



Settings

SANtricity 11.5

NetApp
February 12, 2024

Table of Contents

- Settings 1
 - Alerts 1
 - System 14
 - Access management 66
 - Certificates 98

Settings

Alerts

Concepts

How alerts work

Alerts notify administrators about important events that occur on the storage array. Alerts can be sent through email, SNMP traps, and syslog.

The alerts process works as follows:

1. An administrator configures one or more of the following alerting methods in System Manager:
 - **Email** — Messages are sent to email addresses.
 - **SNMP** — SNMP traps are sent to an SNMP server.
 - **Syslog** — Messages are sent to a syslog server.
2. When the storage array's event monitor detects an issue, it writes information about that issue to the event log (available from **Support > Event Log**). For example, issues can include such events as a battery failure, a component moving from Optimal to Offline, or redundancy errors in the controller.
3. If the event monitor determines that the event is "alertable," it then sends a notification using the configured alerting methods (email, SNMP, and/or syslog). All Critical events are considered "alertable," along with some Warning and Informational events.

Alerts configuration

You can configure alerts from the Initial Setup wizard (for email alerts only) or from the Alerts page. To check the current configuration, go to **Settings > Alerts**.

The Alerts tile displays the alerts configuration, which can be one of the following:

- Not configured.
- Configured; at least one alerting method is set up. To determine which alerting methods are configured, point the cursor at the tile.

Alerts information

Alerts can include the following types of information:

- Name of the storage array.
- Event error type related to an event log entry.
- Date and time when the event occurred.
- Brief description of the event.



Syslog alerts follow the RFC 3164 messaging standard.

Alerts terminology

Learn how the alerts terms apply to your storage array.

Component	Description
Event monitor	The event monitor resides on the storage array and runs as a background task. When the event monitor detects anomalies on the storage array, it writes information about the issues to the event log. Issues can include such events as a battery failure, a component moving from Optimal to Offline, or redundancy errors in the controller. If the event monitor determines that the event is "alertable," it then sends a notification using the configured alerting methods (email, SNMP, and/or syslog). All Critical events are considered "alertable," along with some Warning and Informational events.
Mail server	The mail server is used for sending and receiving email alerts. The server uses Simple Mail Transfer Protocol (SMTP).
SNMP	Simple Network Management Protocol (SNMP) is an Internet-standard protocol used for managing and sharing information between devices on IP networks.
SNMP trap	An SNMP trap is a notification sent to an SNMP server. The trap contains information about significant issues with the storage array.
SNMP trap destination	An SNMP trap destination is an IPv4 or IPv6 address of the server running an SNMP service.
Community name	A community name is a string that acts like a password for the network server(s) in a SNMP environment.
MIB file	The management information base (MIB) file defines the data being monitored and managed in the storage array. It must be copied and compiled on the server with the SNMP service application. This MIB file is available with the System Manager software on the Support site.
MIB variables	Management Information Base (MIB) variables can return values such as the storage array name, array location, and a contact person in response to SNMP GetRequests.

Component	Description
Syslog	Syslog is a protocol used by network devices for sending event messages to a logging server.
UDP	User Datagram Protocol (UDP) is a transport layer protocol that specifies a source and destination port number in their packet headers.

How to

Manage email alerts

Configure mail server and recipients for alerts

To configure email alerts, you must specify a mail server address and the email addresses of the alert recipients. Up to 20 email addresses are allowed.

Before you begin

- The address of the mail server must be available. The address can be an IPv4 or IPv6 address, or a fully qualified domain name.



To use a fully qualified domain name, you must configure a DNS server on both controllers. You can configure a DNS server from the Hardware page.

- Email address to be used as the alert sender must be available. This is the address that appears in the "From" field of the alert message. A sender address is required in the SMTP protocol; without it, an error results.
- Email address(es) of the alert recipient(s) must be available. The recipient is typically an address for a network administrator or storage administrator. You can enter up to 20 email addresses.

About this task

This task describes how to configure the mail server, enter email addresses for the sender and recipients, and test all the email addresses entered from the Alerts page.



Email alerts can also be configured from the Initial Setup wizard.

Steps

1. Select **Settings > Alerts**.
2. Select the **Email** tab.

If an email server is not yet configured, the **Email** tab displays "Configure Mail Server."

3. Select **Configure Mail Server**.

The **Configure Mail Server** dialog box opens.

4. Enter the mail server information, and then click **Save**.

- Mail server address — Enter a fully qualified domain name, IPv4 address, or IPv6 address of the mail

server.



To use a fully qualified domain name, you must configure a DNS server on both controllers. You can configure a DNS server from the Hardware page.

- Email sender address — Enter a valid email address to be used as the sender of the email. This address appears in the "From" field of the email message.
- Include contact information in email — To include the sender's contact information with the alert message, select this option, and then enter a name and phone number. After you click **Save**, the email addresses appear in the **Email** tab of the **Alerts** page.

5. Select **Add Emails**.

The **Add Emails** dialog box opens.

6. Enter one or more email addresses for the alert recipients, and then click **Add**.

The email addresses appear on the **Alerts** page.

7. If you want to make sure the email addresses are valid, click **Test All Emails** to send test messages to the recipients.

Result

After you configure email alerts, the event monitor sends email messages to the specified recipients whenever an alertable event occurs.

Edit email addresses for alerts

You can change the email addresses of the recipients who receive email alerts.

Before you begin

The email address you intend to edit must be defined in the Email tab of the Alerts page.

Steps

1. Select **Settings > Alerts**.
2. Select the **Email** tab.
3. From the **Email Address** table, select the address you want to change, and then click the **Edit** (pencil) icon on the far right.

The row becomes an editable field.

4. Enter a new address, and then click the **Save** (checkmark) icon.



If you want to cancel changes, select the Cancel (X) icon.

Result

The Email tab of the Alerts page displays the updated email addresses.

Add email addresses for alerts

You can add up to 20 recipients for email alerts.

Steps

1. Select **Settings > Alerts**.
2. Select the **Email** tab.
3. Select **Add Emails**.

The Add Emails dialog box opens.

4. In the empty field, enter a new email address. If you want to add more than one address, select **Add another email** to open another field.
5. Click **Add**.

Result

The Email tab of the Alerts page displays the new email addresses.

Delete email addresses for alerts

You can delete the email addresses of the recipients who receive email alerts.

Steps

1. Select **Settings > Alerts**.
2. Select the **Email** tab.
3. From the **Email Address** table, select the email address you want to delete.

The **Delete** button in the upper right of the table becomes available for selection.

4. Click **Delete**.

The **Confirm Delete Email** dialog box opens.

5. Confirm the operation, and then click **Delete**.

Result

Alerts are no longer sent to this email address.

Edit mail server for alerts

You can change the mail server address and email sender address used for email alerts.

Before you begin

The address of the mail server you are changing must be available. The address can be an IPv4 or IPv6 address, or a fully qualified domain name.



To use a fully qualified domain name, you must configure a DNS server on both controllers. You can configure a DNS server from the Hardware page.

Steps

1. Select **Settings > Alerts**.
2. Select the **Email** tab.
3. Select **Configure Mail Server**.

The **Configure Mail Server** dialog opens.

4. Edit the mail server address, sender information, and contact information.

- Mail server address — Edit the fully qualified domain name, IPv4 address, or IPv6 address of the mail server.



To use a fully qualified domain name, you must configure a DNS server on both controllers. You can configure a DNS server from the Hardware page.

- Email sender address — Edit the email address to be used as the sender of the email. This address appears in the "From" field of the email message.
- Include contact information in email — To edit the sender's contact information, select this option, and then edit the name and phone number.

5. Click **Save**.

Manage SNMP alerts

Configure communities and destinations for SNMP alerts

To configure Simple Network Management Protocol (SNMP) alerts, you must identify at least one server where the storage array's event monitor can send SNMP traps. The configuration requires a community name and IP address for the server.

Before you begin

- A network server must be configured with an SNMP service application. You need the network address of this server (either an IPv4 or an IPv6 address), so the event monitor can send trap messages to that address. You can use more than one server (up to 10 servers are allowed).
- A community name must be created, consisting of only printable ASCII characters. The community name, which is a string that acts like a password for the network servers, is typically created by a network administrator. Up to 256 communities can be created.
- The management information base (MIB) file has been copied and compiled on the server with the SNMP service application. This MIB file defines the data being monitored and managed.

If you do not have the MIB file, you can obtain it from the NetApp Support site:

- Go to [NetApp Support](#).
- Click **Downloads**.
- Click **Software**.
- Find your management software (for example, SANtricity System Manager), and then click **Go!** on the right.
- Click **View & Download** on the latest version.
- Click **Continue** at the bottom of the page.
- Accept the EULA.
- Scroll down until you see **MIB file for SNMP traps**, and then click the link to download the file.

About this task

This task describes how to identify the SNMP server for trap destinations, and then test your configuration.

Steps

1. Select **Settings > Alerts**.
2. Select the **SNMP** tab.

If a community is not yet configured, the SNMP tab displays "Configure Communities."

3. Select **Configure Communities**.

The **Configure Communities** dialog box opens.

4. In the **Community Name** field, enter one or more community strings for the network servers, and then click **Save**.

The **Alerts** page displays "Add Trap Destinations."

5. Select **Add Trap Destinations**.

The **Add Trap Destinations** dialog box opens.

6. Enter one or more trap destinations, select their associated community names, and then click **Add**.
 - Trap Destination — Enter an IPv4 or IPv6 address of the server running an SNMP service.
 - Community name — From the drop-down, select the community name for this trap destination. (If you defined only one community name, the name already appears in this field.)
 - Send Authentication Failure Trap — Select this option (the checkbox) if you want to alert the trap destination whenever an SNMP request is rejected because of an unrecognized community name. After you click Add, the trap destinations and associated community names appear in the **SNMP** tab of the **Alerts** page.
7. To make sure a trap is valid, select a trap destination from the table, and then click **Test Trap Destination** to send a test trap to the configured address.

Result

The event monitor sends SNMP traps to the server(s) whenever an alertable event occurs.

Edit community names for SNMP traps

You can edit community names for SNMP traps, and also associate a different community name to an SNMP trap destination.

Before you begin

A community name must be created, consisting of only printable ASCII characters. The community name, which is a string that acts like a password for the network servers, is created by a network administrator.

Steps

1. Select **Settings > Alerts**.
2. Select the **SNMP** tab.

The trap destinations and community names appear in the table.

3. Edit community names as follows:
 - To edit a community name, select **Configure Communities**. Enter the new community name, and then

click **Save**. Community names can consist of only printable ASCII characters.

- To associate a community name to a new trap destination, select the community name from the table, and then click the **Edit** (pencil) icon on the far right. From the **Community Name** drop-down, select a new community name for an SNMP trap destination, and then click the Save (checkmark) icon.



If you want to cancel changes, select the Cancel (X) icon.

Result

The **SNMP** tab of the **Alerts** page displays the updated communities.

Add community names for SNMP traps

You can add up to 256 community names for SNMP traps.

Before you begin

The community name(s) must be created. The community name, which is a string that acts like a password for the network servers, is typically created by a network administrator. It consists of only printable ASCII characters.

Steps

1. Select **Settings > Alerts**.
2. Select the **SNMP** tab.

The trap destinations and community names appear in the table.

3. Select **Configure Communities**.

The **Configure Communities** dialog box opens.

4. Select **Add another community**.
5. Enter the new community name, and then click **Save**.

Result

The new community name appears in the **SNMP** tab of the **Alerts** page.

Remove community name for SNMP traps

You can remove a community name for SNMP traps.

Steps

1. Select **Settings > Alerts**.
2. Select the **SNMP** tab.

The trap destinations and community names appear on the Alerts page.

3. Select **Configure Communities**.

The **Configure Communities** dialog box opens.

4. Select the community name you want to delete, and then click the **Remove** (X) icon on the far right.

If trap destinations are associated with this community name, the **Confirm Remove Community** dialog box shows the affected trap destination addresses.

5. Confirm the operation, and then click **Remove**.

Results

The community name and its associated trap destination are removed from the Alerts page.

Configure SNMP MIB variables

For SNMP alerts, you can optionally configure Management Information Base (MIB) variables that appear in SNMP traps. These variables can return the storage array name, array location, and a contact person.

Before you begin

The MIB file must be copied and compiled on the server with the SNMP service application.

If you do not have a MIB file, you can obtain it as follows:

- Go to [NetApp Support](#).
- Click **Downloads**.
- Click **Software**.
- Find your management software (for example, SANtricity System Manager), and then click **Go!** on the right.
- Click **View & Download** on the latest version.
- Click **Continue** at the bottom of the page.
- Accept the EULA.
- Scroll down until you see **MIB file for SNMP traps**, and then click the link to download the file.

About this task

This task describes how to define MIB variables for SNMP traps. These variables can return the following values in response to SNMP GetRequests:

- *sysName* (name for the storage array)
- *sysLocation* (location of the storage array)
- *sysContact* (name of an administrator)

Steps

1. Select **Settings > Alerts**.
2. Select the **SNMP** tab.
3. Select **Configure SNMP MIB Variables**.

The **Configure SNMP MIB Variables** dialog box opens.

4. Enter one or more of the following values, and then click **Save**.
 - **Name** — The value for the MIB variable *sysName*. For example, enter a name for the storage array.
 - **Location** — The value for the MIB variable *sysLocation*. For example, enter a location of the storage

array.

- **Contact** — The value for the MIB variable *sysContact*. For example, enter an administrator responsible for the storage array.

Result

These values appear in SNMP trap messages for storage array alerts.

Add trap destinations for SNMP alerts

You can add up to 10 servers for sending SNMP traps.

Before you begin

- The network server you want to add must be configured with an SNMP service application. You need the network address of this server (either an IPv4 or an IPv6 address), so the event monitor can send trap messages to that address. You can use more than one server (up to 10 servers are allowed).
- A community name must be created, consisting of only printable ASCII characters. The community name, which is a string that acts like a password for the network servers, is typically created by a network administrator. Up to 256 communities can be created.
- The management information base (MIB) file has been copied and compiled on the server with the SNMP service application. This MIB file defines the data being monitored and managed.

If you do not have the MIB file, you can obtain it from the NetApp Support site:

- Go to [NetApp Support](#).
- Click **Downloads**.
- Click **Software**.
- Find your management software (for example, SANtricity System Manager), and then click **Go!** on the right.
- Click **View & Download** on the latest version.
- Click **Continue** at the bottom of the page.
- Accept the EULA.
- Scroll down until you see **MIB file for SNMP traps**, and then click the link to download the file.

Steps

1. Select **Settings > Alerts**.
2. Select the **SNMP** tab.

The currently defined trap destinations appear in the table.

3. Select **Add Trap Destinations**.

The **Add Trap Destinations** dialog box opens.

4. Enter one or more trap destinations, select their associated community names, and then click **Add**.
 - **Trap Destination** — Enter an IPv4 or IPv6 address of the server running an SNMP service.
 - **Community name** — From the drop-down, select the community name for this trap destination. (If you defined only one community name, the name already appears in this field.)

- Send Authentication Failure Trap — Select this option (the checkbox) if you want to alert the trap destination whenever an SNMP request is rejected because of an unrecognized community name. After you click Add, the trap destinations and associated community names appear in the table.

5. To make sure a trap is valid, select a trap destination from the table, and then click **Test Trap Destination** to send a test trap to the configured address.

Result

The event monitor sends SNMP traps to the server(s) whenever an alertable event occurs.

Delete trap destinations

You can delete a trap destination address so that the storage array's event monitor no longer sends SNMP traps to that address.

Steps

1. Select **Settings > Alerts**.
2. Select the **SNMP** tab.

The trap destination addresses appear in the table.

3. Select a trap destination, and then click **Delete** in the upper right of the page.
4. Confirm the operation, and then click **Delete**.

The destination address no longer appears on the **Alerts** page.

Result

The deleted trap destination no longer receives SNMP traps from the storage array's event monitor.

Manage syslog alerts

Configure syslog server for alerts

To configure syslog alerts, you must enter a syslog server address and a UDP port. Up to five syslog servers are allowed.

Before you begin

- The syslog server address must be available. This address can be a fully qualified domain name, an IPv4 address, or an IPv6 address.
- UDP port number of the syslog server must be available. This port is typically 514.

About this task

This task describes how to enter the address and port for the syslog server, and then test the address you entered.

Steps

1. Select **Settings > Alerts**.
2. Select the **Syslog** tab.

If a syslog server is not yet defined, the **Alerts** page displays "Add Syslog Servers."

3. Click **Add Syslog Servers**.

The **Add Syslog Server** dialog box opens.

4. Enter information for one or more syslog servers (maximum of five), and then click **Add**.
- **Server Address** — Enter a fully qualified domain name, an IPv4 address, or an IPv6 address.
 - **UDP Port** — Typically, the UDP port for syslog is 514. The table displays the configured syslog servers.
5. To send a test alert to the server addresses, select **Test All Syslog Servers**.

Result

The event monitor sends alerts to the syslog server whenever an alertable event occurs.

Edit syslog servers for alerts

You can edit the server address used for receiving syslog alerts.

Steps

1. Select **Settings > Alerts**.
2. Select the **Syslog** tab.
3. From the table, select a syslog server address, and then click the **Edit** (pencil) icon from on the far right.

The row becomes an editable field.

4. Edit the server address and UDP port number, and then click the **Save** (checkmark) icon.

Result

The updated server address appears in the table.

Add syslog servers for alerts

You can add a maximum of five servers for syslog alerts.

Before you begin

- The syslog server address must be available. This address can be a fully qualified domain name, an IPv4 address, or an IPv6 address.
- The UDP port number of the syslog server must be available. This port is typically 514.

Steps

1. Select **Settings > Alerts**.
2. Select the **Syslog** tab.
3. Select **Add Syslog Servers**.

The **Add Syslog Server** dialog box opens.

4. Select **Add another syslog server**.
5. Enter information for the syslog server, and then click **Add**.
 - **Syslog Server Address** — Enter a fully qualified domain name, an IPv4 address, or an IPv6 address.

- UDP Port — Typically, the UDP port for syslog is 514.



You can configure up to five syslog servers.

Result

The syslog server addresses appear in the table.

Delete syslog servers for alerts

You can delete a syslog server so it no longer receives alerts.

Steps

1. Select **Settings** > **Alerts**.
2. Select the **Syslog** tab.
3. Select a syslog server address, and then click **Remove** from the top right.

The **Confirm Delete Syslog Server** dialog box opens.

4. Confirm the operation, and then click **Delete**.

Result

The server you removed no longer receives alerts from the event monitor.

FAQs

What if alerts are disabled?

If you want administrators to receive notifications about important events that occur in the storage array, you must configure an alerting method.

For storage arrays managed with SANtricity System Manager, you configure alerts from the Alerts page. Alert notifications can be sent through email, SNMP traps, or syslog messages. In addition, email alerts can be configured from the Initial Setup Wizard.

How do I configure SNMP or syslog alerts?

In addition to email alerts, you can configure alerts to be sent by Simple Network Management Protocol (SNMP) traps or by syslog messages.

To configure SNMP or syslog alerts, go to **Settings** > **Alerts**.

