



Storage

SANtricity 11.5

NetApp

February 12, 2024

This PDF was generated from <https://docs.netapp.com/us-en/e-series-santricity-115/sm-storage/how-pools-and-volume-groups-work.html> on February 12, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Storage..... 1
 - Pools and volume groups..... 1
 - Volumes 65
 - Hosts 132
 - Performance..... 150
 - Snapshots 162

Storage

Pools and volume groups

Concepts


How pools and volume groups work

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

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

- Pools
- RAID volume groups

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



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

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

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

You can have pools or volume groups on the same storage system, but a drive cannot be part of more than

one pool or volume group. Volumes that are presented to hosts for I/O are then created, using the space on the pool or volume group.

Pools

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

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

Drive matching

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

Managing failed drives

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

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

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

Volume groups

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

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

Drive matching

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

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

Managing failed drives

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

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

Capacity terminology

Learn how the capacity terms apply to your storage array.

Storage objects

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

Storage object	Description
Host	A host is a server that sends I/O to a volume on a storage array.
LUN	<p>A logical unit number (LUN) is the number assigned to the address space that a host uses to access a volume. The volume is presented to the host as capacity in the form of a LUN.</p> <p>Each host has its own LUN address space. Therefore, the same LUN can be used by different hosts to access different volumes.</p>
Mirror consistency group	A mirror consistency group is a container for one or more mirrored pairs. For asynchronous mirroring operations, you must create a mirror consistency group.
Mirrored volume pair	A mirrored pair is comprised of two volumes, a primary volume and a secondary volume.
Pool	A pool is a set of drives that is logically grouped. You can use a pool to create one or more volumes accessible to a host. (You create volumes from either a pool or a volume group.)
Snapshot consistency group	A snapshot consistency group is a collection of volumes that are treated as a single entity when a snapshot image is created. Each of these volumes has its own snapshot image, but all the images are created at the same point in time.

Storage object	Description
Snapshot group	A snapshot group is a collection of snapshot images from a single base volume.
Snapshot volume	A snapshot volume allows the host to access data in the snapshot image. The snapshot volume contains its own reserved capacity, which saves any modifications to the base volume without affecting the original snapshot image.
Volume	A volume is a container in which applications, databases, and file systems store data. It is the logical component created for the host to access storage on the storage array.
Volume group	A volume group is a container for volumes with shared characteristics. A volume group has a defined capacity and RAID level. You can use a volume group to create one or more volumes accessible to a host. (You create volumes from either a volume group or a pool.)

Storage capacity

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

Capacity type	Description
Allocated capacity	<p>Allocated capacity is the physical capacity allocated from the drives in a pool or volume group.</p> <p>You use allocated capacity to create volumes and for copy services operations.</p>
Free capacity	Free capacity is the capacity available in a pool or volume group that has not yet been allocated to volume creation or copy services operations and storage objects.
Pool or volume group capacity	Pool, volume, or volume group capacity is the capacity in a storage array that has been assigned to a pool or volume group. This capacity is used to create volumes and service the various capacity needs of copy services operations and storage objects.
Pool unusable capacity	Pool unusable capacity is the space in a pool that cannot be used due to mismatched drive sizes.

Capacity type	Description
Preservation capacity	Preservation capacity is the amount of capacity (number of drives) that is reserved in a pool to support potential drive failures.
Reported capacity	Reported capacity is the capacity that is reported to the host and can be accessed by the host.
Reserved capacity	Reserved capacity is the physical allocated capacity that is used for any copy service operation and storage object. It is not directly readable by the host.
SSD Cache	SSD Cache is a set of Solid-State Disk (SSD) drives that you logically group together in your storage array. The SSD Cache feature caches the most frequently accessed data ("hot" data) onto lower latency SSD drives to dynamically accelerate application workloads.
Unassigned capacity	Unassigned capacity is the space in a storage array that has not been assigned to a pool or volume group.
Written capacity	Written capacity is the amount of capacity that has been written from the reserved capacity allocated for thin volumes.

How reserved capacity works

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

Copy service objects that use reserved capacity

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

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

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

Thin volumes and reserved capacity

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

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

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

Capacity alerts

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

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

Orphaned reserved capacity volumes

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

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

System Manager alerts you of orphaned reserved capacity volumes with a **Reclaim unused capacity** message in the Notifications area on the Home page. You can click **Reclaim unused capacity** to display the Reclaim Unused Capacity dialog box, where you can delete the orphaned reserved capacity volume.

