



System

SANtricity 11.5

NetApp
February 12, 2024

Table of Contents

- System 1
 - Storage array settings 1
 - iSCSI settings 15
 - System: NVMe settings 27
 - Add-on features 35
 - Security key management 38

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 tos

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
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 CHAPsecret.
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 tos

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.)	If necessary, enter a new port number. 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.
MTU size (Available by clicking Show more port settings.)	If necessary, enter a new size in bytes for the Maximum Transmission Unit (MTU). The default Maximum Transmission Unit (MTU) size is 1500 bytes per frame. You must enter a value between 1500 and 9000.
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	<ol style="list-style-type: none">1. Select Hardware.2. Select Show back of shelf.3. Select a controller.4. Select Configure iSER over InfiniBand ports. or <ol style="list-style-type: none">1. Select Settings > System.2. Scroll down to iSER over InfiniBand settings, and then select Configure iSER over InfiniBand Ports.
View iSER over InfiniBand statistics	<ol style="list-style-type: none">1. Select Settings > System.2. Scroll down to iSER over InfiniBand settings, and then select View iSER over InfiniBand Statistics.

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 tos

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	<p>Select one or both options to enable support for IPv4 and IPv6 networks.</p> <p> If you want to disable port access, deselect both check boxes.</p>
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>

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 data-bbox="844 1344 1477 1470" style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;">  <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).

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.