Why are timestamps inconsistent between the array and alerts?

When the storage array sends alerts, it does not correct for the time zone of the target server or host that receives the alerts. Instead, the storage array uses the local time (GMT) to create the timestamp used for the alert record. As a result, you might see inconsistencies between the timestamps for the storage array and the server or host receiving an alert.

Because the storage array does not correct for time zone when sending alerts, the timestamp on the alerts is

GMT-relative, which has a time-zone offset of zero. To calculate a timestamp appropriate to your local time zone, you should determine your hour offset from GMT, and then add or subtract that value from the timestamps.



To avoid this issue, configure Network Time Protocol (NTP) on your storage array controllers. NTP ensures that the controllers are always synced to the correct time.

System

Storage array settings

Concepts

Cache settings and performance

Cache memory is an area of temporary volatile storage on the controller that has a faster access time than the drive media.

With caching, overall I/O performance can be increased as follows:

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

The default cache settings meet the requirements for most environments, but you can change them if you want.

Storage array cache settings

For all volumes in the storage array, you can specify the following values from the System page:

- **Start value for flushing** — the percentage of unwritten data in the cache that triggers a cache flush (write to disk). When the cache holds the specified start percentage of unwritten data, a flush is triggered. By default, the controller starts flushing the cache when the cache reaches 80 percent full.
- **Cache block size** — the maximum size of each cache block, which is an organizational unit for cache management. The cache block size is by default 8 KiB, but can be set to 4, 8, 16, or 32 KiB. Ideally the cache block size should be set to the predominant I/O size of your applications. File systems or database applications generally use smaller sizes, while a larger size is good for applications requiring large data transfer or sequential I/O.

Volume cache settings

For individual volumes in a storage array, you can specify the following values from the Volumes page (**Storage > Volumes**):

- **Read caching** — The read cache is a buffer that stores data that has been read from the drives. The data for a read operation might already be in the cache from a previous operation, which eliminates the need to access the drives. The data stays in the read cache until it is flushed.
 - **Dynamic read cache prefetch** — Dynamic cache read prefetch allows the controller to copy additional sequential data blocks into the cache while it is reading data blocks from a drive to the cache. This

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

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



Possible loss of data — If you enable the* Write caching without batteries option and do not have a universal power supply for protection, you could lose data. In addition, you could lose data if you do not have controller batteries and you enable the Write caching without batteries option.

- **Write caching without batteries** — The write caching without batteries setting lets write caching continue even when the batteries are missing, failed, discharged completely, or not fully charged. Choosing write caching without batteries is not typically recommended, because data might be lost if power is lost. Typically, write caching is turned off temporarily by the controller until the batteries are charged or a failed battery is replaced.
- **Write caching with mirroring** — Write caching with mirroring occurs when the data written to the cache memory of one controller is also written to the cache memory of the other controller. Therefore, if one controller fails, the other can complete all outstanding write operations. Write cache mirroring is available only if write caching is enabled and two controllers are present. Write caching with mirroring is the default setting at volume creation.

Automatic load balancing overview

Automatic load balancing provides improved I/O resource management by reacting dynamically to load changes over time and automatically adjusting volume controller ownership to correct any load imbalance issues when workloads shift across the controllers.

The workload of each controller is continually monitored and, with cooperation from the multipath drivers installed on the hosts, can be automatically brought into balance whenever necessary. When workload is automatically re-balanced across the controllers, the storage administrator is relieved of the burden of manually adjusting volume controller ownership to accommodate load changes on the storage array.

When Automatic Load Balancing is enabled, it performs the following functions:

- Automatically monitors and balances controller resource utilization.
- Automatically adjusts volume controller ownership when needed, thereby optimizing I/O bandwidth between the hosts and the storage array.

Enabling and disabling Automatic Load Balancing

Automatic Load Balancing is enabled by default on all storage arrays.

You might want to disable Automatic Load Balancing on your storage array for the following reasons:

- You do not want to automatically change a particular volume's controller ownership to balance workload.
- You are operating in a highly tuned environment where load distribution is purposefully set up to achieve a specific distribution between the controllers.

Host types that support the Automatic Load Balancing feature

Even though Automatic Load Balancing is enabled at the storage array level, the host type you select for a host or host cluster has a direct influence on how the feature operates.

When balancing the storage array's workload across controllers, the Automatic Load Balancing feature attempts to move volumes that are accessible by both controllers and that are mapped only to a host or host cluster capable of supporting the Automatic Load Balancing feature.

This behavior prevents a host from losing access to a volume due to the load balancing process; however, the presence of volumes mapped to hosts that do not support Automatic Load Balancing affects the storage array's ability to balance workload. For Automatic Load Balancing to balance the workload, the multipath driver must support TPGS and the host type must be included in the following table.



For a host cluster to be considered capable of Automatic Load Balancing, all hosts in that group must be capable of supporting Automatic Load Balancing.

Host type supporting Automatic Load Balancing	With this multipath driver
Windows or Windows Clustered	MPIO with NetApp E-Series DSM
Linux DM-MP (Kernel 3.10 or later)	DM-MP with <code>scsi_dh_alua</code> device handler
VMware	Native Multipathing Plugin (NMP) with <code>VMW_SATP_ALUA</code> Storage Array Type plug-in



With minor exceptions, host types that do not support Automatic Load Balancing continue to operate normally whether or not the feature is enabled. One exception is that if a system has a failover, storage arrays move unmapped or unassigned volumes back to the owning controller when the data path returns. Any volumes that are mapped or assigned to non-Automatic Load Balancing hosts are not moved.

See the [Interoperability Matrix Tool](#) for compatibility information for specific multipath driver, OS level, and controller-drive tray support.

Verifying OS compatibility with the Automatic Load Balancing feature

Verify OS compatibility with the Automatic Load Balancing feature before setting up a new (or migrating an existing) system.

1. Go to the [Interoperability Matrix Tool](#) to find your solution and verify support.

If your system is running Red Hat Enterprise Linux 6 or SUSE Linux Enterprise Server 11, contact technical support.

2. Update and configure the `/etc/multipath.conf` file.
3. Ensure that both `retain_attached_device_handler` and `detect_prio` are set to `yes` for the applicable vendor and product, or use default settings.

Default host operating system type

The default host type is used by the storage array when hosts are initially connected. It defines how the controllers in the storage array work with the host's operating system when volumes are accessed. You can change the host type if there is a need to change how the storage array operates, relative to the hosts that are connected to it.

Generally, you will change the default host type before you connect hosts to the storage array or when you connect additional hosts.

Keep these guidelines in mind:

- If all of the hosts you plan to connect to the storage array have the same operating system (homogenous host environment), then change the host type to match the operating system.
- If there are hosts with different operating systems that you plan to connect to the storage array (heterogeneous host environment), change the host type to match the majority of the hosts' operating systems.

For example, if you are connecting eight different hosts to the storage array, and six of those hosts are running a Windows operating system, you must select Windows as the default host operating system type.

- If the majority of the connected hosts have a mix of different operating systems, change the host type to Factory Default.

For example, if you are connecting eight different hosts to the storage array, and two of those hosts are running a Windows operating system, three are running an HP-UX operating system, and another three are running a Linux operating system, you must select Factory Default as the default host operating system type.

How to

Edit storage array name

You can change the storage array name that appears in the title bar of SANtricity System Manager.

Steps

1. Select **Settings** > **System**.
2. Under **General**, look for the **Name:** field.

If a storage array name has not been defined, this field displays "Unknown."

3. Click the **Edit** (pencil) icon next to the storage array name.

The field becomes editable.

4. Enter a new name.

A name can contain letters, numbers, and the special characters underscore (_), dash (-), and hash sign (#). A name cannot contain spaces. A name can have a maximum length of 30 characters. The name must be unique.

5. Click the **Save** (check mark) icon.



If you want to close the editable field without making changes, click the Cancel (X) icon.

Result

The new name appears in the title bar of SANtricity System Manager.

Turn on storage array locator lights

To find the physical location of a storage array in a cabinet, you can turn on its locator (LED) lights.

Steps

1. Select **Settings > System**.
2. Under **General**, click **Turn on Storage Array Locator Lights**.

The **Turn On Storage Array Locator Lights** dialog box opens, and the corresponding storage array's locator lights turn on.

3. When you have physically located the storage array, return to the dialog box and select **Turn Off**.

Results

The locator lights turn off, and the dialog box closes.

Synchronize storage array clocks

If Network Time Protocol (NTP) is not enabled, you can manually set the clocks on the controllers so they are synchronized with the management client (the system used to run the browser that accesses SANtricity System Manager).

About this task

Synchronization ensures that event time stamps in the event log match time stamps written to the host log files. During the synchronization process, the controllers remain available and operational.



If NTP is enabled in System Manager, do not use this option to synchronize clocks. Instead, NTP automatically synchronizes the clocks with an external host using SNTP (Simple Network Time Protocol).



After synchronization, you might notice that performance statistics are lost or skewed, schedules are impacted (ASUP, snapshots, etc.), and time stamps in log data are skewed. Using NTP avoids this problem.

Steps

1. Select **Settings > System**.
2. Under **General**, click **Synchronize Storage Array Clocks**.

The **Synchronize Storage Array Clocks** dialog box opens. It shows the current date and time for the controller(s) and the computer used as the management client.



For simplex storage arrays, only one controller is shown.

3. If the times shown in the dialog box do not match, click **Synchronize**.

Results

After synchronization is successful, event time stamps are the same for the event log and host logs.

Save storage array configuration

You can save a storage array's configuration information in a script file to save time setting up additional storage arrays with the same configuration.

Before you begin

The storage array must not be undergoing any operation that changes its logical configuration settings. Examples of these operations include creating or deleting volumes, downloading controller firmware, assigning or modifying hot spare drives, or adding capacity (drives) to a volume group.

About this task

Saving the storage array configuration generates a command line interface (CLI) script that contains storage array settings, volume configuration, host configuration, or host-to-volume assignments for a storage array. You can use this generated CLI script to replicate a configuration to another storage array with the exact same hardware configuration.

However, you should not use this generated CLI script for disaster recovery. Instead, to do a system restore, use the configuration database backup file that you create manually or contact technical support to get this data from the latest Auto-Support data.

This operation *does not* save these settings:

- The life of the battery
- The controller time-of-day
- The nonvolatile static random access memory (NVSRAM) settings
- Any premium features
- The storage array password
- The operating status and states of the hardware components
- The operating status (except Optimal) and states of the volume groups
- Copy services, such as mirroring and volume copy



Risk of application errors — Do not use this option if the storage array is undergoing an operation that will change any logical configuration setting. Examples of these operations include creating or deleting volumes, downloading controller firmware, assigning or modifying hot spare drives, or adding capacity (drives) to a volume group.

Steps

1. Select **Settings > System**.
2. Select **Save Storage Array Configuration**.
3. Select the items of the configuration that you want to save:

- **Storage array settings**
- **Volume configuration**
- **Host configuration**
- **Host-to-volume assignments**



If you select the **Host-to-volume assignments** item, the **Volume configuration** item and the **Host configuration** item are also selected by default. You cannot save **Host-to-volume assignments** without also saving **Volume configuration** and **Host configuration**.

4. Click **Save**.

The file is saved in the Downloads folder for your browser with the name `storage-array-configuration.cfg`.

After you finish

To load a storage array configuration onto another storage array, use the SANtricity Unified Manager.

Clear storage array configuration

Use the Clear Configuration operation when you want to delete all the pools, volume groups, volumes, host definitions, and host assignments from the storage array.

Before you begin

- Before clearing the storage array configuration, back up the data.

About this task

There are two Clear Storage Array Configuration options:

- **Volume** — Typically, you might use the Volume option to reconfigure a test storage array as a production storage array. For example, you might configure a storage array for testing, and then, when you are done testing, remove the test configuration and set up the storage array for a production environment.
- **Storage Array** — Typically, you might use the Storage Array option to move a storage array to another department or group. For example, you might be using a storage array in Engineering, and now Engineering is getting a new storage array, so you want to move the current storage array to Administration where it will be reconfigured.

The Storage Array option deletes some additional settings.

	Volume	Storage Array
Deletes pools and volume groups	X	X
Deletes volumes	X	X
Deletes hosts and host clusters	X	X
Deletes host assignments	X	X

	Volume	Storage Array
Deletes storage array name		X
Resets storage array cache settings to default		X



Risk of data loss — This operation deletes all data from your storage array. (It does not do a secure erase.) You cannot cancel this operation after it starts. Perform this operation only when instructed to do so by technical support.

Steps

1. Select **Settings > System**.
2. Select **Clear Storage Array Configuration**.
3. In the drop-down list, select either **Volume** or **Storage Array**.
4. **Optional:** If you want to save the configuration (not the data), use the links in the dialog box.
5. Confirm that you want to perform the operation.

Results

- The current configuration is deleted, destroying all existing data on the storage array.
- All drives are unassigned.

Configure login banner

You can create a login banner that is presented to users before they establish sessions in SANtricity System Manager. The banner can include an advisory notice and a consent message.

About this task

When you create a banner, it appears before the login screen in a dialog box.

Steps

1. Select **Settings > System**.
2. Under the **General** section, select **Configure Login Banner**.

The **Configure Login Banner** dialog box opens.

3. Enter the text you want to appear in the login banner.



Do not use HTML or other markup tags for formatting.

4. Click **Save**.

Result

The next time users log in to System Manager, the text opens in a dialog box. Users must click **OK** to continue to the login screen.

Manage session timeouts

You can configure timeouts in SANtricity System Manager, so that users' inactive sessions are disconnected after a specified time.

About this task

By default, the session timeout for System Manager is 30 minutes. You can adjust that time or you can disable session timeouts altogether.



If Access Management is configured using the Security Assertion Markup Language (SAML) capabilities embedded in the array, a session timeout might occur when the user's SSO session reaches its maximum limit. This might occur before the System Manager session timeout.

Steps

1. Select **Settings > System**.
2. Under the **General** section, select **Enable/Disable Session Timeout**.

The **Enable/Disable Session Timeout** dialog box opens.

3. Use the spinner controls to increase or decrease the time in minutes.

The minimum timeout you can set for System Manager is 15 minutes.



To disable session timeouts, deselect the **Set the length of time...** checkbox.

4. Click **Save**.

Change cache settings for the storage array

For all volumes in the storage array, you can adjust the cache memory settings for flushing and block size.

About this task

Cache memory is an area of temporary volatile storage on the controller, which has a faster access time than the drive media. To tune cache performance, you can adjust the following settings:

Cache setting	Description
Start demand cache flushing	Start demand cache flushing specifies the percentage of unwritten data in the cache that triggers a cache flush (write to disk). By default, cache flushing starts when unwritten data reaches 80% capacity. A higher percentage is a good choice for environments with primarily write operations, so new write requests can be processed by cache without having to go to the disk. Lower settings are better in environments where the I/O is erratic (with data bursts), so that the system flushes cache frequently between data bursts. However, a start percentage lower than 80% may cause decreased performance.

Cache setting	Description
Cache block size	The cache block size determines the maximum size of each cache block, which is an organizational unit for cache management. By default, the block size is 8 KiB. System Manager allows the cache block size to be 4, 8, 16, or 32 KiBs. Applications use different block sizes, which have an impact on storage performance. A smaller size is a good choice for file systems or database applications. A larger size is ideal for applications that generate sequential I/O, such as multimedia.

Steps

1. Select **Settings > System**.
2. Scroll down to **Additional Settings**, and then click **Change Cache Settings**.

The Change Cache Settings dialog box opens.

3. Adjust the following values:
 - Start demand cache flushing — Choose a percentage that is appropriate for the I/O used in your environment. If you choose a value lower than 80%, you may notice decreased performance.
 - Cache block size — Choose a size that is appropriate for your applications.
4. Click **Save**.

Set host connectivity reporting

You can enable host connectivity reporting so the storage array continuously monitors the connection between the controllers and the configured hosts, and then alerts you if the connection is disrupted. This feature is enabled by default.

About this task

If you disable host connectivity reporting, the system no longer monitors connectivity or multipath driver issues with a host connected to the storage array.



Disabling host connectivity reporting also disables automatic load balancing, which monitors and balances controller resource utilization.

Steps

1. Select **Settings > System**.
2. Scroll down to **Additional Settings**, and then click **Enable/Disable Host Connectivity Reporting**.

The text below this option indicates whether it is currently enabled or disabled.

A confirmation dialog opens.

3. Click **Yes** to continue.

By selecting this option, you toggle the feature between enabled/disabled.

Set automatic load balancing

The **Automatic Load Balancing** feature ensures that incoming I/O traffic from the hosts is dynamically managed and balanced across both controllers. This feature is enabled by default, but you can disable it from System Manager.

About this task

When Automatic Load Balancing is enabled, it performs the following functions:

- Automatically monitors and balances controller resource utilization.
- Automatically adjusts volume controller ownership when needed, thereby optimizing I/O bandwidth between the hosts and the storage array.

You might want to disable Automatic Load Balancing on your storage array for the following reasons:

- You do not want to automatically change a particular volume's controller ownership to balance workload.
- You are operating in a highly tuned environment where load distribution is purposefully set up to achieve a specific distribution between the controllers.

Steps

1. Select **Settings > System**.
2. Scroll down to **Additional Settings**, and then click **Enable/Disable Automatic Load Balancing**.

The text below this option indicates whether the feature is currently enabled or disabled.

A confirmation dialog opens.

3. Confirm by clicking **Yes** to continue.

By selecting this option, you toggle the feature between enabled/disabled.



If this feature is moved from disabled to enabled, the Host Connectivity Reporting feature is automatically enabled as well.

Change default host type

Use the Change Default Host Operating System setting to change the default host type at the storage array level. Generally, you will change the default host type before you connect hosts to the storage array or when you connect additional hosts.

About this task

Keep these guidelines in mind:

- If all of the hosts you plan to connect to the storage array have the same operating system (homogenous host environment), then change the host type to match the operating system.
- If there are hosts with different operating systems that you plan to connect to the storage array (heterogeneous host environment), change the host type to match the majority of the hosts' operating systems.

For example, if you are connecting eight different hosts to the storage array, and six of those hosts are running a Windows operating system, you must select Windows as the default host operating system type.

- If the majority of the connected hosts have a mix of different operating systems, change the host type to Factory Default.

For example, if you are connecting eight different hosts to the storage array, and two of those hosts are running a Windows operating system, three are running an HP-UX operating system, and another three are running a Linux operating system, you must select Factory Default as the default host operating system type.

Steps

1. Select **Settings > System**.
2. Scroll down to **Additional Settings**, and then click **Change Default Host Operating System Type**.
3. Select the host operating system type that you want to use as the default.
4. Click **Change**.

Enable or disable legacy management interface

You can enable or disable the legacy management interface (SYMBOL), which is a method of communication between the storage array and the management client. By default, the legacy management interface is on. If you disable it, the storage array and management client will use a more secure method of communication (REST API over https); however, certain tools and tasks might be affected if it is disabled.