Characteristics of reserved capacity

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

- Reserved capacity exists for each read/write snapshot volume, snapshot group, consistency group member volume, and mirrored pair volume.

How SSD Cache works

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

SSD Cache versus primary cache

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

SSD Cache operates differently than primary cache:

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

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

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

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

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

How SSD Cache is used

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

When a host program accesses the storage array's drives, the data is stored in the SSD Cache. When the same data is accessed by the host program again, it is read from the SSD Cache instead of the hard drives. The commonly accessed data is stored in the SSD Cache. The hard drives are only accessed when the data cannot be read from the SSD Cache.

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

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

Steps

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

Improved performance

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

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

SSD Cache and the Drive Security feature

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

Implement SSD Cache

To implement SSD Cache, do the following:

Steps

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



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

SSD Cache restrictions

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

- Any volume assigned to use a controller's SSD Cache is not eligible for an automatic load balance transfer.
- Currently, only one SSD Cache is supported per storage array.
- The maximum usable SSD Cache capacity on a storage array depends on the controller's primary cache capacity.
- SSD Cache is not supported on snapshot images.
- If you import or export volumes that are SSD Cache enabled or disabled, the cached data is not imported or exported.
- You cannot remove the last drive in an SSD Cache without first deleting the SSD Cache.

SSD Cache restrictions with Drive Security

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

Decide whether to use a pool or a volume group

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

Reasons to choose a pool or volume group

Choose a pool

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

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

Choose a volume group

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



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

Feature differences between pools and volume groups

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

Use	Pool	Volume group
Workload random	Better	Good
Workload sequential	Good	Better
Drive rebuild time	Faster	Slower
Performance (optimal mode)	Good: Best for small block, random workload.	Good: Best for large block, sequential workloads
Performance (drive rebuild mode)	Better: Usually better than RAID 6	Degraded: Up to 40% drop in performance
Multiple drive failures	Greater data protection: Faster, prioritized rebuilds	Less data protection: Slow rebuilds, greater risk of data loss

Use	Pool	Volume group
Adding drives	Faster: Add to pool on the fly	Slower: Requires Dynamic Capacity Expansion operation
Thin volumes support	Yes	No
Solid State Disk (SSD) support	Yes	Yes
Simplified administration	Yes: No hot spares or RAID settings to configure	No: Must allocate hot spares, configure RAID
Tunable performance	No	Yes

Functional comparison of pools and volume groups

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

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

Function	Pool	Volume Group
Different RAID level supported	No. Always RAID 6.	Yes. RAID 0, 1, 10, 5, and 6 available.
Thin volumes supported	Yes	No
Full disk encryption (FDE) supported	Yes	Yes
Data Assurance (DA) supported	Yes	Yes
Shelf loss protection supported	Yes	Yes
Drawer loss protection supported	Yes	Yes
Mixed drive speeds supported	Recommended to be the same, but not required. Slowest drive determines speed for all drives.	Recommended to be the same, but not required. Slowest drive determines speed for all drives.
Mixed drive capacity supported	Recommended to be the same, but not required. Smallest drive determines capacity for all drives.	Recommended to be the same, but not required. Smallest drive determines capacity for all drives.

Function	Pool	Volume Group
Minimum number of drives	11	Depends on RAID level. RAID 0 needs 1. RAID 1 or 10 needs 2 (requires an even number). RAID 5 minimum is 3. RAID 6 minimum is 5.
Maximum number of drives	Up to the maximum limit for the storage array	RAID 1 and 10—up to the maximum limit of the storage array RAID 5, 6—30 drives
Can choose individual drives when creating a volume	No	Yes
Can specify segment size when creating a volume	Yes. 128K supported.	Yes
Can specify I/O characteristics when creating a volume	No	Yes. File system, database, multimedia, and custom supported.
Drive failure protection	Uses preservation capacity on each drive in the pool making reconstruction faster.	Uses a hot spare drive. Reconstruction is limited by the IOPs of the drive.
Warning when reaching capacity limit	Yes. Can set an alert when used capacity reaches a percentage of the maximum capacity.	No
Migration to a different storage array supported	No. Requires that you migrate to a volume group first.	Yes
Dynamic Segment Size (DSS)	No	Yes
Can change RAID level	No	Yes
Volume expansion (increase capacity)	Yes	Yes
Capacity expansion (add capacity)	Yes	Yes
Capacity reduction	Yes	No



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

Automatic versus manual pool creation

You create pools automatically or manually to allow physical storage to be grouped, and

then dynamically allocated as needed. When a pool is created, you can add physical drives.

