet 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 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.

Why is the storage system in non-optimal mode?

A storage system in non-optimal mode is due to an Invalid System Configuration state. Despite this state, normal I/O access to existing volumes is fully supported; however, System Manager will prohibit some operations.

A storage system might transition to Invalid System Configuration for one of these reasons:

- The controller is out of compliance, possibly because it has an incorrect submodel ID (SMID) code or it has exceeded the limit of premium features.

- An internal service operation is in progress, such as a drive firmware download.
- The controller exceeded the parity error threshold and went into lockdown.
- A general lockdown condition occurred.

iSCSI 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.

Action	Location
Configure iSER 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 iSER over InfiniBand ports. <p>or</p> <ol style="list-style-type: none"> 1. Select Settings > System. 2. Scroll down to iSER over InfiniBand settings, and then select Configure iSER over InfiniBand Ports.
View iSER over InfiniBand statistics	<ol style="list-style-type: none"> 1. Select Settings > System. 2. Scroll down to iSER over InfiniBand settings, and then select View iSER over InfiniBand Statistics.

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.

Action	Location
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.
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.

NVMe 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.

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.

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.
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.

Drive 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 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, the system 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.

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.