About this task

The setting affects operations as follows:

- **On** (default) — Required setting for mirroring, for CLI commands that operate only on E5700 and E5600 storage arrays, and some other tools like the QuickConnect utility and the OCI adapter.
- **Off** — Required setting to enforce confidentiality in communications between the storage array and the management client, and to access external tools. Recommended setting when configuring a Directory Server (LDAP).

Steps

1. Select **Settings > System**.
2. Scroll down to **Additional Settings**, and then click **Change Management Interface**.
3. In the dialog box, click **Yes** to continue.

FAQs

What is controller cache?

The controller cache is a physical memory space that streamlines two types of I/O (input/output) operations: between the controllers and hosts, and between the controllers and disks.

For read and write data transfers, the hosts and controllers communicate over high-speed connections. However, communications from the back-end of the controller to the disks is slower, because disks are relatively slow devices.

When the controller cache receives data, the controller acknowledges to the host applications that it is now

holding the data. This way, the host applications do not need to wait for the I/O to be written to disk. Instead, applications can continue operations. The cached data is also readily accessible by server applications, eliminating the need for extra disk reads to access the data.

The controller cache affects the overall performance of the storage array in several ways:

- The cache acts as a buffer, so that host and disk data transfers do not need to be synchronized.
- The data for a read or write operation from the host might be in cache from a previous operation, which eliminates the need to access the disk.
- If write caching is used, the host can send subsequent write commands before the data from a previous write operation is written to disk.
- If cache prefetch is enabled, sequential read access is optimized. Cache prefetch makes a read operation more likely to find its data in the cache, instead of reading the data from disk.



Possible loss of data — If you enable the **Write caching without batteries** option and do not have a universal power supply for protection, you could lose data. In addition, you could lose data if you do not have controller batteries and you enable the **Write caching without batteries** option.

What is cache flushing?

When the amount of unwritten data in the cache reaches a certain level, the controller periodically writes cached data to a drive. This write process is called "flushing."

The controller uses two algorithms for flushing cache: demand-based and age-based. The controller uses a demand-based algorithm until the amount of cached data drops below the cache flush threshold. By default, a flush begins when 80 percent of the cache is in use.

In System Manager, you can set the "Start demand cache flushing" threshold to best support the type of I/O used in your environment. In an environment that is primarily write operations, you should set the "Start demand cache flushing" percentage high to increase the probability that any new write requests can be processed by cache without having to go to the disk. A high percentage setting limits the number of cache flushes so that more data remains in cache, which increases the chance of more cache hits.

In an environment where the I/O is erratic (with data bursts), you can use low cache flushing so that the system flushes cache frequently between data bursts. In a diverse I/O environment that processes a variety of loads, or when the type of loads are unknown, set the threshold at 50 percent as a good middle ground. Be aware that if you choose a start percentage lower than 80 percent, you might see decreased performance because data needed for a host read might not be available. Choosing a lower percentage also increases the number of disk writes necessary to maintain the cache level, which increases system overhead.

The age-based algorithm specifies the period of time during which write data can remain in the cache before it is eligible to be flushed to the disks. The controllers use the age-based algorithm until the cache flush threshold is reached. The default is 10 seconds, but this time period is counted only during periods of inactivity. You cannot modify the flush timing in System Manager; instead, you must use the Set Storage Array command in the command-line interface (CLI).



Possible loss of data — If you enable the **Write caching without batteries** option and do not have a universal power supply for protection, you could lose data. In addition, you could lose data if you do not have controller batteries and you enable the **Write caching without batteries** option.

What is cache block size?

The storage array's controller organizes its cache into "blocks," which are chunks of memory that can be 4, 8, 16, or 32 KiBs in size. All volumes on the storage system share the same cache space; therefore, the volumes can have only one cache block size.



Cache blocks are not the same as the 512-byte blocks that are used by the logical block system of the disks.

Applications use different block sizes, which can have an impact on storage performance. By default, the block size in System Manager is 8 KiB, but you can set the value to 4, 8, 16, or 32 KiBs. A smaller size is a good choice for file systems or database applications. A larger size is a good choice for applications that require large data transfer, sequential I/O, or high bandwidth, such as multimedia.

When should I synchronize storage array clocks?

You should manually synchronize the controller clocks in the storage array if you notice that the time stamps shown in System Manager are not aligned with time stamps shown in your management client (the computer that is accessing System Manager through the browser). This task is only necessary if NTP (Network Time Protocol) is not enabled in System Manager.



We highly recommend that you use an NTP server instead of manually synchronizing the clocks. NTP automatically synchronizes the clocks with an external server using SNTP (Simple Network Time Protocol).

You can check synchronization status from the **Synchronize Storage Array Clocks** dialog box, which is available from the System page. If the times shown in the dialog box do not match, run a synchronization. You can periodically view this dialog box, which indicates whether the controller clocks' time displays have drifted apart and are no longer synchronized.

What is host connectivity reporting?

When host connectivity reporting is enabled, the storage array continuously monitors the connection between the controllers and the configured hosts, and then alerts you if the connection is disrupted.

Disruptions to the connection might occur if there is a loose, damaged, or missing cable, or another problem with the host. In these situations, the system might open a Recovery Guru message:

- **Host Redundancy Lost** — Opens if either controller cannot communicate with the host.
- **Host Type Incorrect** — Opens if the host's type is incorrectly specified on the storage array, which could result in failover problems.

You might want to disable host connectivity reporting in situations where rebooting a controller might take longer than the connection timeout. Disabling this feature suppresses Recovery Gurus messages.



Disabling host connectivity reporting also disables automatic load balancing, which monitors and balances controller resource use. However, if you re-enable host connectivity reporting, the automatic load balancing feature is not automatically re-enabled.

iSCSI settings

Concepts

iSCSI terminology

Learn how the iSCSI terms apply to your storage array.

Term	Description
CHAP	The Challenge Handshake Authentication Protocol (CHAP) method validates the identity of targets and initiators during the initial link. Authentication is based on a shared security key called a <i>CHAPsecret</i> .
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the System Manager functions.
DHCP	Dynamic Host Configuration Protocol (DHCP) is a protocol used on Internet Protocol (IP) networks for dynamically distributing network configuration parameters, such as IP addresses.
IB	InfiniBand (IB) is a communications standard for data transmission between high-performance servers and storage systems.
ICMP PING response	Internet Control Message Protocol (ICMP) is a protocol used by operating systems of networked computers to send messages. ICMP messages determine whether a host is reachable and how long it takes to get packets to and from that host.
IQN	An iSCSI Qualified Name (IQN) identifier is a unique name for an iSCSI initiator or iSCSI target.
iSER	iSCSI Extensions for RDMA (iSER) is a protocol that extends the iSCSI protocol for operation over RDMA transports, such as InfiniBand or Ethernet.
iSNS	Internet Storage Name Service (iSNS) is a protocol that allows automated discovery, management, and configuration of iSCSI and Fibre Channel devices on TCP/IP networks.
MAC address	Media access control identifiers (MAC addresses) are used by Ethernet to distinguish between separate logical channels connecting two ports on the same physical transport network interface.
Management client	A management client is the computer where a browser is installed for accessing System Manager.
MTU	A Maximum Transmission Unit (MTU) is the largest size packet or frame that can be sent in a network.

Term	Description
RDMA	Remote Direct Memory Access (RDMA) is a technology that allows network computers to exchange data in main memory without involving the operating system of either computer.
Unnamed discovery session	When the option for unnamed discovery sessions is enabled, iSCSI initiators are not required to specify the target IQN to retrieve the controller's information.

How to

Configure iSCSI ports

If your controller includes an iSCSI host connection, you can configure the iSCSI port settings from the Hardware page or the System page.

Before you begin

- Your controller must include iSCSI ports; otherwise, the iSCSI settings are not available.
- You must know the network speed (the data transfer rate between the ports and the host).

About this task

This task describes how to access the iSCSI port configuration from the Hardware page. You can also access the configuration from the System page (**Settings** › **System**).



The iSCSI settings and functions only appear if your storage array supports iSCSI.

Steps

1. Select **Hardware**.
2. If the graphic shows the drives, click **Show back of shelf**.

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

3. Click the controller with the iSCSI ports you want to configure.

The controller's context menu appears.

4. Select **Configure iSCSI ports**.



The **Configure iSCSI ports** option appears only if System Manager detects iSCSI ports on the controller.

The Configure iSCSI Ports dialog box opens.

5. In the drop-down list, select the port you want to configure, and then click **Next**.
6. Select the configuration port settings, and then click **Next**.

To see all port settings, click the Show more port settings link on the right of the dialog box.

Field Details

Port Setting	Description
Enable IPv4 / Enable IPv6	Select one or both options to enable support for IPv4 and IPv6 networks. NOTE: If you want to disable port access, deselect both check boxes.
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.

7. Configure the IPv4 and/or IPv6 settings, either automatically or manually. To see all port settings, click the **Show more settings** link on the right of the dialog box.

Field Details

Port setting	Description
Automatically obtain configuration	Select this option to obtain the configuration automatically.
Manually specify static configuration	Select this option, and then enter a static address in the fields. (If desired, you can cut and paste addresses into the fields.) For IPv4, include the network subnet mask and gateway. For IPv6, include the routable IP address and router IP address.
Enable VLAN support (Available by clicking Show more settings.)	Select this option to enable a VLAN and enter its ID. A VLAN is a logical network that behaves like it is physically separate from other physical and virtual local area networks (LANs) supported by the same switches, the same routers, or both.
Enable ethernet priority (Available by clicking Show more settings.)	<p>Select this option to enable the parameter that determines the priority of accessing the network. Use the slider to select a priority between 1 (lowest) and 7 (highest).</p> <p>In a shared local area network (LAN) environment, such as Ethernet, many stations might contend for access to the network. Access is on a first-come, first-served basis. Two stations might try to access the network at the same time, which causes both stations to back off and wait before trying again. This process is minimized for switched Ethernet, where only one station is connected to a switch port.</p>

8. Click **Finish**.

Configure 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 *CHAPsecret*.

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**.

Result

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**.

Result

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 subnet.
- **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.

End iSCSI session

You can end an iSCSI session that is no longer needed. iSCSI sessions can occur with hosts or remote storage arrays in an asynchronous mirror relationship.

About this task

You might want to end an iSCSI session for these reasons:

- **Unauthorized access** — If an iSCSI initiator is logged on and should not have access, you can end the iSCSI session to force the iSCSI initiator off the storage array. The iSCSI initiator could have logged on because the None authentication method was available.
- **System downtime** — If you need to take down a storage array and you see that iSCSI initiators are still logged on, you can end the iSCSI sessions to get the iSCSI initiators off the storage array.

Steps

1. Select **Support > Support Center > Diagnostics** tab.
2. Select **View/End iSCSI Sessions**.

A list of the current iSCSI sessions appears.

3. Select the session that you want to end.
4. Click **End Session**, and confirm that you want to perform the operation.

View iSCSI sessions

You can view detailed information about the iSCSI connections to your storage array. iSCSI sessions can occur with hosts or remote storage arrays in an asynchronous mirror relationship.

Steps

1. Select **Support > Support Center > Diagnostics** tab.
2. Select **View/End iSCSI Sessions**.

A list of the current iSCSI sessions appears.

3. To see additional information about a specific iSCSI session, select a session, and then click **View Details**.

Field Details

Item	Description
Session Identifier (SSID)	A hexadecimal string that identifies a session between an iSCSI initiator and an iSCSI target. The SSID is composed of the ISID and the TPGT.
Initiator Session ID (ISID)	The initiator part of the session identifier. The initiator specifies the ISID during login.
Target Portal Group	The iSCSI target.
Target Portal Group Tag (TPGT)	The target part of the session identifier. A 16-bit numerical identifier for an iSCSI target portal group.
Initiator iSCSI name	The worldwide unique name of the initiator.
Initiator iSCSI label	The user label set in System Manager.
Initiator iSCSI alias	A name that also can be associated with an iSCSI node. The alias allows an organization to associate a user-friendly string with the iSCSI name. However, the alias is not a substitute for the iSCSI name. The initiator iSCSI alias only can be set at the host, not in System Manager
Host	A server that sends input and output to the storage array.
Connection ID (CID)	A unique name for a connection within the session between the initiator and the target. The initiator generates this ID and presents it to the target during login requests. The connection ID is also presented during logouts that close connections.
Ethernet port identifier	The controller port associated with the connection.
Initiator IP address	The IP address of the initiator.
Negotiated login parameters	The parameters that are transacted during the login of the iSCSI session.
Authentication method	The technique to authenticate users who want access to the iSCSI network. Valid values are CHAP and None .
Header digest method	The technique to show possible header values for the iSCSI session. HeaderDigest and DataDigest can be either None or CRC32C . The default value for both is None .
Data digest method	The technique to show possible data values for the iSCSI session. HeaderDigest and DataDigest can be either None or CRC32C . The default value for both is None .

Item	Description
Maximum connections	The greatest number of connections allowed for the iSCSI session. The maximum number of connections can be 1 through 4. The default value is 1 .
Target alias	The label associated with the target.
Initiator alias	The label associated with the initiator.
Target IP address	The IP address of the target for the iSCSI session. DNS names are not supported.
Initial R2T	The initial ready to transfer status. The status can be either Yes or No .
Maximum burst length	The maximum SCSI payload in bytes for this iSCSI session. The maximum burst length can be from 512 to 262,144 (256 KB). The default value is 262,144 (256 KB) .
First burst length	The SCSI payload in bytes for unsolicited data for this iSCSI session. The first burst length can be from 512 to 131,072 (128 KB). The default value is 65,536 (64 KB) .
Default time to wait	The minimum number of seconds to wait before you attempt to make a connection after a connection termination or a connection reset. The default time to wait value can be from 0 to 3600. The default is 2 .
Default time to retain	The maximum number of seconds that connection is still possible following a connection termination or a connection reset. The default time to retain can be from 0 to 3600. The default value is 20 .
Maximum outstanding R2T	The maximum number of "ready to transfers" outstanding for this iSCSI session. The maximum outstanding ready to transfer value can be from 1 to 16. The default is 1 .
Error recovery level	The level of error recovery for this iSCSI session. The error recovery level value is always set to 0 .
Maximum receive data segment length	The maximum amount of data that either the initiator or the target can receive in any iSCSI payload data unit (PDU).
Target name	The official name of the target (not the alias). The target name with the <i>iqn</i> format.
Initiator name	The official name of the initiator (not the alias). The initiator name that uses either the <i>iqn</i> or <i>eui</i> format.

4. To save the report to a file, click **Save**.

The file is saved in the Downloads folder for your browser with the filename `iscsi-session-connections.txt`.

Configure iSER over InfiniBand ports

If your controller includes an iSER over InfiniBand port, you can configure the network connection to the host. The configuration settings are available from the **Hardware** page or the **System** page.

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.

About this task

You can access the iSER over InfiniBand configuration from the **Hardware** page or from **Settings > System**. This task describes how to configure the ports from the **Hardware** page.



The iSER over InfiniBand settings and functions appear only if your storage array's controller includes an iSER 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 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.

You can access iSER over InfiniBand statistics from the System page (**Settings > System**) or from the Support page. These instructions describe how to access the statistics from the Support page.

Steps

1. Select **Support > Support Center > Diagnostics** tab.
2. Select **View iSER over InfiniBand Statistics**.
3. Click a tab to view the different sets of statistics.
4. To set the baseline, click **Set new baseline**.

Setting the baseline sets a new starting point for the collection of the statistics. The same baseline is used for all iSER over InfiniBand statistics.

FAQs

What happens when I use an iSNS server for registration?

When Internet Storage Name Service (iSNS) server information is used, the hosts (initiators) can be configured to query the iSNS server to retrieve information from the target (controllers).

This registration provides the iSNS server with the controller's iSCSI Qualified Name (IQN) and port information, and allows for queries between the initiators (iSCSI hosts) and targets (controllers).

Which registration methods are automatically supported for iSCSI?

The iSCSI implementation supports either the Internet Storage Name Service (iSNS) discovery method or the use of the Send Targets command.

The iSNS method allows for iSNS discovery between the initiators (iSCSI hosts) and targets (the controllers). You register the target controller to provide the iSNS server with the controller's iSCSI Qualified Name (IQN) and port information.

If you do not configure iSNS, the iSCSI host can send the Send Targets command during an iSCSI discovery session. In response, the controller returns the port information (for example, the Target IQN, port IP address, listening port, and Target Port Group). This discovery method is not required if you use iSNS, because the host initiator can retrieve the target IPs from the iSNS server.

How do I interpret iSER over InfiniBand statistics?

The **View iSER over InfiniBand Statistics** dialog box displays local target (protocol) statistics and iSER over InfiniBand (IB) interface statistics. All statistics are read-only, and


cannot be set.

- **Local Target (Protocol) statistics** — Provides statistics for the iSER over InfiniBand target, which shows block-level access to its storage media.
- **iSER over InfiniBand Interface statistics** — Provides statistics for all iSER over InfiniBand ports on the InfiniBand interface, which includes performance statistics and link error information associated with each switch port.

You can view each of these statistics as raw statistics or as baseline statistics. Raw statistics are all of the statistics that have been gathered since the controllers were started. Baseline statistics are point-in-time statistics that have been gathered since you set the baseline time.

What else do I need to do to configure or diagnose iSER over InfiniBand?

The following table lists the System Manager functions that you can use to configure and manage iSER over InfiniBand sessions.



The iSER over InfiniBand settings are available only if your storage array’s controller includes an iSER over InfiniBand host management port.

Configure and diagnose iSER over InfiniBand

Action	Location
Configure iSER over InfiniBand ports	<div>1. Select Hardware.</div> <div>2. Select Show back of shelf.</div> <div>3. Select a controller.</div> <div>4. Select Configure iSER over InfiniBand ports.</div> <div>or</div> <div>1. Select Settings > System.</div> <div>2. Scroll down to iSER over InfiniBand settings, and then select Configure iSER over InfiniBand Ports.</div>
View iSER over InfiniBand statistics	<div>1. Select Settings > System.</div> <div>2. Scroll down to iSER over InfiniBand settings, and then select View iSER over InfiniBand Statistics.</div>

System: NVMe settings

Concepts

NVMe overview

Some controllers include a port for implementing NVMe (Non-Volatile Memory Express) over an InfiniBand fabric or over a RoCE (RDMA over Converged Ethernet) fabric. 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. For the SANtricity OS 11.40 release and higher, an NVMe storage array (called a *subsystem*) can be accessed by a host using an InfiniBand or RDMA fabric. NVMe commands are enabled and encapsulated in transport abstraction layers on both the host side and the subsystem side. This extends the high performance NVMe interface end-to-end from the host to the storage and standardizes and simplifies the command set.

NVMe-oF storage is presented to a host as a local block storage device. The volume (called a *namespace*) can be mounted to a file system as with any other block storage device. You can use the REST API, the SMcli, or SANtricity System Manager to provision your storage as needed.

What is an NVMe Qualified Name (NQN)?

The NVMe Qualified Name (NQN) is used to identify the remote storage target. The NVMe qualified name for the storage array is always assigned by the subsystem and may not be modified. There is only one NVMe Qualified Name for the entire array. The NVMe Qualified Name is limited to 223 characters in length. You can compare it to an iSCSI Qualified Name.

What is a namespace and a namespace ID?