Automatic creation

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

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

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

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

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

Manual creation

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

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

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

How tos

Create pools and volume groups

Create pool automatically

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

Before you begin

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

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

About this task

Keep in mind the following:

- When you add drives to a storage array, System Manager automatically detects the drives and prompts you to create a single pool or multiple pools based on the drive type and the current configuration.
- If pools were previously defined, System Manager automatically prompts you with the option of adding the compatible drives to an existing pool. When new drives are added to an existing pool, System Manager automatically redistributes the data across the new capacity, which now includes the new drives that you added.

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

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

Steps

1. Select **Storage > Pools & Volume Groups**.
2. Select **More > Launch pool auto-configuration**. The results table lists new pools, existing pools with drives added, or both. A new pool is named with a sequential number by default.

Notice that System Manager will do the following:

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

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

5. Click **Accept**.

Create pool manually

You can create a pool manually (from a set of candidates) if the Pool Auto Configuration

feature does not provide a pool that meets your needs. A pool provides the logical storage capacity necessary from which you can create individual volumes that can then be used to host your applications.

Before you begin

- You must have a minimum of 11 drives with the same drive type (HDD or SSD).
- Shelf loss protection requires that the drives comprising the pool are located in at least six different drive shelves and there are no more than two drives in a single drive shelf.
- Drawer loss protection requires that the drives comprising the pool are located in at least five different drawers and the pool includes an equal number of drive shelves from each drawer.

Steps

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


The **Create Pool** dialog box appears.

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

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

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

Characteristic	Use
Free Capacity	<p>Shows the free capacity of the pool candidate in GiB. Select a pool candidate with the capacity for your application’s storage needs.</p> <p>Preservation (spare) capacity is also distributed throughout the pool and is not part of the free capacity amount.</p>
Total Drives	<p>Shows the number of drives available in the pool candidate.</p> <p>System Manager automatically reserves as many drives as possible for preservation capacity (for every six drives in a pool, System Manager reserves one drive for preservation capacity).</p> <p>When a drive failure occurs, the preservation capacity is used to hold the reconstructed data.</p>

Characteristic	Use
Secure-Capable	<p>Indicates whether this pool candidate is comprised entirely of secure-capable drives, which can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.</p> <ul style="list-style-type: none"> You can protect your pool with Drive Security, but all drives must be secure-capable to use this feature. If you want to create an FDE-only pool, look for Yes - FDE in the Secure-Capable column. If you want to create a FIPS-only pool, look for Yes - FIPS in the Secure-Capable column. You can create a pool comprised of drives that may or may not be secure-capable or are a mix of security levels. If the drives in the pool include drives that are not secure-capable, you cannot make the pool secure.
Enable Security?	<p>Provides the option for enabling the Drive Security feature with secure-capable drives. If the pool is secure-capable and you have created a security key, you can enable security by selecting the check box.</p> <div>  <p>The only way to remove Drive Security after it is enabled is to delete the pool and erase the drives.</p> </div>
DA Capable	<p>Indicates if Data Assurance (DA) is available for this pool candidate. DA checks for and corrects errors that might occur as data is communicated between a host and a storage array.</p> <p>If you want to use DA, select a pool that is DA capable. This option is available only when the DA feature has been enabled.</p> <p>A pool can contain drives that are DA-capable or not DA-capable, but all drives must be DA capable for you to use this feature.</p>
Shelf Loss Protection	<p>Shows if shelf loss protection is available.</p> <p>Shelf loss protection guarantees accessibility to the data on the volumes in a pool if a total loss of communication occurs with a single drive shelf.</p>

Characteristic	Use
Drawer Loss Protection	Shows if drawer loss protection is available, which is provided only if you are using a drive shelf that contains drawers. Drawer loss protection guarantees accessibility to the data on the volumes in a pool if a total loss of communication occurs with a single drawer in a drive shelf.

Create volume group

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

About this task

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

Follow these guidelines when you create a volume group.

- You need at least one unassigned drive.
- Limits exist as to how much drive capacity you can have in a single volume group. These limits vary according to your host type.
- To enable shelf/drawer loss protection, you must create a volume group that uses drives located in at least three shelves or drawers, unless you are using RAID 1, where two shelves/drawers is the minimum.

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

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

Steps

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

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

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

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

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

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

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

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

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




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

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

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

Automatic method drive characteristics

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

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

Manual method drive characteristics

Characteristic	Use
Media Type	<p>Indicates the media type. The following media types are supported:</p> <ul style="list-style-type: none">• Hard drive• Solid State Disk (SSD) All drives in a volume group must be of the same media type (either all SSDs or all hard drives). Volume groups cannot have a mixture of media types or interface types.
Drive Capacity	<p>Indicates the drive capacity.</p> <ul style="list-style-type: none">• Whenever possible, select drives that have a capacity equal to the capacities of the current drives in the volume group.• If you must add unassigned drives with a smaller capacity, be aware that the usable capacity of each drive currently in the volume group is reduced. Therefore, the drive capacity is the same across the volume group.• If you must add unassigned drives with a larger capacity, be aware that the usable capacity of the unassigned drives that you add is reduced so that they match the current capacities of the drives in the volume group.
Tray	Indicates the tray location of the drive.
Slot	Indicates the slot location of the drive.
Speed (rpm)	Indicates the speed of the drive.
Logical sector size	Indicates the sector size and format.

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

Create SSD Cache

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

Before you begin

Your storage array must contain some SSD drives.

About this task

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

Follow these guidelines when you create SSD Cache.

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


Steps

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

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

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

Characteristic	Use
Capacity	<p>Shows the available capacity in GiB. Select the capacity for your application's storage needs.</p> <p>The maximum capacity for SSD Cache depends on the controller's primary cache capacity. If you allocate more than the maximum amount to SSD Cache, then any extra capacity is unusable.</p> <p>SSD Cache capacity counts towards your overall allocated capacity.</p>
Total drives	<p>Shows the number of drives available for this SSD cache. Select the SSD candidate with the number of drives that you want.</p>
Secure-capable	<p>Indicates whether the SSD Cache candidate is comprised entirely of secure-capable drives, which can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.</p> <p>If you want to create a secure-enabled SSD Cache, look for Yes - FDE or Yes - FIPS in the Secure-capable column.</p>

Characteristic	Use
Enable security?	<p>Provides the option for enabling the Drive Security feature with secure-capable drives. If you want to create a secure-enabled SSD Cache, select the Enable Security check box.</p> <div>  <p>Once enabled, security cannot be disabled. You can enable security on the SSD Cache only when you are creating it, not later.</p> </div>
DA capable	<p>Indicates if Data Assurance (DA) is available for this SSD Cache candidate. Data Assurance (DA) checks for and corrects errors that might occur as data is communicated between a host and a storage array.</p> <p>If you want to use DA, select an SSD Cache candidate that is DA capable. This option is available only when the DA feature has been enabled.</p> <p>SSD Cache can contain both DA-capable and non-DA-capable drives, but all drives must be DA-capable for you to use DA.</p>

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

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

6. Click **Create**.

Add capacity to a pool or volume group

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

Before you begin

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

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

Standard (FIPS) drives.

About this task

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



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



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

Steps

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

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

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

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

Field Details

Field	Description
Shelf	Indicates the shelf location of the drive.
Bay	Indicates the bay location of the drive.
Capacity (GiB)	<p>Indicates the drive capacity.</p> <ul style="list-style-type: none">• Whenever possible, select drives that have a capacity equal to the capacities of the current drives in the pool or volume group.• If you must add unassigned drives with a smaller capacity, be aware that the usable capacity of each drive currently in the pool or volume group is reduced. Therefore, the drive capacity is the same across the pool or volume group.• If you must add unassigned drives with a larger capacity, be aware that the usable capacity of the unassigned drives that you add is reduced so that they match the current capacities of the drives in the pool or volume group.
Secure-Capable	<p>Indicates whether the drive is secure-capable.</p> <ul style="list-style-type: none">• You can protect your pool or volume group with the Drive Security feature, but all drives must be secure-capable to use this feature.• You can mix secure-capable and non-secure-capable drives, but the encryption abilities of the secure-capable drives cannot be used.• Secure-capable drives can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.