A namespace is the equivalent of a logical unit in SCSI, which relates to a volume in the array. The namespace ID (NSID) is equivalent to a logical unit number (LUN) in SCSI. You create the NSID at namespace creation time, and can set it to a value between 1 and 255.

What is an NVMe controller?

Similar to a SCSI I_T nexus, which represents the path from the host's initiator to the storage system's target, an NVMe controller created during the host connection process provides an access path between a host and the namespaces in the storage array. An NQN for the host plus a host port identifier uniquely identify an NVMe controller. While an NVMe controller can only be associated with a single host, it can access multiple namespaces.

You configure which hosts can access which namespaces and set the namespace ID for the host using SANtricity System Manager. Then, when the NVMe controller is created, the list of namespace IDs accessible by the NVMe controller is created and used to configure the permissible connections.

NVMe terminology

Learn how the NVMe terms apply to your storage array.

Term	Description
InfiniBand	InfiniBand (IB) is a communications standard for data transmission between high-performance servers and storage systems.

Term	Description
Namespace	A namespace is NVM storage that is formatted for block access. It is analogous to a logical unit in SCSI, which relates to a volume in the storage array.
Namespace ID	The namespace ID is the NVMe controller's unique identifier for the namespace, and can be set to a value between 1 and 255. It is analogous to a logical unit number (LUN) in SCSI.
NQN	NVMe Qualified Name (NQN) is used to identify the remote storage target (the storage array).
NVM	Non-Volatile Memory (NVM) is persistent memory used in many types of storage devices.
NVMe	Non-Volatile Memory Express (NVMe) is an interface designed for flash-based storage devices, such as SSD drives. NVMe reduces I/O overhead and includes performance improvements, as compared to previous logical-device interfaces.
NVMe-oF	Non-Volatile Memory Express over Fabrics (NVMe-oF) is a specification that enables NVMe commands and data to transfer over a network between a host and storage.
NVMe controller	An NVMe controller is created during the host connection process. It provides an access path between a host and the namespaces in the storage array.
NVMe queue	A queue is used for passing commands and messages over the NVMe interface.
NVMe subsystem	The storage array with an NVMe host connection.
RDMA	Remote direct memory access (RDMA) enables more direct data movement in and out of a server by implementing a transport protocol in the network interface card (NIC) hardware.
RoCE	RDMA over Converged Ethernet (RoCE) is a network protocol that allows remote direct memory access (RDMA) over an Ethernet network.
SSD	Solid-state disks (SSDs) are data storage devices that use solid state memory (flash) to store data persistently. SSDs emulate conventional hard drives, and are available with the same interfaces that hard drives use.

How to

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 or the System page.

Before you begin

- Your controller must include an NVMe over InfiniBand host port; otherwise, the NVMe over InfiniBand settings are not available in System Manager.
- You must know the IP address of the host connection.

About this task

You can access the NVMe over InfiniBand configuration from the **Hardware** page or from **Settings > System**. This task describes how to configure the ports from the **Hardware** page.



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. 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 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 or the System page.

Before you begin

- Your controller must include an NVMe over RoCE host port; otherwise, the NVMe over RoCE settings are not available in System Manager.
- You must know the IP address of the host connection.

About this task

You can access the NVMe over RoCE configuration from the **Hardware** page or from **Settings > System**. This task describes how to configure the ports from the Hardware page.



The NVMe over RoCE settings and functions appear only if your storage array's controller includes an NVMe over RoCE 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 RoCE port you want to configure.

The controller's context menu appears.

4. Select **Configure NVMe over RoCE ports**.


The Configure NVMe over RoCE Ports dialog box opens.

5. In the drop-down list, select the HIC port you want to configure.

6. Click **Next**.

To see all port settings, click the **Show more port settings** link on the right of the dialog box.

Field Details

Port Setting	Description
Configured ethernet port speed	Select the speed that matches the speed capability of the SFP on the port.
Enable IPv4 / Enable IPv6	<div>Select one or both options to enable support for IPv4 and IPv6 networks.</div> <div> If you want to disable port access, deselect both check boxes.</div>
MTU size (Available by clicking Show more port settings.)	<div>If necessary, enter a new size in bytes for the Maximum Transmission Unit (MTU).</div> <div>The default Maximum Transmission Unit (MTU) size is 1500 bytes per frame. You must enter a value between 1500 and 9000.</div>

If you selected Enable IPv4, a dialog box opens for selecting IPv4 settings after you click Next. If you selected Enable IPv6, a dialog box opens for selecting IPv6 settings after you click Next. If you selected both options, the dialog box for IPv4 settings opens first, and then after you click Next, the dialog box for IPv6 settings opens.

7. Configure the IPv4 and/or IPv6 settings, either automatically or manually.

Field Details

Port setting	Description
Automatically obtain configuration	Select this option to obtain the configuration automatically.
Manually specify static configuration	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.

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** — Provides statistics for the NVMe controller, including timeouts and connection failures.
- **RDMA Interface statistics** — Provides statistics for the RDMA interface, including received and transmitted packet information.

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.

You can access NVMe over Fabrics statistics from the System page (**Settings > System**) or from the Support page. These instructions describe how to access the statistics from the Support page.

Steps

1. Select **Support > Support Center > Diagnostics** tab.
2. Select **View NVMe over Fabrics Statistics**.
3. To set the baseline, click **Set new baseline**.

Setting the baseline sets a new starting point for the collection of the statistics. The same baseline is used for all NVMe statistics.

FAQs

How do I interpret NVMe over InfiniBand statistics?

The **View NVMe over Fabrics Statistics** dialog box displays statistics for the NVMe

subsystem and the NVMe over InfiniBand 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. 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.

How do I interpret NVMe over Fabrics statistics?


The **View NVMe over Fabrics Statistics** dialog box displays statistics for the NVMe subsystem and the NVMe over RoCE 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. For more information about the statistics, click **View legend for table headings**.

You can view each of these statistics as raw statistics or as baseline statistics. Raw statistics are all of the statistics that have been gathered since the controllers were started. Baseline statistics are point-in-time statistics that have been gathered since you set the baseline time.

What else do I need to do to configure or diagnose NVMe over InfiniBand?

The following table lists the System Manager functions that you can use to configure and manage NVMe over InfiniBand sessions.



The NVMe over InfiniBand settings are available only if your storage array’s controller includes an NVMe over InfiniBand port.

Configure and diagnose NVMe over InfiniBand

Action	Location
Configure NVMe over InfiniBand ports	<ol style="list-style-type: none"> 1. Select Hardware. 2. Select Show back of shelf. 3. Select a controller. 4. Select Configure NVMe over InfiniBand ports. <p>or</p> <ol style="list-style-type: none"> 1. Select Settings > System. 2. Scroll down to NVMe over InfiniBand settings, and then select Configure NVMe over InfiniBand Ports.
View NVMe over InfiniBand statistics	<ol style="list-style-type: none"> 1. Select Settings > System. 2. Scroll down to NVMe over InfiniBand settings, and then select View NVMe over Fabrics Statistics.

What else do I need to do to configure or diagnose NVMe over RoCE?

You can configure and manage NVMe over RoCE from the Hardware and Settings pages.



The NVMe over RoCE settings are available only if your storage array's controller includes an NVMe over RoCE port.

Configure and diagnose NVMe over RoCE

Action	Location
Configure NVMe over RoCE ports	<ol style="list-style-type: none"> 1. Select Hardware. 2. Select Show back of shelf. 3. Select a controller. 4. Select Configure NVMe over RoCE ports. <p>or</p> <ol style="list-style-type: none"> 1. Select Settings > System. 2. Scroll down to NVMe over RoCE settings, and then select Configure NVMe over RoCE Ports.
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.

Add-on features

Concepts

How add-on features work

Add-ons are features that are not included in the standard configuration of System Manager and require a key to enable. An add-on feature can be either a single premium feature or a bundled feature pack.

The following steps provide an overview for enabling a premium feature or feature pack:

1. Obtain the following information:
 - Chassis serial number and the Feature Enable Identifier, which identify the storage array for the feature to be installed. These items are available in System Manager.
 - Feature Activation Code, which is available from the Support site when you purchase the feature.
2. Obtain the feature key by contacting your storage provider or by accessing the Premium Feature Activation site. Provide the chassis serial number, Feature Enable Identifier, and Feature Activation Code.
3. Using System Manager, enable the premium feature or feature pack using the feature key file.

Add-on feature terminology

Learn how the add-on feature terms apply to your storage array.

Term	Description
Feature Enable Identifier	A Feature Enable Identifier is a unique string that identifies the specific storage array. This identifier ensures that when you obtain the premium feature, it is associated with only that particular storage array. This string is displayed under Add-Ons on the System page.
Feature key file	A feature key file is a file you receive for unlocking and enabling a premium feature or feature pack.
Feature pack	A feature pack is a bundle that changes storage array attributes (for example, changing the protocol from Fibre Channel to iSCSI). Feature packs require a special key to enable them.
Premium feature	A premium feature is an extra option that requires a key to enable it. It is not included in the standard configuration of System Manager.

How tos

Obtain a feature key file

To enable a premium feature or feature pack on your storage array, you must first obtain a feature key file. A key is associated with only one storage array.

About this task

This task describes how to gather required information for the feature, and then send a request for a feature key file. Required information includes:

- Chassis serial number
- Feature Enable Identifier
- Feature Activation Code

Steps

1. In System Manager, locate and record the chassis serial number. You can view this serial number by hovering your mouse over the Support Center tile.
2. In System Manager, locate the Feature Enable Identifier. Go to **Settings > System**, and then scroll down to **Add-ons**. Look for the **Feature Enable Identifier**. Record the number for the Feature Enable Identifier.
3. Locate and record the Feature Activation Code. For features packs, this activation code is provided in the appropriate instructions for performing the conversion.

NetApp instructions are available from [NetApp E-Series Systems Documentation Center](#).

For premium features, you can access the activation code from the Support site, as follows:

- a. Log in to [NetApp Support](#).
 - b. Go to **Products > Manage Products > Software Licenses**.
 - c. Enter the serial number for the storage array chassis, and then click **Go**.
 - d. Look for the Feature Activation Codes in the **License Key** column.
 - e. Record the Feature Activation Code for the feature you want.
4. Request a feature key file by sending an email or a text document to your storage supplier with the following information: chassis serial number, the Feature Activation Code, and the Feature Enable Identifier.

You can also go to [NetApp License Activation: Storage Array Premium Feature Activation](#) and enter the required information to obtain the feature or feature pack. (The instructions on this site are for premium features, not feature packs.)

After you finish

When you have a feature key file, you can enable the premium feature or feature pack.

Enable a premium feature

A premium feature is an extra option that requires a key to enable.

Before you begin

- You have obtained a feature key. If necessary, contact technical support for a key.
- You have loaded the key file on the management client (the system with a browser for accessing System

Manager).

About this task

This task describes how to use System Manager to enable a premium feature.



If you want to disable a premium feature, you must use the Disable Storage Array Feature command (`disable storageArray (featurePack | feature=featureAttributeList)`) in the Command Line Interface (CLI).

Steps

1. Select **Settings > System**.
2. Under **Add-ons**, select **Enable Premium Feature**.

The Enable a Premium Feature dialog box opens.

3. Click **Browse**, and then select the key file.

The file name is displayed in the dialog box.

4. Click **Enable**.

Enable feature pack

A feature pack is a bundle that changes storage array attributes (for example, changing the protocol from Fibre Channel to iSCSI). Feature packs require a special key for enablement.

Before you begin

- You have followed the appropriate instructions for performing the conversion and for preparing your system for the new storage array attributes.



Conversion instructions are available from [NetApp E-Series Systems Documentation Center](#).

- The storage array is offline, so no hosts or applications are accessing it.
- All data is backed up.
- You have obtained a feature pack file.

The feature pack file is loaded on the management client (the system with a browser for accessing System Manager).



You must schedule a downtime maintenance window and stop all I/O operations between the host and controllers. In addition, be aware that you cannot access data on the storage array until you have successfully completed the conversion.

About this task

This task describes how to use System Manager to enable a feature pack. When you are done, you must restart the storage array.

Steps

1. Select **Settings > System**.
2. Under **Add-ons**, select **Change Feature Pack**.
3. Click **Browse**, and then select the key file.

The file name is displayed in the dialog box.

4. Type **CHANGE** in the field.
5. Click **Change**.

The feature pack migration begins and the controllers reboot. Unwritten cache data is deleted, which ensures no I/O activity. Both controllers automatically reboot for the new feature pack to take effect. The storage array returns to a responsive state after the reboot is complete.

Security key management

Concepts

How the Drive Security feature works

Drive Security is a storage array feature that provides an extra layer of security with either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives. When these drives are used with the Drive Security feature, they require a security key for access to their data. When the drives are physically removed from the array, they cannot operate until they are installed in another array, at which point, they will be in a Security Locked state until the correct security key is provided.

How to implement Drive Security

To implement Drive Security, you perform the following steps.

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

How Drive Security works at the drive level

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

The Drive Security feature provides an extra layer of protection with secure-capable drives. When volume groups or pools on these drives are selected for Drive Security, the drives look for a security key before allowing access to the data. You can enable Drive Security for pools and volume groups at any time, without affecting existing data on the drive. However, you cannot disable Drive Security without erasing all data on the drive.

How Drive Security works at the storage array level

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

If a secure-enabled drive is removed from the storage array and re-installed in a different storage array, the drive will be in a Security Locked state. The re-located drive looks for the security key before it makes the data accessible again. To unlock the data, you apply the security key from the source storage array. After a successful unlock process, the re-located drive will then use the security key already stored in the target storage array, and the imported security key file is no longer needed.



For internal key management, the actual security key is stored on the controller in a non-accessible location. It is not in human-readable format, nor is it user-accessible.

How Drive Security works at the volume level

When you create a pool or volume group from secure-capable drives, you can also enable Drive Security for those pools or volume groups. The Drive Security option makes the drives and associated volume groups and pools *secure-enabled*.

Keep the following guidelines in mind before creating secure-enabled volume groups and pools:

- Volume groups and pools must be comprised entirely of secure-capable drives. (For volumes that require FIPS support, use only FIPS drives. Mixing FIPS and FDE drives in a volume group or pool will result in all drives being treated as FDE drives. Also, an FDE drive cannot be added to or used as a spare in an all-FIPS volume group or pool.)
- Volume groups and pools must be in an optimal state.

How security key management works

When you implement the Drive Security feature, the secure-enabled drives (FIPS or FDE) require a security key for data access. A security key is a string of characters that is shared between these types of drives and the controllers in a storage array.

Whenever power to the drives is turned off and on, the secure-enabled drives change to a Security Locked state until the controller applies the security key. If a secure-enabled drive is removed from the storage array, the drive's data is locked. When the drive is re-installed in a different storage array, it looks for the security key before it makes the data accessible again. To unlock the data, you must apply the original security key.

You can create and manage security keys using one of the following methods:

- Internal key management on the controller's persistent memory.
- External key management on an external key management server.

Internal key management

Internal keys are maintained on the controller's persistent memory. To implement internal key management, you perform the following steps:

1. Install secure-capable drives in the storage array. These drives can be Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.
2. Make sure the Drive Security feature is enabled. If necessary, contact your storage vendor for instructions on enabling the Drive Security feature.
3. Create an internal security key, which involves defining an identifier and a pass phrase. The identifier is a string that is associated with the security key, and is stored on the controller and on all drives associated with the key. The pass phrase is used to encrypt the security key for backup purposes. To create an internal key, go to **Settings > System > Security key management > Create Internal Key**.

The security key is stored on the controller in a non-accessible location. You can then create secure-enabled volume groups or pools, or you can enable security on existing volume groups and pools.

External key management


External keys are maintained on a separate key management server, using a Key Management Interoperability Protocol (KMIP). To implement external key management, you perform the following steps:

1. Install secure-capable drives in the storage array. These drives can be Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.
2. Make sure the Drive Security feature is enabled. If necessary, contact your storage vendor for instructions on enabling the Drive Security feature.
3. Complete and download a client Certificate Signing Request (CSR) for authentication between the storage array and the key management server. Go to **Settings > Certificates > Key Management > Complete CSR**.
4. Create and download a client certificate from the key management server using the downloaded CSR file.
5. Ensure that the client certificate and a copy of the certificate for the key management server are available on your local host.
6. Create an external key, which involves defining the IP address of the key management server and the port number used for KMIP communications. During this process, you also load certificate files. To create an external key, go to **Settings > System > Security key management > Create External Key**.

The system connects to the key management server with the credentials you entered. You can then create secure-enabled volume groups or pools, or you can enable security on existing volume groups and pools.

Drive Security terminology

Learn how the Drive Security terms apply to your storage array.

Term	Description
Drive Security feature	Drive Security is a storage array feature that provides an extra layer of security with either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives. When these drives are used with the Drive Security feature, they require a security key for access to their data. When the drives are physically removed from the array, they cannot operate until they are installed in another array, at which point, they will be in a Security Locked state until the correct security key is provided.
FDE drives	Full Disk Encryption (FDE) drives perform encryption on the disk drive at the hardware level. The hard drive contains an ASIC chip that encrypts data during writes, and then decrypts data during reads.
FIPS drives	FIPS drives use Federal Information Processing Standards (FIPS) 140-2 level 2. They are essentially FDE drives that adhere to United States government standards for ensuring strong encryption algorithms and methods. FIPS drives have higher security standards than FDE drives.
Management client	A local system (computer, tablet, etc.) that includes a browser for accessing System Manager.
Pass phrase	<p>The pass phrase is used to encrypt the security key for backup purposes. The same pass phrase used to encrypt the security key must be provided when the backed up security key is imported as the result of a drive migration or head swap. A pass phrase can have between 8 and 32 characters.</p> <div>  <p>The pass phrase for Drive Security is independent from the storage array's Administrator password.</p> </div>
Secure-capable drives	Secure-capable drives can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives, which encrypt data during writes and decrypt data during reads. These drives are considered <i>secure-capable</i> because they can be used for additional security using the Drive Security feature. If the Drive Security feature is enabled for volume groups and pools used with these drives, the drives become <i>secure-enabled</i> .

Term	Description
Secure-enabled drives	Secure-enabled drives are used with the Drive Security feature. When you enable the Drive Security feature and then apply Drive Security to a pool or volume group on secure- <i>capable</i> drives, the drives become secure- <i>enabled</i> . Read and write access is available only through a controller that is configured with the correct security key. This added security prevents unauthorized access to the data on a drive that is physically removed from the storage array.
Security key	<p>A security key is a string of characters that is shared between the secure-enabled drives and controllers in a storage array. Whenever power to the drives is turned off and on, the secure-enabled drives change to a Security Locked state until the controller applies the security key. If a secure-enabled drive is removed from the storage array, the drive's data is locked. When the drive is re-installed in a different storage array, it looks for the security key before it makes the data accessible again. To unlock the data, you must apply the original security key. You can create and manage security keys using one of the following methods:</p> <ul style="list-style-type: none"> • Internal key management — Create and maintain security keys on the controller's persistent memory. • External key management — Create and maintain security keys on an external key management server.
Security key identifier	The security key identifier is a string that is associated with the security key during key creation. The identifier is stored on the controller and on all drives associated with the security key.