Field	Description
DA Capable	<p>Indicates whether the drive is Data Assurance (DA) capable.</p> <ul style="list-style-type: none"> Using drives that are not Data Assurance (DA) capable to add capacity to a DA-capable pool or volume group is not recommended. The pool or volume group no longer has DA capabilities, and you no longer have the option to enable DA on newly created volumes within the pool or volume group. Using drives that are Data Assurance (DA) capable to add capacity to a pool or volume group that is non DA-capable is not recommended, because that pool or volume group cannot take advantage of the capabilities of the DA-capable drive (the drive attributes do not match). Consider using drives that are not DA-capable in this situation.

4. Click **Add**.

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

- Shelf loss protection
- Drawer loss protection
- Full Disk Encryption capability
- Data Assurance capability To continue, click **Yes**; otherwise click **Cancel**.

Results

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

Manage pools, volume groups, and SSD Cache

Change configuration settings for a pool

You can edit the settings for a pool if you want to change its name, edit the capacity alerts settings, its modification priorities, or preservation capacity.

Steps

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

The **Pool Settings** dialog box appears.

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

Field Details

Setting	Description
Name	You can change the user-supplied name of the pool. Specifying a name for a pool is required.
Capacity alerts	<p>You can send alert notifications when the free capacity in a pool reaches or exceeds a specified threshold. When the data stored in the pool exceeds the specified threshold, System Manager sends a message, allowing you time to add more storage space or to delete unnecessary objects.</p> <p>Alerts are shown in the Notifications area on the Dashboard and can be sent from the server to administrators by email and SNMP trap messages.</p> <p>You can define the following capacity alerts:</p> <ul style="list-style-type: none">• Critical alert — This critical alert notifies you when the free capacity in the pool reaches or exceeds the specified threshold. Use the spinner controls to adjust the threshold percentage. Select the check box to disable this notification.• Early alert — This early alert notifies you when the free capacity in a pool is reaching a specified threshold. Use the spinner controls to adjust the threshold percentage. Select the check box to disable this notification.

Setting	Description
Modification priorities	<p data-bbox="841 157 1451 430">You can specify the priority levels for modification operations in a pool relative to system performance. A higher priority for modification operations in a pool causes an operation to complete faster, but can slow the host I/O performance. A lower priority causes operations to take longer, but host I/O performance is less affected.</p> <p data-bbox="841 464 1451 598">You can choose from five priority levels: lowest, low, medium, high, and highest. The higher the priority level, the larger is the impact on host I/O and system performance.</p> <ul data-bbox="867 632 1451 1312" style="list-style-type: none"> <li data-bbox="867 632 1451 835">• Critical reconstruction priority — This slider bar determines the priority of a data reconstruction operation when multiple drive failures result in a condition where some data has no redundancy and an additional drive failure might result in loss of data. <li data-bbox="867 856 1451 1060">• Degraded reconstruction priority — This slider bar determines the priority of the data reconstruction operation when a drive failure has occurred, but the data still has redundancy and an additional drive failure does not result in loss of data. <li data-bbox="867 1081 1451 1312">• Background operation priority — This slider bar determines the priority of the pool background operations that occur while the pool is in an optimal state. These operations include Dynamic Volume Expansion (DVE), Instant Availability Format (IAF), and migrating data to a replaced or added drive.

Setting	Description
Preservation capacity	<p>You can define the number of drives to determine the capacity that is reserved on the pool to support potential drive failures. When a drive failure occurs, the preservation capacity is used to hold the reconstructed data. Pools use preservation capacity during the data reconstruction process instead of hot spare drives, which are used in volume groups.</p> <p>Use the spinner controls to adjust the number of drives. Based on the number of drives, the preservation capacity in the pool appears next to the spinner box.</p> <p>Keep the following information in mind about preservation capacity.</p> <ul style="list-style-type: none"> Because preservation capacity is subtracted from the total free capacity of a pool, the amount of capacity that you reserve affects how much free capacity is available to create volumes. If you specify 0 for the preservation capacity, all of the free capacity on the pool is used for volume creation. If you decrease the preservation capacity, you increase the capacity that can be used for pool volumes.

4. Click **Save**.

Change SSD Cache settings

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

Steps

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

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

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

Field Details

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

Setting	Description
Associated objects	Shows the volumes and drives associated with the SSD Cache.

4. Click **Save**.

Change RAID level for a volume group

You can change the RAID level for a volume group to accommodate the performance needs of the applications that are accessing the volume group. This operation changes a volume group's RAID level without impacting data I/O.

Before you begin

- The volume group must be in Optimal status.
- You must have enough capacity in the volume group to convert to the new RAID level.
- You cannot change the RAID level of a pool. System Manager automatically configures pools as RAID 6.

About this task

You cannot cancel this operation after it begins. Your data remains available during this operation.

More about RAID levels

RAID Level	Description
RAID 0 striping	<p>Offers high performance, but does not provide any data redundancy. If a single drive fails in the volume group, all of the associated volumes fail, and all data is lost.</p> <p>A striping RAID group combines two or more drives into one large, logical drive.</p>
RAID 1 mirroring	<p>Offers high performance and the best data availability, and is suitable for storing sensitive data on a corporate or personal level.</p> <p>Protects your data by automatically mirroring the contents of one drive to the second drive in the mirrored pair. It provides protection in the event of a single drive failure.</p>
RAID 10 striping/mirroring	<p>Provides a combination of RAID 0 (striping) and RAID 1 (mirroring), and is achieved when four or more drives are selected.</p> <p>RAID 10 is suitable for high volume transaction applications, such as a database, that require high performance and fault tolerance.</p>

RAID Level	Description
RAID 5	Optimal for multi-user environments (such as database or file system storage) where typical I/O size is small and there is a high proportion of read activity.
RAID 6	Optimal for environments requiring redundancy protection beyond RAID 5, but not requiring high write performance.

RAID 3 can be assigned only to volume groups using the command line interface (CLI).

Steps

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

The Volume Group Settings dialog box appears.

3. Select the RAID level from the drop-down list, and then click **Save**.

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

Results

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

View SSD Cache statistics

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

About this task

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

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

Steps

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

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

Field Details

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

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

Check volume redundancy

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

Before you begin

- The status of the pool or volume group must be Optimal.

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



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

About this task

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

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



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

Steps

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

The **Check Redundancy** dialog box appears.

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

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

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

More about the status of the redundancy check

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

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

Delete pool or volume group

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

Before you begin

- You must have backed up the data on all of the volumes in the pool or volume group.
- You must have stopped all input/output (I/O).
- Unmount any file systems on the volumes.
- You must have deleted any mirror relationships in the pool or volume group.
- You must have stopped any volume copy operation in progress for the pool or volume group.

- The pool or volume group must not be participating in an asynchronous mirroring operation.
- The drives in the pool or volume group must have a persistent reservation.

Steps

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

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

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

Results

System Manager performs the following actions:

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

Consolidate free capacity for a volume group

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

Before you begin

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

About this task

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

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

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

More about free capacity areas

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

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

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

Steps

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

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

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

Results

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

After you finish

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

Export/Import volume groups

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

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

Manage drives

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

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

Steps

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

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

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

Remove capacity from a pool or SSD Cache

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

About this task

Follow these guidelines when you remove capacity:

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

Read about potential performance impacts

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

Read about impacts to secure-capable drives

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

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

Steps

1. Select **Storage > Pools & Volume Groups**.
2. Select the pool or SSD Cache, and then click **More > Remove capacity**.

The **Remove Capacity** dialog box appears.

3. Select one or more drives in the list.

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

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

Results

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

Enable security for a pool or volume group

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

Before you begin

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

About this task

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

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

Steps

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

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

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

Assign hot spares

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

Before you begin

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

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

About this task

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



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

Steps

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

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

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

The drive's context menu opens.

4. Select **Assign hot spare**.

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



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

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

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

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

Results

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

Replace drive logically

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

About this task

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

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

The replacement drive must have the following characteristics:

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

Steps

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

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

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

The drive's context menu appears.

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

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

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

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

7. Click **Replace**.

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

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

Manage reserved capacity

Increase reserved capacity

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

Before you begin

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

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

About this task

You can increase reserved capacity only in increments of 4 GiB for the following storage objects:

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

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



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

Steps

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

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

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

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

5. Click **Increase**.

Results

System Manager performs the following actions:

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

Decrease reserved capacity

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

Before you begin

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

About this task

Review the following guidelines:

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

Steps

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

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

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

Results

System Manager performs the following actions:

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

Change the reserved capacity settings for a snapshot group

You can change the settings for a snapshot group to change its name, auto-delete settings, the maximum number of allowed snapshot images, the percentage point at

which System Manager sends a reserved capacity alert notification, or the policy to use when the reserved capacity reaches its maximum defined percentage.