How tos

Create internal security key

To use the Drive Security feature, you can create an internal security key that is shared by the controllers and secure-capable drives in the storage array. Internal keys are maintained on the controller's persistent memory.

Before you begin

- Secure-capable drives must be installed in the storage array. These drives can be Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.
- The Drive Security feature must be enabled. Otherwise, a **Cannot Create Security Key** dialog box opens during this task. If necessary, contact your storage vendor for instructions on enabling the Drive Security

feature.



If both FDE and FIPS drives are installed in the storage array, they all share the same security key.

About this task

In this task, you define an identifier and a pass phrase to associate with the internal security key.



The pass phrase for Drive Security is independent from the storage array's Administrator password.

Steps

1. Select **Settings > System**.
2. Under **Security key management**, select **Create Internal Key**.

If you have not yet generated a security key, the **Create Security Key** dialog box opens.

3. Enter information in the following fields:

- Define a security key identifier — You can either accept the default value (storage array name and time stamp, which is generated by the controller firmware) or enter your own value. You can enter up to 189 alphanumeric characters without spaces, punctuation, or symbols.



Additional characters are generated automatically, appended to both ends of the string you enter. The generated characters ensure that the identifier is unique.

- Define a pass phrase/Re-enter pass phrase — Enter and confirm a pass phrase. The value can have between 8 and 32 characters, and must include each of the following:
 - An uppercase letter (one or more). Keep in mind that the pass phrase is case sensitive.
 - A number (one or more).
 - A non-alphanumeric character, such as !, *, @ (one or more).



Be sure to record your entries for later use. If you need to move a secure-enabled drive from the storage array, you must know the identifier and pass phrase to unlock drive data.

4. Click **Create**.

The security key is stored on the controller in a non-accessible location. Along with the actual key, there is an encrypted key file that is downloaded from your browser.



The path for the downloaded file might depend on the default download location of your browser.

5. Record your key identifier, pass phrase, and the location of the downloaded key file, and then click **Close**.

Result

You can now create secure-enabled volume groups or pools, or you can enable security on existing volume groups and pools.



Whenever power to the drives is turned off and then on again, all the secure-enabled drives change to a Security Locked state. In this state, the data is inaccessible until the controller applies the correct security key during drive initialization. If someone physically removes a locked drive and installs it in another system, the Security Locked state prevents unauthorized access to its data.

After you finish

You should validate the security key to make sure the key file is not corrupted.

Create external security key

To use the Drive Security feature with a key management server, you must create an external key that is shared by the key management server and the secure-capable drives in the storage array.

Before you begin

- Secure-capable drives must be installed in the array. These drives can be Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.



If both FDE and FIPS drives are installed in the storage array, they all share the same security key.

- The Drive Security feature must be enabled. Otherwise, a **Cannot Create Security Key** dialog box opens during this task. If necessary, contact your storage vendor for instructions on enabling the Drive Security feature.
- The client and server certificates are available on your local host so the storage array and key management server can authenticate each other. The client certificate validates the controllers, while the server certificate validates the key management server.

About this task

In this task, you define the IP address of the key management server and the port number it uses, and then load certificates for external key management.

Steps

1. Select **Settings > System**.
2. Under **Security key management**, select **Create External Key**.



If internal key management is currently configured, a dialog box opens and asks you to confirm that you want to switch to external key management.

The **Create External Security Key** dialog box opens.

3. Under **Connect to Key Server**, enter information in the following fields:
 - Key management server address — Enter the fully qualified domain name or the IP address (IPv4 or IPv6) of the server used for key management.
 - Key management port number — Enter the port number used for the Key Management Interoperability Protocol (KMIP) communications. The most common port number used for key management server communications is 5696.
 - Select client certificate — Click the first Browse button to select the certificate file for the storage array's

controllers.

- Select key management server's server certificate — Click the second Browse button to select the certificate file for the key management server.

4. Click **Next**.

5. Under **Create/Backup Key**, enter information in the following field:

- Define a pass phrase/Re-enter pass phrase — Enter and confirm a pass phrase. The value can have between 8 and 32 characters, and must include each of the following:
 - An uppercase letter (one or more). Keep in mind that the pass phrase is case sensitive.
 - A number (one or more).
 - A non-alphanumeric character, such as !, *, @ (one or more).



Be sure to record your entries for later use. If you need to move a secure-enabled drive from the storage array, you must know the pass phrase to unlock drive data.

6. Click **Finish**.

The system connects to the key management server with the credentials you entered. A copy of the security key is then stored on your local system.



The path for the downloaded file might depend on the default download location of your browser.

7. Record your pass phrase and the location of the downloaded key file, and then click **Close**.

The page displays the following message with additional links for external key management:

Current key management method: External

8. Test the connection between the storage array and the key management server by selecting **Test Communication**.

Test results display in the dialog box.

Results

When external key management is enabled, you can create secure-enabled volume groups or pools, or you can enable security on existing volume groups and pools.



Whenever power to the drives is turned off and then on again, all the secure-enabled drives change to a Security Locked state. In this state, the data is inaccessible until the controller applies the correct security key during drive initialization. If someone physically removes a locked drive and installs it in another system, the Security Locked state prevents unauthorized access to its data.

After you finish

- You should validate the security key to make sure the key file is not corrupted.

Change security key

At any time, you can replace a security key with a new key. You might need to change a

security key in cases where you have a potential security breach at your company and want to make sure unauthorized personnel cannot access the drives' data.

Before you begin

A security key already exists.

About this task

This task describes how to change a security key and replace it with a new one. After this process, the old key is invalidated.

Steps

1. Select **Settings > System**.
2. Under **Security key management**, select **Change Key**.

The **Change Security Key** dialog box opens.

3. Enter information in the following fields.

- Define a security key identifier — (For internal security keys only.) Accept the default value (storage array name and time stamp, which is generated by the controller firmware) or enter your own value. You can enter up to 189 alphanumeric characters without spaces, punctuation, or symbols.



Additional characters are generated automatically and are appended to both ends of the string you enter. The generated characters help to ensure that the identifier is unique.

- Define a pass phrase/Re-enter pass phrase — In each of these fields, enter your pass phrase. The value can have between 8 and 32 characters, and must include each of the following:
 - An uppercase letter (one or more). Keep in mind that the pass phrase is case sensitive.
 - A number (one or more).
 - A non-alphanumeric character, such as !, *, @ (one or more).



Be sure to record your entries for later use — If you need to move a secure-enabled drive from the storage array, you must know the identifier and pass phrase to unlock drive data.

4. Click **Change**.

The new security key overwrites the previous key, which is no longer valid.



The path for the downloaded file might depend on the default download location of your browser.

5. Record your key identifier, pass phrase, and the location of the downloaded key file, and then click **Close**.

After you finish

You should validate the security key to make sure the key file is not corrupted.

Switch from external to internal key management

You can change the management method for Drive Security from an external key server to the internal method used by the storage array. The security key previously defined for

external key management is then used for internal key management.

Before you begin

An external key was created.

About this task

In this task, you disable external key management and download a new backup copy to your local host. The existing key is still used for Drive Security, but will be managed internally in the storage array.

Steps

1. Select **Settings > System**.
2. Under **Security key management**, select **Disable External Key Management**.

The **Disable External Key Management** dialog box opens.

3. In **Define a pass phrase/Re-enter pass phrase**, enter and confirm a pass phrase for the backup of the key. The value can have between 8 and 32 characters, and must include each of the following:
 - An uppercase letter (one or more). Keep in mind that the pass phrase is case sensitive.
 - A number (one or more).
 - A non-alphanumeric character, such as !, *, @ (one or more).



Be sure to record your entries for later use. If you need to move a secure-enabled drive from the storage array, you must know the identifier and pass phrase to unlock drive data.

4. Click **Disable**.

The backup key is downloaded to your local host.

5. Record your key identifier, pass phrase, and the location of the downloaded key file, and then click **Close**.

Results

Drive Security is now managed internally through the storage array.

After you finish

- You should validate the security key to make sure the key file is not corrupted.

Edit key management server settings

If you configured external key management, you can view and edit the key management server settings at any time.

Before you begin

External key management must be configured.

Steps

1. Select **Settings > System**.
2. Under **Security key management**, select **View/Edit Key Management Server Settings**.
3. Edit information in the following fields:
 - Key management server address — Enter the fully qualified domain name or the IP address (IPv4 or

IPv6) of the server used for key management.

- KMIP port number — Enter the port number used for the Key Management Interoperability Protocol (KMIP) communications.

4. Click **Save**.

Back up security key

After creating or changing a security key, you can create a backup copy of the key file in case the original gets corrupted.

Before you begin

- A security key already exists.

About this task

This task describes how to back up a security key you previously created. During this procedure, you create a new pass phrase for the backup. This pass phrase does not need to match the pass phrase that was used when the original key was created or last changed. The pass phrase is applied only to the backup you are creating.

Steps

1. Select **Settings > System**.
2. Under **Security key management**, select **Back Up Key**.

The **Back Up Security Key** dialog box opens.

3. In the **Define a pass phrase/Re-enter pass phrase** fields, enter and confirm a pass phrase for this backup.

The value can have between 8 and 32 characters, and must include each of the following:

- An uppercase letter (one or more)
- A number (one or more)
- A non-alphanumeric character, such as !, *, @ (one or more)



Be sure to record your entry for later use. You need the pass phrase to access the backup of this security key.

4. Click **Back Up**.

A backup of the security key is downloaded to your local host, and then the **Confirm/Record Security Key Backup** dialog box opens.



The path for the downloaded security key file might depend on the default download location of your browser.

5. Record your pass phrase in a secure location, and then click **Close**.

After you finish

You should validate the backup security key.

Validate security key

You can validate the security key to make sure it has not been corrupted and to verify that you have a correct pass phrase.

Before you begin

A security key has been created.

About this task

This task describes how to validate the security key you previously created. This is an important step to make sure the key file is not corrupted and the pass phrase is correct, which ensures that you can later access drive data if you move a secure-enabled drive from one storage array to another.

Steps

1. Select **Settings > System**.
2. Under **Security key management**, select **Validate Key**.

The **Validate Security Key** dialog box opens.

3. Click **Browse**, and then select the key file (for example, `drivesecurity.slk`).
4. Enter the pass phrase associated with the key you selected.

When you select a valid key file and pass phrase, the **Validate** button becomes available.

5. Click **Validate**.

The results of the validation are displayed in the dialog box.

6. If the results show "The security key validated successfully," click **Close**. If an error message appears, follow the suggested instructions displayed in the dialog box.

Unlock drives using a security key

If you move secure-enabled drives from one storage array to another, you must import the appropriate security key to the new storage array. Importing the key unlocks the data on the drives.

Before you begin

- The target storage array (where you are moving the drives) must already have a security key configured. The migrated drives will be re-keyed to the target storage array.
- You must know the security key that is associated with the drives you want to unlock.
- The security key file is available on the management client (the system with a browser used for accessing System Manager). If you are moving the drives to a storage array that is managed by a different system, you need to move the security key file to that management client.

About this task

This task describes how to unlock data in secure-enabled drives that have been removed from a storage array and reinstalled in another. Once the array discovers the drives, a "Needs Attention" condition appears along with a status of "Security Key Needed" for these re-located drives. You can unlock drive data by importing their security key into the storage array. During this process, you select the security key file and enter the pass phrase for the key.



The pass phrase is not the same as the storage array's Administrator password.

If other secure-enabled drives are installed in the new storage array, they might use a different security key than the one you are importing. During the import process, the old security key is used only to unlock the data for the drives you are installing. When the unlock process is successful, the newly installed drives are re-keyed to the target storage array's security key.

Steps

1. Select **Settings > System**.
2. Under **Security key management**, select **Unlock Secure Drives**.

The **Unlock Secure Drives** dialog box opens. Any drives that require a security key are shown in the table.

3. Optionally, hover the mouse over a drive number to see the location of the drive (shelf number and bay number).
4. Click **Browse**, and then select the security key file that corresponds to the drive you want to unlock.

The key file you selected appears in the dialog box.

5. Enter the pass phrase associated with this key file.

The characters you enter are masked.

6. Click **Unlock**.

If the unlock operation is successful, the dialog box displays: "The associated secure drives have been unlocked."

Results

When all drives are locked and then unlocked, each controller in the storage array will reboot. However, if there are already some unlocked drives in the target storage array, then the controllers will not reboot.

FAQs

What do I need to know before creating a security key?

A security key is shared by controllers and secure-enabled drives within a storage array. If a secure-enabled drive is removed from the storage array, the security key protects the data from unauthorized access.

You can create and manage security keys using one of the following methods:

- Internal key management on the controller's persistent memory.
- External key management on an external key management server.

Before creating an internal security key, you must do the following:

1. Install secure-capable drives in the storage array. These drives can be Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.
2. Make sure the Drive Security feature is enabled. If necessary, contact your storage vendor for instructions on enabling the Drive Security feature.

You can then create an internal security key, which involves defining an identifier and a pass phrase. The identifier is a string that is associated with the security key, and is stored on the controller and on all drives associated with the key. The pass phrase is used to encrypt the security key for backup purposes. When you are finished, the security key is stored on the controller in a non-accessible location. You can then create secure-enabled volume groups or pools, or you can enable security on existing volume groups and pools.

Before creating an external security key, you must do the following:

1. Install secure-capable drives in the storage array. These drives can be Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.
2. Make sure the Drive Security feature is enabled. If necessary, contact your storage vendor for instructions on enabling the Drive Security feature.
3. Complete and download a client Certificate Signing Request (CSR) for authentication between the storage array and the key management server. Go to **Settings > Certificates > Key Management > Complete CSR**.
4. Create and download a client certificate from the key management server using the downloaded CSR file.
5. Ensure that the client certificate and a copy of the certificate for the key management server are available on your local host.

You can then create an external key, which involves defining the IP address of the key management server and the port number used for KMIP communications. During this process, you also load certificate files. When you are finished, the system connects to the key management server with the credentials you entered. You can then create secure-enabled volume groups or pools, or you can enable security on existing volume groups and pools.

Why do I need to define a pass phrase?

The pass phrase is used to encrypt and decrypt the security key file stored on the local management client. Without the pass phrase, the security key cannot be decrypted and used to unlock data from a secure-enabled drive if it is re-installed in another storage array.

Why is it important to record security key information?

If you lose the security key information and do not have a backup, you could lose data when relocating secure-enabled drives or upgrading a controller. You need the security key to unlock data on the drives.

Be sure to record the security key identifier, the associated pass phrase, and the location on the local host where the security key file was saved.

What do I need to know before backing up a security key?

If your original security key becomes corrupted and you do not have a backup, you will lose access to the data on drives if they are migrated from one storage array to another.

Before backing up a security key, keep these guidelines in mind:

- Make sure you know the security key identifier and pass phrase of the original key file.



Only internal keys use identifiers. When you created the identifier, additional characters were generated automatically and appended to both ends of the identifier string. The generated characters ensure that the identifier is unique.

- You create a new pass phrase for the backup. This pass phrase does not need to match the pass phrase that was used when the original key was created or last changed. The pass phrase is only applied to the backup you are creating.



The pass phrase for Drive Security should not be confused with the storage array's Administrator password. The pass phrase for Drive Security protects backups of a security key. The Administrator password protects the entire storage array from unauthorized access.

- The backup security key file is downloaded to your management client. The path for the downloaded file might depend on the default download location of your browser. Be sure to make a record of where your security key information is stored.

What do I need to know before unlocking secure drives?

To unlock the data from a secure-enabled drive that is migrated to a new storage array, you must import its security key.

Before unlocking secure-enabled drives, keep the following guidelines in mind:

- The target storage array (where you are moving the drives) must already have a security key. The migrated drives will be re-keyed to the target storage array.
- For the drives you are migrating, you know the security key identifier and the pass phrase that corresponds to the security key file.
- The security key file is available on the management client (the system with a browser used for accessing System Manager).

What is read/write accessibility?

The **Drive Settings** window includes information about the **Drive Security** attributes. "Read/Write Accessible" is one of the attributes that displays if a drive's data has been locked.

To view **Drive Security** attributes, go to the Hardware page. Select a drive, click **View settings**, and then click **Show more settings**. At the bottom of the page, the Read/Write Accessible attribute value is **Yes** when the drive is unlocked. The Read/Write Accessible attribute value is **No, invalid security key** when the drive is locked. You can unlock a secure drive by importing a security key (go to **Settings > System > Unlock Secure Drives**).

What do I need to know about validating the security key?

After creating a security key, you should validate the key file to make sure it is not corrupt.

If the validation fails, do the following:

- If the security key identifier does not match the identifier on the controller, locate the correct security key file and then try the validation again.

- If the controller cannot decrypt the security key for validation, you might have incorrectly entered the pass phrase. Double-check the pass phrase, re-enter it if necessary, and then try the validation again. If the error message appears again, select a backup of the key file (if available) and re-try validation.
- If you still cannot validate the security key, the original file might be corrupted. Create a new backup of the key and validate that copy.

What is the difference between internal security key and external security key management?

When you implement the **Drive Security** feature, you can use an internal security key or an external security key to lock down data when a secure-enabled drive is removed from the storage array.

A security key is a string of characters, which is shared between the secure-enabled drives and controllers in a storage array. Internal keys are maintained on the controller's persistent memory. External keys are maintained on a separate key management server, using a Key Management Interoperability Protocol (KMIP).

Access management

Concepts

How Access Management works

Access Management is a method of establishing user authentication in SANtricity System Manager. Authentication requires users to log in to these systems with their assigned credentials.

Access Management configuration and user authentication works as follows:

1. An administrator logs in to System Manager with a user profile that includes Security Admin permissions.



For first-time login, the username `admin` is automatically displayed and cannot be changed. The `admin` user has full access to all functions in the system.

2. The administrator navigates to Access Management in the user interface. The storage array is pre-configured to use local user roles, which are an implementation of RBAC (role-based access control) capabilities.
3. The administrator configures one or more of the following authentication methods:
 - **Local user roles** — Authentication is managed through RBAC capabilities enforced in the storage array. Local user roles include pre-defined user profiles and roles with specific access permissions. Administrators can use these local user roles as the single method of authentication, or use them in combination with a directory service. No configuration is necessary, other than setting passwords for users.
 - **Directory services** — Authentication is managed through an LDAP (Lightweight Directory Access Protocol) server and directory service, such as Microsoft's Active Directory. An administrator connects to the LDAP server, and then maps the LDAP users to the local user roles embedded in the storage array.
 - **SAML** — Authentication is managed through an Identity Provider (IdP) using the Security Assertion Markup Language (SAML) 2.0. An administrator establishes communication between the IdP system and the storage array, and then maps IdP users to the local user roles embedded in the storage array.

4. The administrator provides users with login credentials for System Manager.
5. Users log in to the system by entering their credentials.



If authentication is managed with SAML and an SSO (single sign-on), the system might bypass the System Manager login dialog.

During login, the system performs the following background tasks:

- Authenticates the user name and password against the user account.
- Determines the user's permissions based on the assigned roles.
- Provides the user with access to tasks in the user interface.
- Displays the user name in the upper right of the interface.