Before you begin

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

Steps

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

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

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

Field Details

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

Setting	Description
Associated objects	
Base volume	The name of the base volume used for the group. A base volume is the source from which a snapshot image is created. It can be a thick or thin volume and is typically assigned to a host. The base volume can reside in either a volume group or disk pool.
Snapshot images	The number of images created from this group. A snapshot image is a logical copy of volume data, captured at a particular point-in-time. Like a restore point, snapshot images allow you to roll back to a known good data set. Although the host can access the snapshot image, it cannot directly read or write to it.

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

Change the reserved capacity settings for a snapshot volume

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

Steps

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

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

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

Field Details

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

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

Change the reserved capacity settings for a consistency group member volume

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

About this task

Changing the reserved capacity settings for an individual member volume also changes the reserved capacity settings for all member volumes associated with a consistency group.


Steps

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

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

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

Field Details

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

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

Results

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

Change the reserved capacity settings for a mirrored pair volume

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


Steps

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

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

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

Field details

Setting	Description
Alert me when...	<p>Use the spinner box to adjust the percentage point at which System Manager sends an alert notification when the reserved capacity for a mirrored pair is nearing full.</p> <p>When the reserved capacity for the mirrored pair exceeds the specified threshold, System Manager sends an alert, allowing you time to increase reserved capacity.</p> <div><p>Changing the Alert setting for one mirrored pair changes the Alert setting for all mirrored pairs that belong to the same mirror consistency group.</p></div>

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

Cancel pending snapshot image

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

About this task

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

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

Steps

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

Delete snapshot group

You delete a snapshot group when you want to permanently delete its data and remove it

from the system. Deleting a snapshot group reclaims reserved capacity for reuse in the pool or volume group.

About this task

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

Steps

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

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

4. Type `delete` to confirm.

Results

System Manager performs the following actions:

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

FAQs

What is a hot spare drive?

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

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

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



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

What is a volume group?

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

What is a pool?

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

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

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

What is reserved capacity?

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

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

What is FDE/FIPS security?

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

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



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

What is redundancy check?

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

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

- Scans the data blocks in a RAID 3 volume, a RAID 5 volume, or a RAID 6 volume, and then checks the redundancy information for each block. (RAID 3 can only be assigned to volume groups using the command line interface.)

- Compares the data blocks on RAID 1 mirrored drives.
- Returns redundancy errors if the data is determined to be inconsistent by the controller firmware.



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

What are the differences between pools and volume groups?

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

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

Why would I want to manually configure a pool?

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

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

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

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

- If there are multiple pools of the same drive type, a message appears indicating that System Manager cannot recommend the drives for a pool automatically. However, you can manually add the drives to an existing pool.

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

Why are capacity alerts important?

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

You set this percentage when you create a pool using either the **Pool auto-configuration** option or the **Create pool** option. If you choose the automatic option, default settings automatically determine when you receive alert notifications. If you choose to manually create the pool, you can determine the alert notification settings;

or if you prefer, you can accept the default settings. You can adjust these settings later in **Settings > Alerts**.



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

Why can't I increase my preservation capacity?

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

Preservation capacity is the amount of capacity (number of drives) that is reserved on a pool to support potential drive failures. When a pool is created, System Manager automatically reserves a default amount of preservation capacity depending on the number of drives in the pool. If you have created volumes on all available usable capacity, you cannot increase preservation capacity without adding capacity to the pool by either adding drives or deleting volumes.

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



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

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

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

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

What media types are supported for a drive?

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

Why are some drives not showing up?

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

Drives are not eligible for any of the following reasons:

- A drive must be unassigned and not secure-enabled. Drives already part of another pool, another volume group, or configured as a hot spare are not eligible. If a drive is unassigned but is secure-enabled, you must manually erase that drive for it to become eligible.
- A drive that is in a non-optimal state is not eligible.
- If the capacity of a drive is too small, it is not eligible.

- The drive media type must match within a pool or volume group. You cannot mix Hard Disk Drives (HDDs) with Solid State Disks (SSDs).
- If a pool or volume group contains all secure-capable drives, non-secure-capable drives are not listed.
- If a pool or volume group contains all Federal Information Processing Standards (FIPS) drives, non-FIPS drives are not listed.
- If a pool or volume group contains all Data Assurance (DA)-capable drives and there is at least one DA-enabled volume in the pool or volume group, a drive that is not DA capable is not eligible, so it cannot be added to that pool or volume group. However, if there is no DA-enabled volume in the pool or volume group, a drive that is not DA capable can be added to that pool or volume group. If you decide to mix these drives, keep in mind that you cannot create any DA-enabled volumes.