Tasks available in System Manager

Access to tasks depends on a user's assigned roles, which include the following:

- **Storage admin** — Full read/write access to the storage objects (for example, volumes and disk pools), but no access to the security configuration.
- **Security admin** — Access to the security configuration in Access Management, certificate management, audit log management, and the ability to turn the legacy management interface (SYMbol) on or off.
- **Support admin** — Access to all hardware resources on the storage array, failure data, MEL events, and controller firmware upgrades. No access to storage objects or the security configuration.
- **Monitor** — Read-only access to all storage objects, but no access to the security configuration.

An unavailable task is either grayed out or does not display in the user interface. For example, a user with the Monitor role can view all information about volumes, but cannot access functions for modifying that volume. The tabs for features such as **Copy Services** and **Add to Workload** will be grayed out; only View/Edit Settings is available.

User access to SANtricity Storage Manager

When local user roles and directory services are configured, users must enter credentials before performing any of the following functions in the Enterprise Management Window (EMW):

- Renaming the storage array
- Upgrading controller firmware
- Loading a storage array configuration
- Executing a script
- Attempting to perform an active operation when an unused session has timed out

If SAML is configured for a storage array, users cannot use the EMW to discover or manage storage for that array.

Access Management terminology

Learn how the Access Management terms apply to your storage array.

Term	Description
Active Directory	Active Directory (AD) is a Microsoft directory service that uses LDAP for Windows domain networks.
Binding	Bind operations are used to authenticate clients to the directory server. Binding usually requires account and password credentials, but some servers allow for anonymous bind operations.
CA	A certificate authority (CA) is a trusted entity that issues electronic documents, called digital certificates, for Internet security. These certificates identify website owners, which allows for secure connections between clients and servers.
Certificate	A certificate identifies the owner of a site for security purposes, which prevents attackers from impersonating the site. The certificate contains information about the site owner and the identity of the trusted entity who certifies (signs) this information.
IdP	An Identity Provider (IdP) is an external system used to request credentials from a user and to determine if that user is successfully authenticated. The IdP can be configured to provide multi-factor authentication and to use any user database, such as Active Directory. Your security team is responsible for maintaining the IdP.
LDAP	Lightweight Directory Access Protocol (LDAP) is an application protocol for accessing and maintaining distributed directory information services. This protocol allows many different applications and services to connect to the LDAP server for validating users.
RBAC	Role-based access control (RBAC) is a method of regulating access to computer or network resources based on the roles of individual users. RBAC controls are enforced on the storage array and include predefined roles.

Term	Description
SAML	Security Assertion Markup Language (SAML) is an XML-based standard for authentication and authorization between two entities. SAML allows for multi-factor authentication, in which users must provide two or more items for proving their identity (for example, a password and fingerprint). The storage array's embedded SAML feature is SAML2.0 compliant for identity assertion, authentication, and authorization.
SP	A Service Provider (SP) is a system that controls user authentication and access. When Access Management is configured with SAML, the storage array acts as the Service Provider for requesting authentication from the Identity Provider.
SSO	Single sign-on (SSO) is an authentication service that allows for one set of login credentials to access multiple applications.

Permissions for mapped roles

The RBAC (role-based access control) capabilities enforced on the storage array include pre-defined user profiles with one or more roles mapped to them. Each role includes permissions for accessing tasks in SANtricity System Manager.

User profiles and mapped roles are accessible from **Settings > Access Management > Local User Roles** in the user interface of either System Manager.

The roles provide user access to tasks, as follows:

- **Storage admin** — Full read/write access to the storage objects (for example, volumes and disk pools), but no access to the security configuration.
- **Security admin** — Access to the security configuration in Access Management, certificate management, audit log management, and the ability to turn the legacy management interface (SYMBOL) on or off.
- **Support admin** — Access to all hardware resources on the storage array, failure data, MEL events, and controller firmware upgrades. No access to storage objects or the security configuration.
- **Monitor** — Read-only access to all storage objects, but no access to the security configuration.

If a user does not have permissions for a certain task, that task is either grayed out or does not display in the user interface.

Access Management with local user roles

For Access Management, administrators can use RBAC (role-based access control) capabilities enforced in the storage array. These capabilities are referred to as "local user roles."

Configuration workflow

Local user roles are pre-configured for the storage array. To use local user roles for authentication, administrators can do the following:

1. An administrator logs in to SANtricity System Manager with a user profile that includes Security Admin permissions.



The `admin` user has full access to all functions in the system.

2. An administrator reviews the user profiles, which are predefined and cannot be modified.
3. Optionally, the administrator assigns new passwords for each user profile.
4. Users log in to the system with their assigned credentials.

Management

When using only local user roles for authentication, administrators can perform the following management tasks:

- Change passwords.
- Set a minimum length for passwords.
- Allow users to log in without passwords.

Access Management with directory services

For Access Management, administrators can use an LDAP (Lightweight Directory Access Protocol) server and a directory service, such as Microsoft's Active Directory.

Configuration workflow

If an LDAP server and directory service are used in the network, configuration works as follows:

1. An administrator logs in to SANtricity System Manager with a user profile that includes Security Admin permissions.



The `admin` user has full access to all functions in the system.

2. The administrator enters the configuration settings for the LDAP server. Settings include the domain name, URL, and Bind account information.
3. If the LDAP server uses a secure protocol (LDAPS), the administrator uploads a Certificate Authority (CA) certificate chain for authentication between the LDAP server and the storage array.
4. After the server connection is established, the administrator maps the user groups to the storage array's roles. These roles are predefined and cannot be modified.
5. The administrator tests the connection between the LDAP server and the storage array.
6. Users log in to the system with their assigned LDAP/Directory Services credentials.

Management

When using directory services for authentication, administrators can perform the following management tasks:

- Add a directory server.
- Edit directory server settings.
- Map LDAP users to local user roles.
- Remove a directory server.

Access Management with SAML

For Access Management, administrators can use the Security Assertion Markup Language (SAML) 2.0 capabilities embedded in the array.

Configuration workflow

SAML configuration works as follows:

1. An administrator logs in to System Manager with a user profile that includes Security Admin permissions.



The `admin` user has full access to all functions in System Manager.

2. The administrator goes to the **SAML** tab under Access Management.
3. An administrator configures communications with the Identity Provider (IdP). An IdP is an external system used to request credentials from a user and determine if the user is successfully authenticated. To configure communications with the storage array, the administrator downloads the IdP metadata file from the IdP system, and then uses System Manager to upload the file to the storage array.
4. An administrator establishes a trust relationship between the Service Provider and the IdP. A Service Provider controls user authorization; in this case, the controller in the storage array acts as the Service Provider. To configure communications, the administrator uses System Manager to export a Service Provider metadata file for each controller. From the IdP system, the administrator then imports those metadata files to the IdP.



Administrators should also make sure that the IdP supports the ability to return a Name ID on authentication.

5. The administrator maps the storage array's roles to user attributes defined in the IdP. To do this, the administrator uses System Manager to create the mappings.
6. The administrator tests the SSO login to the IdP URL. This test ensures the storage array and IdP can communicate.



Once SAML is enabled, you *cannot* disable it through the user interface, nor can you edit the IdP settings. If you need to disable or edit the SAML configuration, contact Technical Support for assistance.

7. From System Manager, the administrator enables SAML for the storage array.
8. Users log in to the system with their SSO credentials.

Management

When using SAML for authentication, administrators can perform the following management tasks:

- Modify or create new role mappings

- Export Service Provider files

Access restrictions

When SAML is enabled, the following clients cannot access storage array services and resources:

- Enterprise Management Window (EMW)
- Command-line interface (CLI)
- Software Developer Kits (SDK) clients
- In-band clients
- HTTP Basic Authentication REST API clients
- Login using standard REST API endpoint

How tos

View local user roles

From the Local User Roles tab, you can view the mappings of the user profiles to the default roles. These mappings are part of the RBAC (role-based access controls) enforced in the storage array.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.

About this task

The user profiles and mappings cannot be changed. Only passwords can be modified.

Steps

1. Select **Settings** > **Access Management**.
2. Select the **Local User Roles** tab.

The user profiles are shown in the table:

- **Root admin** (admin) — Super administrator who has access to all functions in the system. This user profile includes all roles.
- **Storage admin** (storage) — The administrator responsible for all storage provisioning. This user profile includes the following roles: Storage Admin, Support Admin, and Monitor.
- **Security admin** (security) — The user responsible for security configuration, including access management, certificate management, and secure-enabled drive functions. This user profile includes the following roles: Security Admin and Monitor.
- **Support admin** (support) — The user responsible for hardware resources, failure data, and firmware upgrades. This user profile includes the following roles: Support Admin and Monitor.
- **Monitor** (monitor) — A user with read-only access to the system. This user profile includes only the Monitor role.

Change passwords

You can change the user passwords for each user profile in Access Management.

Before you begin

- You must be logged in as the local administrator, which includes Root admin permissions.
- You must know the local administrator password.

About this task

Keep these guidelines in mind when choosing a password:

- Any new local user passwords must meet or exceed the current setting for a minimum password (in View/Edit Settings).
- Passwords are case sensitive.
- Trailing spaces are not stripped from passwords when they are set. Be careful to include spaces if they were included in the password.
- For increased security, use at least 15 alphanumeric characters and change the password frequently.



Changing the password in System Manager also changes it in the command line interface (CLI). In addition, password changes cause the user's active session to terminate.

Steps

1. Select **Settings > Access Management**.
2. Select the **Local User Roles** tab.
3. Select a user from the table.

The **Change Password** button becomes available.

4. Select **Change Password**.

The **Change Password** dialog box opens.

5. If no minimum password length is set for local user passwords, you can check the box to require the selected user to enter a password to access the storage array, and then you can type the new password for the selected user.
6. Enter your local administrator password, and then click **Change**.

Result

If the user is currently logged in, the password change causes the user's active session to terminate.

Change local user password settings

You can set the minimum required length for all new or updated local user passwords on the storage array. You can also allow local users to access the storage array without entering a password.

Before you begin

- You must be logged in as the local administrator, which includes Root admin permissions.

About this task

Keep these guidelines in mind when setting the minimum length for local user passwords:

- Setting changes will not affect existing local user passwords.
- The minimum required length setting for local user passwords must be between 0 and 30 characters.
- Any new local user passwords must meet or exceed the current minimum length setting.
- Do not set a minimum length for the password if you want local users to access the storage array without entering a password.

Steps

1. Select **Settings** > **Access Management**.
2. Select the **Local User Roles** tab.
3. Select the **View/Edit Settings** button.

The **Local User Password Settings** dialog box opens.

4. Do one of the following:
 - To allow local users to access the storage array *without* entering a password, uncheck the "Require all local user passwords to be at least" checkbox.
 - To set a minimum password length for all local user passwords, check the "Require all local user passwords to be at least" checkbox and then use the spinner box to set the minimum required length for all local user passwords.

Any new local user passwords must meet or exceed the current setting.

5. Click **Save**.

Add directory server

To configure authentication for Access Management, you can establish communications between the storage array and an LDAP server, and then map the LDAP user groups to the array's predefined roles.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.
- User groups must be defined in your directory service.
- LDAP server credentials must be available, including the domain name, server URL, and optionally the bind account user name and password.
- For LDAPS servers using a secure protocol, the LDAP server's certificate chain must be installed on your local machine.

About this task

Adding a directory server is a two-step process. First you enter the domain name and URL. If your server uses a secure protocol, you must also upload a CA certificate for authentication if it is signed by a non-standard signing authority. If you have credentials for a bind account, you can also enter your user account name and password. Next, you map the LDAP server's user groups to the storage array's predefined roles.





During the procedure to add an LDAP server, the legacy management interface will be disabled. The legacy management interface (SYMbol) is a method of communication between the storage array and the management client. When disabled, the storage array and management client use a more secure method of communication (REST API over https).

Steps

1. Select **Settings** > **Access Management**.
2. From the **Directory Services** tab, select **Add Directory Server**.

The **Add Directory Server** dialog box opens.
3. In the **Server Settings** tab, enter the credentials for the LDAP server.

Field Details

Setting	Description
Configuration settings	
Domain(s)	Enter the domain name of the LDAP server. For multiple domains, enter the domains in a comma separated list. The domain name is used in the login (<i>username@domain</i>) to specify which directory server to authenticate against.
Server URL	Enter the URL for accessing the LDAP server in the form of <code>ldap[s]://host:port</code> .
Upload certificate (optional)	<div>  <p>This field appears only if an LDAPS protocol is specified in the Server URL field above.</p> </div> <p>Click Browse and select a CA certificate to upload. This is the trusted certificate or certificate chain used for authenticating the LDAP server.</p>
Bind account (optional)	Enter a read-only user account for search queries against the LDAP server and for searching within the groups. Enter the account name in an LDAP-type format. For example, if the bind user is called "bindacct," then you might enter a value such as "CN=bindacct,CN=Users,DC=cpoc,DC=local."
Bind password (optional)	<div>  <p>This field appears when you enter a bind account above.</p> </div> <p>Enter the password for the bind account.</p>
Test server connection before adding	Select this checkbox if you want to make sure the storage array can communicate with the LDAP server configuration you entered. The test occurs after you click Add at the bottom of the dialog box. If this checkbox is selected and the test fails, the configuration is not added. You must resolve the error or de-select the checkbox to skip the testing and add the configuration.
Privilege settings	

Setting	Description
Search base DN	Enter the LDAP context to search for users, typically in the form of <code>CN=Users, DC=copc, DC=local</code> .
Username attribute	Enter the attribute that is bound to the user ID for authentication. For example: <code>sAMAccountName</code> .
Group attribute(s)	Enter a list of group attributes on the user, which is used for group-to-role mapping. For example: <code>memberOf, managedObjects</code> .

4. Click the **Role Mapping** tab.
5. Assign LDAP groups to the predefined roles. A group can have multiple assigned roles.

Field Details

Setting	Description
Mappings	
Group DN	Specify the group distinguished name (DN) for the LDAP user group to be mapped.
Roles	<p>Click in the field and select one of the storage array's roles to be mapped to the Group DN. You must individually select each role you want to include for this group. The Monitor role is required in combination with the other roles to log in to SANtricity System Manager.</p> <p>The mapped roles include the following permissions:</p> <ul style="list-style-type: none"> • Storage admin — Full read/write access to the storage objects (for example, volumes and disk pools), but no access to the security configuration. • Security admin — Access to the security configuration in Access Management, certificate management, audit log management, and the ability to turn the legacy management interface (SYMBOL) on or off. • Support admin — Access to all hardware resources on the storage array, failure data, MEL events, and controller firmware upgrades. No access to storage objects or the security configuration. • Monitor — Read-only access to all storage objects, but no access to the security configuration.



The Monitor role is required for all users, including the administrator. System Manager will not operate correctly for any user without the Monitor role present.

6. If desired, click **Add another mapping** to enter more group-to-role mappings.
7. When you are finished with the mappings, click **Add**.

The system performs a validation, making sure that the storage array and LDAP server can communicate. If an error message appears, check the credentials entered in the dialog box and re-enter the information if necessary.

Edit directory server settings and role mappings

If you previously configured a directory server in Access Management, you can change its settings at any time. Settings include the server connection information and the group-to-role mappings.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.
- A directory server must be defined.

Steps

1. Select **Settings > Access Management**.
2. Select the **Directory Services** tab.
3. If more than one server is defined, select the server you want to edit from the table.
4. Select **View/Edit Settings**.

The **Directory Server Settings** dialog box opens.

5. In the **Server Settings** tab, change the desired settings.

Setting	Description
Configuration settings	
Domain(s)	The domain name(s) of the LDAP server(s). For multiple domains, enter the domains in a comma separated list. The domain name is used in the login (<i>username@domain</i>) to specify which directory server to authenticate against.
Server URL	The URL for accessing the LDAP server in the form of <code>ldap[s]://host:port</code> .
Bind account (optional)	The read-only user account for search queries against the LDAP server and for searching within the groups.
Bind password (optional)	The password for the bind account. (This field appears when a bind account is entered.)
Test server connection before saving	Checks that the storage array can communicate with the LDAP server configuration. The test occurs after you click Save at the bottom of the dialog box. If this checkbox is selected and the test fails, the configuration is not changed. You must resolve the error or de-select the checkbox to skip the testing and re-edit the configuration.

Setting	Description
Privilege settings	
Search base DN	The LDAP context to search for users, typically in the form of CN=Users, DC=copc, DC=local.
Username attribute	The attribute that is bound to the user ID for authentication. For example: sAMAccountName.
Group attribute(s)	A list of group attributes on the user, which is used for group-to-role mapping. For example: memberOf, managedObjects.

6. In the **Role Mapping** tab, change the desired mapping.

Setting	Description
Mappings	
Group DN	The domain name for the LDAP user group to be mapped.
Roles	<p>The storage array's roles to be mapped to the Group DN. You must individually select each role you want to include for this group. The Monitor role is required in combination with the other roles to log in to SANtricity System Manager.</p> <p>The storage array's roles include the following:</p> <ul style="list-style-type: none"> • Storage admin — Full read/write access to the storage objects (for example, volumes and disk pools), but no access to the security configuration. • Security admin — Access to the security configuration in Access Management, certificate management, audit log management, and the ability to turn the legacy management interface (SYMBOL) on or off. • Support admin — Access to all hardware resources on the storage array, failure data, MEL events, and controller firmware upgrades. No access to storage objects or the security configuration. • Monitor — Read-only access to all storage objects, but no access to the security configuration.



The Monitor role is required for all users, including the administrator. System Manager will not operate correctly for any user without the Monitor role present.

7. If desired, click **Add another mapping** to enter more group-to-role mappings.
8. Click **Save**.

Result

After you complete this task, any active user sessions are terminated. Only your current user session is retained.

Remove directory server

To break the connection between a directory server and the storage array, you can remove the server information from the Access Management page. You might want to perform this task if you configured a new server, and then want to remove the old one.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.

About this task

After you complete this task, any active user sessions are terminated. Only your current user session is retained.

Steps

1. Select **Settings > Access Management**.
2. Select the **Directory Services** tab.
3. From the list, select the directory server you want to delete.
4. Click **Remove**.

The **Remove Directory Server** dialog box opens.

5. Type `remove` in the field, and then click **Remove**.

The directory server configuration settings, privilege settings, and role mappings are removed. Users can no longer log in with credentials from this server.

Configure SAML

To configure authentication for Access Management, you can use the Security Assertion Markup Language (SAML) capabilities embedded in the storage array. This configuration establishes a connection between an Identity Provider and the Storage Provider.

About this task

An Identity Provider (IdP) is an external system used to request credentials from a user and to determine if that user is successfully authenticated. The IdP can be configured to provide multi-factor authentication and to use any user database, such as Active Directory. Your security team is responsible for maintaining the IdP. A Service Provider (SP) is a system that controls user authentication and access. When Access Management is configured with SAML, the storage array acts as the Service Provider for requesting authentication from the

Identity Provider. To establish a connection between the IdP and storage array, you share metadata files between these two entities. Next, you map the IdP user entities to the storage array roles. And finally, you test the connection and SSO logins before enabling SAML.



SAML and Directory Services. If you enable SAML when Directory Services is configured as the authentication method, SAML supersedes Directory Services in System Manager. If you disable SAML later, the Directory Services configuration returns to its previous configuration.



Editing and Disabling. Once SAML is enabled, you *cannot* disable it through the user interface, nor can you edit the IdP settings. If you need to disable or edit the SAML configuration, contact Technical Support for assistance.

Configuring SAML authentication is a multi-step procedure:

- [Step 1: Upload the IdP metadata file](#)
- [Step 2: Export Service Provider files](#)
- [Step 3: Map roles](#)
- [Step 4: Test SSO login](#)
- [Step 5: Enable SAML](#)

Step 1: Upload the IdP metadata file

To provide the storage array with IdP connection information, you import IdP metadata into System Manager.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.
- An IdP administrator has configured an IdP system.
- An IdP administrator has ensured that the IdP supports the ability to return a Name ID on authentication.
- An administrator has ensured that the IdP server and controller clocks are synchronized (either through an NTP server or by adjusting the controller clock settings).
- An IdP metadata file is downloaded from the IdP system and is available on the local system used for accessing System Manager.

About this task

In this task, you upload a metadata file from the IdP into System Manager. The IdP system needs this metadata to redirect authentication requests to the correct URL and to validate responses received. You only need to upload one metadata file for the storage array, even if there are two controllers.

Steps

1. Select **Settings** > **Access Management**.
2. Select the **SAML** tab.

The page displays an overview of configuration steps.

3. Click the **Import Identity Provider (IdP) file** link.

The **Import Identity Provider File** dialog opens.

4. Click **Browse** to select and upload the IdP metadata file you copied to your local system.

After you select the file, the IdP Entity ID is displayed.

5. Click **Import**.

Step 2: Export Service Provider files

To establish a trust relationship between the IdP and the storage array, you import the Service Provider metadata into the IdP.

Before you begin

- You know the IP address or domain name of each controller in the storage array.

About this task

In this task, you export metadata from the controllers (one file for each controller). The IdP needs this metadata to establish a trust relationship with the controllers and to process authorization requests. The file includes information such as the controller domain name or IP address, so that the IdP can communicate with the Service Providers.

Steps

1. Click the **Export Service Provider files** link.

The **Export Service Provider Files** dialog opens.

2. Enter the controller IP address or DNS name in the **Controller A** field, and then click **Export** to save the metadata file to your local system. If the storage array includes two controllers, repeat this step for the second controller in the **Controller B** field.

After you click Export, the Service Provider metadata is downloaded to your local system. Make a note of where the file is stored.

3. From the local system, locate the Service Provider metadata file(s) you exported.

There is one XML-formatted file for each controller.

4. From the IdP server, import the Service Provider metadata file(s) to establish the trust relationship. You can either import the files directly or you can manually enter the controller information from the files.

Step 3: Map roles

To provide users with authorization and access to System Manager, you must map the IdP user attributes and group memberships to the storage array's predefined roles.

Before you begin

- An IdP administrator has configured user attributes and group membership in the IdP system.
- The IdP metadata file is imported into System Manager.
- A Service Provider metadata file for each controller is imported into the IdP system for the trust relationship.

About this task

In this task, you use System Manager to map IdP groups to local user roles.

Steps

1. Click the link for mapping System Manager roles.

The **Role Mapping** dialog box opens.

2. Assign IdP user attributes and groups to the predefined roles. A group can have multiple assigned roles.

Field Details

Setting	Description
Mappings	
User Attribute	Specify the attribute (for example, "member of") for the SAML group to be mapped.
Attribute Value	Specify the attribute value for the group to be mapped.
Roles	<p>Click in the field and select one of the storage array's roles to be mapped to the Attribute. You must individually select each role you want to include. The Monitor role is required in combination with the other roles to log in to System Manager. The Security Admin role is also required for at least one group. The mapped roles include the following permissions:</p> <ul style="list-style-type: none">• Storage admin — Full read/write access to the storage objects (for example, volumes and disk pools), but no access to the security configuration.• Security admin — Access to the security configuration in Access Management, certificate management, audit log management, and the ability to turn the legacy management interface (SYMBOL) on or off.• Support admin — Access to all hardware resources on the storage array, failure data, MEL events, and controller firmware upgrades. No access to storage objects or the security configuration.• Monitor — Read-only access to all storage objects, but no access to the security configuration.



The Monitor role is required for all users, including the administrator. System Manager will not operate correctly for any user without the Monitor role present.

3. If desired, click **Add another mapping** to enter more group-to-role mappings.



Role mappings can be modified after SAML is enabled.

4. When you are finished with the mappings, click **Save**.

Step 4: Test SSO login

To ensure that the IdP system and storage array can communicate, you can optionally test an SSO login. This test is also performed during the final step for enabling SAML.

Before you begin

- The IdP metadata file is imported into System Manager.
- A Service Provider metadata file for each controller is imported into the IdP system for the trust relationship.

Steps

1. Select the **Test SSO Login** link.

A dialog opens for entering SSO credentials.

2. Enter login credentials for a user with both Security Admin permissions and Monitor permissions.

A dialog opens while the system tests the login.

3. Look for a Test Successful message. If the test completes successfully, go to the next step for enabling SAML.

If the test does not complete successfully, an error message appears with further information. Make sure that:

- The user belongs to a group with permissions for Security Admin and Monitor.
- The metadata you uploaded for the IdP server is correct.
- The controller addresses in the SP metadata files are correct.

Step 5: Enable SAML

Your final step is to enable SAML user authentication.

Before you begin

- The IdP metadata file is imported into System Manager.
- A Service Provider metadata file for each controller is imported into the IdP system for the trust relationship.
- At least one Monitor and one Security Admin role mapping is configured.

About this task

This task describes how to finish the SAML configuration for user authentication. During this process, the system also prompts you to test an SSO login. The SSO Login test process is described in the previous step.



Editing and Disabling. Once SAML is enabled, you *cannot* disable it through the user interface, nor can you edit the IdP settings. If you need to disable or edit the SAML configuration, contact Technical Support for assistance.

Steps

1. From the **SAML** tab, select the **Enable SAML** link.

The **Confirm Enable SAML** dialog opens.

2. Type `enable`, and then click **Enable**.
3. Enter user credentials for an SSO login test.

Result

After the system enables SAML, it terminates all active sessions and begins authenticating users through SAML.

Change SAML role mappings

If you previously configured SAML for Access Management, you can change the role mappings between the IdP groups and the storage array's predefined roles.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.
- An IdP administrator has configured user attributes and group membership in the IdP system.
- SAML is configured and enabled.

Steps

1. Select **Settings > Access Management**.
2. Select the **SAML** tab.
3. Select **Role Mapping**.

The **Role Mapping** dialog box opens.

4. Assign IdP user attributes and groups to the predefined roles. A group can have multiple assigned roles.



Be careful that you do not remove your permissions while SAML is enabled, or you will lose access to System Manager.

Field Details

Setting	Description
Mappings	
User Attribute	Specify the attribute (for example, "member of") for the SAML group to be mapped.
Attribute Value	Specify the attribute value for the group to be mapped.
Roles	<p>Click in the field and select one of the storage array's roles to be mapped to the attribute. You must individually select each role you want to include for this group. The Monitor role is required in combination with the other roles to log in to System Manager. A Security Admin role must be assigned to at least one group.</p> <p>The mapped roles include the following permissions:</p> <ul style="list-style-type: none">• Storage admin — Full read/write access to the storage objects (for example, volumes and disk pools), but no access to the security configuration.• Security admin — Access to the security configuration in Access Management, certificate management, audit log management, and the ability to turn the legacy management interface (SYMBOL) on or off.• Support admin — Access to all hardware resources on the storage array, failure data, MEL events, and controller firmware upgrades. No access to storage objects or the security configuration.• Monitor — Read-only access to all storage objects, but no access to the security configuration.



The Monitor role is required for all users, including the administrator. System Manager will not operate correctly for any user without the Monitor role present.

5. **Optionally:** click **Add another mapping** to enter more group-to-role mappings.

6. Click **Save**.

Result

After you complete this task, any active user sessions are terminated. Only your current user session is retained.

Export SAML Service Provider files

If necessary, you can export Service Provider metadata for the storage array and re-import the file(s) into the Identity Provider (IdP) system.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.
- SAML is configured and enabled.

About this task

In this task, you export metadata from the controllers (one file for each controller). The IdP needs this metadata to establish a trust relationship with the controllers and to process authentication requests. The file includes information such as the controller domain name or IP address that the IdP can use for sending requests.

Steps

1. Select **Settings** > **Access Management**.
2. Select the **SAML** tab.
3. Select **Export**.

The **Export Service Provider Files** dialog opens.

4. For each controller, click **Export** to save the metadata file to your local system.



The domain name fields for each controller are read-only.

Make a note of where the file is stored.

5. From the local system, locate the Service Provider metadata file(s) you exported.

There is one XML-formatted file for each controller.

6. From the IdP server, import the Service Provider metadata file(s). You can either import the files directly or you can manually enter the controller information from them.
7. Click **Close**.

View audit log activity

By viewing audit logs, users with Security Admin permissions can monitor user actions, authentication failures, invalid login attempts, and the user session lifespan.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.

Steps

1. Select **Settings** > **Access Management**.




2. Select the **Audit Log** tab.

Audit log activity appears in tabular format, which includes the following columns of information:

- **Date/Time** — Timestamp of when the storage array detected the event (in GMT).
- **Username** — The user name associated with the event. For any non-authenticated actions on the storage array, "N/A" appears as the user name. Non-authenticated actions might be triggered by the internal proxy or some other mechanism.
- **Status Code** — HTTP status code of the operation (200, 400, etc.) and descriptive text associated with the event.
- **URL Accessed** — Full URL (including host) and query string.
- **Client IP Address** — IP address of the client associated with the event.
- **Source** — Logging source associated with the event, which can be System Manager, CLI, Web Services, or Support Shell.

3. Use the selections on the Audit Log page to view and manage events.

Selection Details

Selection	Description
Show events from the...	Limit events shown by date range (last 24 hours, last 7 days, last 30 days, or a custom date range).
Filter	Limit events shown by the characters entered in the field. Use quotes (") for an exact word match, enter OR to return one or more words, or enter a dash (--) to omit words.
Refresh	Select Refresh to update the page to the most current events.
View/Edit Settings	Select View/Edit Settings to open a dialog box that allows you to specify a full log policy and level of actions to be logged.
Delete events	Select Delete to open a dialog box that allows you to remove old events from the page.
Show/hide columns	<p>Click the Show/Hide column icon  to select additional columns for display in the table. Additional columns include:</p> <ul style="list-style-type: none"> • Method — The HTTP method (for example, POST, GET, DELETE, etc.). • CLI Command Executed — The CLI command (grammar) executed for Secure CLI requests. • CLI Return Status — A CLI status code or a request for input files from the client. • SYMBOL Procedure — The SYMBOL procedure executed. • SSH Event Type — Secure Shell (SSH) events type, such as login, logout, and login_fail. • SSH Session PID — Process ID number of the SSH session. • SSH Session Duration(s) — The number of seconds the user was logged in.
Toggle column filters	Click the Toggle icon  to open filtering fields for each column. Enter characters within a column field to limit events shown by those characters. Click the icon again to close the filtering fields.
Undo changes	Click the Undo icon  to return the table to the default configuration.
Export	Click Export to save the table data to a comma separated value (CSV) file.

Define audit log policies

You can change the overwrite policy and the types of events recorded in the audit log.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.

About this task

This task describes how to change the audit log settings, which include the policy for overwriting old events and the policy for recording event types.


Steps

1. Select **Settings** > **Access Management**.
2. Select the **Audit Log** tab.
3. Select **View/Edit Settings**.

The **Audit Log Settings** dialog box opens.

4. Change the overwrite policy or types of events recorded.

Field Details

Setting	Description
Overwrite policy	<p>Determines the policy for overwriting old events when the maximum capacity is reached:</p> <ul style="list-style-type: none">• Allow the oldest events in the audit log to be overwritten when the audit log is full — Overwrites the old events when the audit log reaches 50,000 records.• Require audit log events to be manually deleted — Specifies that events will not be automatically deleted; instead, a threshold warning appears at the set percentage. Events must be deleted manually. <div><p>If the overwrite policy is disabled and the audit log entries reach the maximum limit, access to System Manager is denied to users without Security Admin permissions. To restore system access to users without Security Admin permissions, a user assigned to the Security Admin role must delete the old event records.</p></div> <div><p>Overwrite policies do not apply if a syslog server is configured for archiving audit logs.</p></div>
Level of actions to be logged	<p>Determines types of events to be logged:</p> <ul style="list-style-type: none">• Record modification events only — Shows only the events where a user action involves making a change in the system.• Record all modification and read-only events — Shows all events, including a user action that involves reading or downloading information.

5. Click **Save**.

Delete events from the audit log

You can clear the audit log of old events, which makes searching through events more

manageable. You have the option of saving old events to a CSV (comma-separated values) file upon deletion.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.

About this task

This task describes how to remove old events from the audit log.

Steps

1. Select **Settings** > **Access Management**.
2. Select the **Audit Log** tab.
3. Select **Delete**.

The **Delete Audit Log dialog** box opens.

4. Select or enter the number of oldest events that you want to delete.
5. If you want to export the deleted events to a CSV file (recommended), keep the checkbox selected. You will be prompted to enter a file name and location when you click **Delete** in the next step. Otherwise, if you do not want to save events to a CSV file, click the checkbox to deselect it.
6. Click **Delete**.

A confirmation dialog box opens.

7. Type `delete` in the field, and then click **Delete**.

The oldest events are removed from the Audit Log page.

Configure syslog server for audit logs

If you want to archive audit logs onto an external syslog server, you can configure communications between that server and the storage array. After the connection is established, audit logs are automatically saved to the syslog server.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.
- The syslog server address, protocol, and port number must be available. The server address can be a fully qualified domain name, an IPv4 address, or an IPv6 address.
- If your server uses a secure protocol (for example, TLS), a Certificate Authority (CA) certificate must be available on your local system. CA certificates identify website owners for secure connections between servers and clients.

Steps

1. Select **Settings** > **Access Management**.
2. From the **Audit Log** tab, select **Configure Syslog Servers**.

The **Configure Syslog Servers** dialog box opens.

3. Click **Add**.

The **Add Syslog Server** dialog box opens.

4. Enter information for the server, and then click **Add**.
 - Server address — Enter a fully qualified domain name, an IPv4 address, or an IPv6 address.
 - Protocol — Select a protocol from the drop-down list (for example, TLS, UDP, or TCP).
 - Upload certificate (optional) — If you selected the TLS protocol and have not yet uploaded a signed CA certificate, click Browse to upload a certificate file. Audit logs are not archived to a syslog server without a trusted certificate.



If the certificate becomes invalid later, the TLS handshake will fail. As a result, an error message is posted to the audit log and messages are no longer sent to the syslog server. To resolve this issue, you must fix the certificate on the syslog server and then go to **Settings > Audit Log > Configure Syslog Servers > Test All**.

- Port — Enter the port number for the syslog receiver. After you click **Add**, the **Configure Syslog Servers** dialog box opens and displays your configured syslog server on the page.

5. To test the server connection with the storage array, select **Test All**.

Result

After configuration, all new audit logs are sent to the syslog server. Previous logs are not transferred.

Edit syslog server settings for audit log records

You can change the settings for the syslog server used for archiving audit logs, and also upload a new Certificate Authority (CA) certificate for the server.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, the Access Management functions do not appear.
- The syslog server address, protocol, and port number must be available. The server address can be a fully qualified domain name, an IPv4 address, or an IPv6 address.
- If you are uploading a new CA certificate, the certificate must be available on your local system.

Steps

1. Select **Settings > Access Management**.
2. From the **Audit Log** tab, select **Configure Syslog Servers**.

Configured syslog servers are displayed on the page.

3. To edit the server information, select the **Edit** (pencil) icon to the right of the server name, and then make desired changes in the following fields:
 - Server Address — Enter a fully qualified domain name, an IPv4 address, or an IPv6 address.
 - Protocol — Select a protocol from the drop-down list (for example, TLS, UDP, or TCP).
 - Port — Enter the port number for the syslog receiver.
4. If you changed the protocol to the secure TLS protocol (from either UDP or TCP), click **Import Trusted Certificate** to upload a CA certificate.

5. To test the new connection with the storage array, select **Test All**.

Result

After configuration, all new audit logs are sent to the syslog server. Previous logs are not transferred.

FAQs

Why can't I log in?

If you receive an error when attempting to log in to System Manager, review these possible causes.

Login errors to System Manager might occur for one of these reasons:

- You entered an incorrect username or password.
- You have insufficient privileges.
- The directory server (if configured) might be unavailable. If this is the case, try logging in with a local user role.
- You attempted to log in unsuccessfully multiple times, which triggered the lockout mode. Wait 10 minutes to re-login.
- A lockout condition was triggered and your audit log might be full. Go to Access Management and delete old events from the audit log.
- SAML authentication is enabled. Refresh your browser to log in.

Login errors to a remote storage array for mirroring tasks might occur for one of these reasons:

- You have entered an incorrect password.
- You attempted to log in unsuccessfully multiple times, which triggered the lockout mode. Wait 10 minutes to log in again.
- The maximum number of client connections used on the controller has been reached. Check for multiple users or clients.

What do I need to know before adding a directory server?

Before adding a directory server in Access Management, make sure you meet the following requirements.

- User groups must be defined in your directory service.
- LDAP server credentials must be available, including the domain name, server URL, and optionally the bind account user name and password.
- For LDAPS servers using a secure protocol, the LDAP server's certificate chain must be installed on your local machine.

What do I need to know about mapping to storage array roles?

Before mapping groups to roles, review the following guidelines.

The storage array's embedded RBAC (role-based access control) capabilities include the following roles:

- **Storage admin** — Full read/write access to the storage objects (for example, volumes and disk pools), but no access to the security configuration.
- **Security admin** — Access to the security configuration in Access Management, certificate management, audit log management, and the ability to turn the legacy management interface (SYMBOL) on or off.
- **Support admin** — Access to all hardware resources on the storage array, failure data, MEL events, and controller firmware upgrades. No access to storage objects or the security configuration.
- **Monitor** — Read-only access to all storage objects, but no access to the security configuration.

Directory Services

If you are using an LDAP (Lightweight Directory Access Protocol) server and Directory Services, make sure that:

- An administrator has defined user groups in the directory service.
- You know the group domain names for the LDAP user groups.
- The Monitor role is required for all users, including the administrator. System Manager will not operate correctly for any user without the Monitor role present.

SAML

If you are using the Security Assertion Markup Language (SAML) capabilities embedded in the storage array, make sure that:

- An Identity Provider (IdP) administrator has configured user attributes and group membership in the IdP system.
- You know the group membership names.
- The Monitor role is required for all users, including the administrator. System Manager will not operate correctly for any user without the Monitor role present.

Which external management tools may be affected by this change?

When you make certain changes in System Manager, such as switching the management interface or using SAML for an authentication method, some external tools and features might be restricted from use.

Management interface

Tools that communicate directly with the legacy management interface (SYMBOL), such as the SANtricity SMI-S Provider or OnCommand Insight (OCI), do not work unless the Legacy Management Interface setting is enabled. In addition, you cannot use legacy CLI commands or perform mirroring operations if this setting is disabled.

Contact technical support for more information.

SAML authentication

When SAML is enabled, the following clients cannot access storage array services and resources:

- Enterprise Management Window (EMW)
- Command-line interface (CLI)

- Software Developer Kits (SDK) clients
- In-band clients
- HTTP Basic Authentication REST API clients
- Login using standard REST API endpoint

Contact technical support for more information.

What do I need to know before configuring and enabling SAML?

Before configuring and enabling the Security Assertion Markup Language (SAML) capabilities for authentication, make sure you meet the following requirements and understand SAML restrictions.

Requirements

Before you begin, make sure that:

- An Identity Provider (IdP) is configured in your network. An IdP is an external system used to request credentials from a user and determine if the user is successfully authenticated. Your security team is responsible for maintaining the IdP.
- An IdP administrator has configured user attributes and groups in the IdP system.
- An IdP administrator has ensured that the IdP supports the ability to return a Name ID on authentication.
- An administrator has ensured that the IdP server and controller clocks are synchronized (either through an NTP server or by adjusting the controller clock settings).
- An IdP metadata file is downloaded from the IdP system and available on the local system used for accessing System Manager.
- You know the IP address or domain name of each controller in the storage array.

Restrictions

In addition to the requirements above, make sure you understand the following restrictions:

- Once SAML is enabled, you *cannot* disable it through the user interface, nor can you edit the IdP settings. If you need to disable or edit the SAML configuration, contact Technical Support for assistance. We recommend that you test the SSO logins before you enable SAML in the final configuration step. (The system also performs an SSO login test before enabling SAML.)
- If you disable SAML in the future, the system automatically restores the previous configuration (Local User Roles and/or Directory Services).
- If Directory Services are currently configured for user authentication, SAML overrides that configuration.
- When SAML is configured, the following clients cannot access storage array resources:
 - Enterprise Management Window (EMW)
 - Command-line interface (CLI)
 - Software Developer Kits (SDK) clients
 - In-band clients
 - HTTP Basic Authentication REST API clients
 - Login using standard REST API endpoint

What types of events are recorded in the audit log?

The audit log can record modification events, or both modification and read-only events.

Depending on the policy settings, the following types of events are shown:

- **Modification events** — User actions from within System Manager that involve changes to the system, such as provisioning storage.
- **Modification and read-only events** — User actions that involve changes to the system, as well as events that involve viewing or downloading information, such as viewing volume assignments.

What do I need to know before configuring a syslog server?

You can archive audit logs onto an external syslog server.

Before configuring a syslog server, keep the following guidelines in mind.

- Make sure you know the server address, protocol, and port number. The server address can be a fully qualified domain name, an IPv4 address, or an IPv6 address.
- If your server uses a secure protocol (for example, TLS), a Certificate Authority (CA) certificate must be available on your local system. CA certificates identify website owners for secure connections between servers and clients.
- After configuration, all new audit logs are sent to the syslog server. Previous logs are not transferred.
- The **Overwrite Policy** settings (available from View/Edit Settings) do not affect how logs are managed with a syslog server configuration.
- Audit logs follow the RFC 5424 messaging format.

The syslog server is no longer receiving audit logs. What do I do?

If you configured a syslog server with a TLS protocol, the server cannot receive messages if the certificate becomes invalid for any reason. An error message about the invalid certificate is posted to the audit log.

To resolve this issue, you must first fix the certificate for the syslog server. Once a valid certificate chain is in place, go to **Settings > Audit Log > Configure Syslog Servers > Test All**.

Certificates

Concepts

How CA certificates work

A certificate authority (CA) is a trusted entity that issues electronic documents, called digital certificates, for Internet security. These certificates identify website owners, which allows for secure connections between clients and servers.

When you open a browser and try connecting to System Manager through the controller management port, the browser attempts to verify that the storage array's controller is a trusted source. If the browser cannot locate a digital certificate for the controller, it alerts you that the certificate is not signed by a recognized authority and asks if you want to continue. If you no longer want to see this alert, you must get a signed, digital certificate

from a CA.

If you are using an external key management server with the Drive Security feature, you can also create certificates for authentication between that server and the controllers or you can accept the self-signed certificate(s) from the storage array.

The following steps are required for using a digital certificate from a trusted authority:

1. Go to **Settings > Certificates**. Your user login must include Security Admin permissions; otherwise, **Certificates** does not appear on the page.
2. Create a Certificate Signing Request (CSR) for each controller or for a key management server.
3. Send the .CSR file(s) to a CA, and then wait for them to send you the certificates.
4. Import the trusted (intermediate and root) certificate from the CA. These certificates establish a point of trust for a CA hierarchy.
5. Import the signed, management certificates for each controller or the key management server.

Certificate terminology

Learn how the certificate terms apply to your storage array.

Term	Description
CA	A certificate authority (CA) is a trusted entity that issues electronic documents, called digital certificates, for Internet security. These certificates identify website owners, which allows for secure connections between clients and servers.
CSR	A Certificate Signing Request (CSR) is a message that is sent from an applicant to a certificate authority (CA). The CSR validates the information the CA requires to issue a certificate.
Certificate	A certificate identifies the owner of a site for security purposes, which prevents attackers from impersonating the site. The certificate contains information about the site owner and the identity of the trusted entity who certifies (signs) this information.
Client certificate	For security key management, a client certificate validates the storage array's controllers, so the key management server can trust their IP addresses.
Key management server certificate	For security key management, a key management server certificate validates the server, so the storage array can trust its IP address.
Management certificate	A management certificate is approved by a certificate authority (CA) and allows secure access to the web application. Also referred to as a "signed certificate."

Term	Description
OCSP server	The Online Certificate Status Protocol (OCSP) server determines if the certificate authority (CA) has revoked any certificates before their scheduled expiration date, and then blocks the user from accessing a server if the certificate is revoked.
Self-signed certificate	A self-signed certificate is pre-loaded on the controller. If the site connection is self-signed, a warning message opens before you can proceed to the web application.
Trusted certificate	A trusted certificate from a certificate authority (CA) is a known certificate at the top of the certificate hierarchy. Also referred to as a "root certificate."

How to

Complete a CA certificate signing request (CSR) for the controllers

To receive a certificate authority (CA) certificate for the storage array's controllers, you must first generate a certificate signing request (CSR) file for each controller in the storage array.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.

About this task

This task describes how to generate the .CSR files (certificate signing requests) that you send to a CA to receive signed, management certificates for the controllers. You must provide information about your organization, plus the IP address or DNS name of the controller(s). During this task, one .CSR file is generated if there is only one controller in the storage array and two .CSR files if there are two controllers.

Steps

1. Select **Settings > Certificates**.
2. From the **Array Management** tab, select **Complete CSR**.



If you see a dialog box prompting you to accept a self-signed certificate for the second controller, click **Accept Self-Signed Certificate** to proceed.

3. Enter the following information, and then click **Next**:
 - **Organization** — The full, legal name of your company or organization. Include suffixes, such as Inc. or Corp.
 - **Organizational unit (optional)** — The division of your organization that is handling the certificate.
 - **City/Locality** — The city where your storage array or business is located.

- **State/Region (optional)** — The state or region where your storage array or business is located.
- **Country ISO code** — Your country's two-digit ISO (International Organization for Standardization) code, such as US.



Some fields might be pre-populated with the appropriate information, such as the IP address of the controller. Do not change prepopulated values unless you are certain they are incorrect. For example, if you have not yet completed a CSR, the controller IP address is set to "localhost." In this case, you must change "localhost" to the DNS name or IP address of the controller.

4. Verify or enter the following information about controller A in your storage array:

- **Controller A common name** — The IP address or DNS name of controller A is displayed by default. Make sure this address is correct; it must match exactly what you enter to access System Manager in the browser.
- **Controller A alternate IP addresses** — If the common name is an IP address, you can optionally enter any additional IP addresses or aliases for controller A. For multiple entries, use a comma-delimited format.
- **Controller A alternate DNS names** — If the common name is a DNS name, enter any additional DNS names for controller A. For multiple entries, use a comma-delimited format. If there are no alternate DNS names, but you entered a DNS name in the first field, copy that name here. If the storage array has only one controller, the **Finish** button is available. If the storage array has two controllers, the **Next** button is available.



Do not click the **Skip this step** link when you are initially creating a CSR request. This link is provided in error-recovery situations. In rare cases, a CSR request might fail on one controller, but not on the other. This link allows you to skip the step for creating a CSR request on controller A if it is already defined, and continue to the next step for re-creating a CSR request on controller B.

5. If there is only one controller, click **Finish**. If there are two controllers, click **Next** to enter information for controller B (same as above), and then click **Finish**.

For a single controller, one .CSR file is downloaded to your local system. For dual controllers, two .CSR files are downloaded. The folder location of the download depends on your browser.

6. Send the .CSR file(s) to your CA.

After you finish

When you receive the digital certificates, import the appropriate certificate files that the CA sent to you.

Import trusted certificates for controllers

After receiving digital certificates from a certificate authority (CA), you can import the certificate chain (intermediate and root) for the controllers.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.
- You have generated a certificate signing request (.CSR file) and sent it to the CA.
- The CA returned trusted certificate files.

- The certificate files are installed on your local system.

About this task

This task describes how to upload the trusted certificates for the storage array's controllers.

Steps

1. Select **Settings > Certificates**.
2. From the **Trusted** tab, select **Import**.

A dialog box opens for importing the trusted certificate files.

3. Click **Browse** to select the certificate files for the controllers.

The file names display in the dialog box.

4. Click **Import**.

Results

The files are uploaded and validated.

After you finish

Import the management certificate.

Import a management certificate for controllers

After importing the trusted certificate chain, you import a management (signed) certificate file for each controller in the storage array.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.
- The trusted certificates have been imported.
- The CA returned a management certificate file for each controller.
- The management certificate file(s) are available on your local system.

About this task

This task describes how to upload the management certificate files for controller authentication.

Steps

1. Select **Settings > Certificates**.
2. From the **Array Management** tab, select **Import**.

A dialog box opens for importing the certificate file(s).

3. Click **Browse** to select the file for controller A. If there are two controllers, click the second **Browse** button to select the file for controller B.

The file names are displayed in the dialog box.

4. Click **Import**.

The file(s) are uploaded and validated.

Results

The session is automatically terminated. You must log in again for the certificate(s) to take effect. When you log in again, the new CA-signed certificate is used for your session.

View imported certificate information

From the Certificates page, you can view the certificate type, issuing authority, and the valid date range of certificates you previously imported.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.

About this task

This task describes how to view information for user-installed or pre-installed certificates.

Steps

1. Select **Settings > Certificates**.
2. Select one of the tabs to view information about management certificates for the controllers, trusted certificates, and certificates for a key management server.

Tab	Description
Array Management	View information about all the server certificates imported for the controllers.
Trusted	<p>View information about all the trusted (root) certificates imported for the controllers. Use the filter field under Show certificates that are... to view either user-installed or pre-installed certificates.</p> <ul style="list-style-type: none">• User-installed. Certificates that a user uploaded to the storage array (includes trusted certificates, LDAPS certificates, and Identity Federation certificates).• Pre-installed. Certificates included with the storage array.
Key Management	View information about all the management (signed) certificates imported for an external key management server.

Delete trusted certificates

You can delete any of the user-imported certificates.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.
- If you are updating a trusted certificate with a new version, the updated certificate must be imported before you delete the old certificate.



You might lose access to the system if you delete a certificate used to authenticate the storage array's management certificates or LDAP server before you import a replacement certificate.

About this task

This task describes how to delete user-imported certificates. Predefined certificates cannot be deleted.

Steps

1. Select **Settings** > **Certificates**.
2. Select the **Trusted** tab.

The table shows the storage array's trusted certificates.

3. From the table, select the certificate you want to remove.
4. Click **Uncommon Tasks** > **Delete**.

A Confirm Delete Trusted Certificate dialog box opens.

5. Type `delete` in the field, and then click **Delete**.

Reset management certificates

You can revert the management certificates on the storage array back to the factory self-signed state.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.
- Certificates must be previously imported.

About this task

Resetting the management certificates on the storage array deletes the current management certificates from each of the controllers. After the certificates are reset, the controllers revert to using self-signed certificates.

Steps

1. Select **Settings** > **Certificates**.
2. From the **Array Management** tab, select **Reset**.

A **Confirm Reset Management Certificates** dialog box opens.

3. Type `reset` in the field, and click **Reset**.

Results

After your browser refreshes, the controllers revert to using self-signed certificates. As a result, the system prompts users to manually accept the self-signed certificate for their sessions.

Complete CA certificate signing request (CSR) for a key server

To receive a certificate authority (CA) certificate for a key management server, you must first generate a certificate signing request (CSR) file.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.

About this task

This task describes how to generate the .CSR files (certificate signing requests) that you send to a CA to receive signed certificates for a key management server. During this task, you must provide information about your organization.

Steps

1. Select **Settings > Certificates**.
2. From the **Key Management** tab, select **Complete CSR**.
3. Enter the following information:
 - **Common name** — A name that identifies this CSR, such as the storage array name, which will be displayed in the certificate files.
 - **Organization** — The full, legal name of your company or organization. Include suffixes, such as Inc. or Corp.
 - **Organizational unit (optional)** — The division of your organization that is handling the certificate.
 - **City/Locality** — The city or locality where your organization is located.
 - **State/Region (optional)** — The state or region where your organization is located.
 - **Country ISO code** — The two-digit ISO (International Organization for Standardization) code, such as US, where your organization is located.
4. Click **Download**.

A .CSR file is saved to your local system.
5. Send the .CSR file(s) to your CA.

After you finish

When you obtain the client and server certificates from the key management server, import them for authentication with the storage array controllers.

Import key management server certificates

For external key management, you import certificates for authentication between the storage array and the key management server so the two entities can trust each other. There are two types of certificates: the client certificate validates the controllers, while the key management server certificate validates the server.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.
- A client certificate is available for the storage array.



A client certificate validates the storage array's controllers, so the key management server can trust their IP addresses. To obtain a client certificate, you must complete a CSR for the storage array and then upload it to the key management server. From the server, generate a client certificate.

- The key management server certificate is available.



A key management server certificate validates the server, so the storage array can trust its IP address. To obtain a key management server certificate, you must generate it from the key management server.

About this task

This task describes how to upload certificate files for authentication between the storage array controllers and the key management server.

Steps

1. Select **Settings** > **Certificates**.
2. From the **Key Management** tab, select **Import**.

A dialog box opens for importing the certificate files.

3. Click the **Browse** buttons to select the files.

The file names display in the dialog box.

4. Click **Import**.

The file(s) are uploaded and validated.

Export key management server certificates

You can save a certificate for a key management server to your local machine.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.
- Certificates must be previously imported.

Steps

1. Select **Settings** > **Certificates**.
2. Select the **Key Management** tab.
3. From the table, select the certificate you want to export, and then click **Export**.

A Save dialog box opens.

4. Enter a filename and click **Save**.

Enable certificate revocation checking

You can enable automatic checks for revoked certificates, so that an Online Certificate Status Protocol (OCSP) server blocks users from making non-secure connections. Automatic revocation checking is helpful in cases where the Certificate Authority (CA) improperly issued a certificate or if a private key is compromised.

Before you begin

- You must be logged in with a user profile that includes Security admin permissions. Otherwise, certificate functions do not appear.
- A DNS server is configured on both controllers, which enables use of a fully qualified domain name for the OCSP server. This task is available from the Hardware page.
- If you want to specify your own OCSP server, you must know the URL of that server.

About this task

During this task, you can configure an OCSP server or use the server specified in the certificate file. The OCSP server determines if the CA has revoked any certificates before their scheduled expiration date, and then blocks the user from accessing a site if the certificate is revoked.

Steps

1. Select **Settings > Certificates**.
2. Select the **Trusted** tab.



You can also enable revocation checking from the Key Management tab.

3. Click **Uncommon Tasks**, and then select **Enable Revocation Checking** from the drop-down menu.
4. Select **I want to enable revocation checking**, so that a checkmark appears in the checkbox and additional fields appear in the dialog box.
5. In the **OCSP responder address** field, you can optionally enter a URL for an OCSP responder server. If you do not enter an address, the system uses the OCSP server's URL from the certificate file.
6. Click **Test Address** to make certain the system can open a connection to the specified URL.
7. Click **Save**.

Result

If the storage array attempts to connect to a server with a revoked certificate, the connection is denied and an event is logged.

FAQs

Why does the Cannot Access Other Controller dialog box appear?

When you perform certain operations related to CA certificates (for example, importing a certificate), you might see a dialog box prompting you to accept a self-signed certificate for the second controller.

In storage arrays with two controllers (duplex configurations), this dialog box sometimes appears if SANtricity System Manager cannot communicate with the second controller or if your browser cannot accept the certificate during a certain point in an operation.

If this dialog box opens, click **Accept Self-Signed Certificate** to proceed. If another dialog box prompts you for a password, enter your Administrator password used for accessing System Manager.

If this dialog box appears again and you cannot complete a certificate task, try one of the following procedures:

- Use a different browser type to access this controller, accept the certificate, and continue.
- Access the second controller with System Manager, accept the self-signed certificate, and then return to the first controller and continue.

How do I know what certificates need to be uploaded to System Manager?

For external key management, you import two types of certificates for authentication between the storage array and the key management server so the two entities can trust each other.

A client certificate validates the storage array's controllers, so the key management server can trust their IP addresses. To obtain a client certificate, you must complete a CSR for the storage array and then upload it to the key management server. From the server, generate a client certificate, and then use System Manager to import it.

A key management server certificate validates the key management server, so the storage array can trust its IP address. To obtain a key management server certificate, you must generate it from the key management server.

What do I need to know about certificate revocation checking?

System Manager allows you to check for revoked certificates by using an Online Certificate Status Protocol (OCSP) server, instead of uploading Certificate Revocation Lists (CRLs).

Revoked certificates should no longer be trusted. A certificate might be revoked for several reasons; for example, if the Certificate Authority (CA) improperly issued the certificate, a private key was compromised, or the identified entity did not adhere to policy requirements.

After you establish a connection to an OCSP server in System Manager, the storage array performs revocation checking whenever it connects to an AutoSupport server, External Key Management Server (EKMS), Lightweight Directory Access Protocol over SSL (LDAPS) server, or a Syslog server. The storage array attempts to validate these servers' certificates to ensure that they have not been revoked. The server then returns a value of "good," "revoked," or "unknown" for that certificate. If the certificate is revoked or the array cannot contact the OCSP server, the connection is refused.



Specifying an OCSP responder address in System Manager or in the command line interface (CLI) overrides the OCSP address found in the certificate file.

What types of servers will revocation checking be enabled for?

The storage array performs revocation checking whenever it connects to an AutoSupport server, External Key Management Server (EKMS), Lightweight Directory Access Protocol over SSL (LDAPS) server, or a Syslog server.

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.