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

What is shelf loss protection and drawer loss protection?

Shelf loss protection and drawer loss protection are attributes of pools and volume groups that allow you to maintain data access in the event of a single shelf or drawer failure.

Shelf loss protection

A shelf is the enclosure that contains either the drives or the drives and the controller. Shelf loss protection guarantees accessibility to the data on the volumes in a pool or volume group if a total loss of communication occurs with a single drive shelf. An example of total loss of communication might be loss of power to the drive shelf or failure of both I/O modules (IOMs).



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

The criteria for shelf loss protection depends on the protection method, as described in the following table:

Level	Criteria for Shelf Loss Protection	Minimum number of shelves required
Pool	The pool must include drives from at least five shelves and there must be an equal number of drives in each shelf. Shelf loss protection is not applicable to high-capacity shelves; if your system contains high-capacity shelves, refer to Drawer Loss Protection.	5
RAID 6	The volume group contains no more than two drives in a single shelf.	3

Level	Criteria for Shelf Loss Protection	Minimum number of shelves required
RAID 3 or RAID 5	Each drive in the volume group is located in a separate shelf.	3
RAID 1	Each drive in a RAID 1 pair must be located in a separate shelf.	2
RAID 0	Cannot achieve Shelf Loss Protection.	Not applicable

Drawer loss protection

A drawer is one of the compartments of a shelf that you pull out to access the drives. Only the high-capacity shelves have drawers. Drawer loss protection guarantees accessibility to the data on the volumes in a pool or volume group if a total loss of communication occurs with a single drawer. An example of total loss of communication might be loss of power to the drawer or failure of an internal component within the drawer.



Drawer loss protection is not guaranteed if a drive has already failed in the pool or volume group. In this situation, losing access to a drawer (and consequently another drive in the pool or volume group) causes loss of data.

The criteria for drawer loss protection depends on the protection method, as described in the following table:

Level	Criteria for drawer loss protection	Minimum number of drawers required
Pool	<p>Pool candidates must include drives from all drawers, and there must be an equal number of drives in each drawer. The pool must include drives from at least five drawers and there must be an equal number of drives in each drawer.</p> <p>A 60-drive shelf can achieve Drawer Loss Protection when the pool contains 15, 20, 25, 30, 35, 40, 45, 50, 55, or 60 drives. Increments in multiples of 5 can be added to the pool after initial creation.</p>	5
RAID 6	The volume group contains no more than two drives in a single drawer.	3
RAID 3 or RAID 5	Each drive in the volume group is located in a separate drawer.	3

Level	Criteria for drawer loss protection	Minimum number of drawers required
RAID 1	Each drive in a mirrored pair must be located in a separate drawer.	2
RAID 0	Cannot achieve Drawer Loss Protection.	Not applicable

How do I maintain shelf/drawer loss protection?

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

Level	Criteria for shelf/drawer loss protection	Minimum number of shelves/drawers required
Pool	For shelves, the pool must contain no more than two drives in a single shelf. For drawers, the pool must include an equal number of drives from each drawer.	6 for shelves 5 for drawers
RAID 6	The volume group contains no more than two drives in a single shelf or drawer.	3
RAID 3 or RAID 5	Each drive in the volume group is located in a separate shelf or drawer.	3
RAID 1	Each drive in a mirrored pair must be located in a separate shelf or drawer.	2
RAID 0	Cannot achieve shelf/drawer loss protection.	Not applicable



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

What RAID level is best for my application?

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

Performance page to obtain these percentages.

RAID levels and application performance

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

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

Select a RAID level based on the following information.

RAID 0

- **Description**
 - Non-redundant, striping mode.
- **How it works**
 - RAID 0 stripes data across all of the drives in the volume group.
- **Data protection features**
 - RAID 0 is not recommended for high availability needs. RAID 0 is better for non-critical data.
 - If a single drive fails in the volume group, all of the associated volumes fail, and all data is lost.
- **Drive number requirements**
 - A minimum of one drive is required for RAID Level 0.
 - RAID 0 volume groups can have more than 30 drives.
 - You can create a volume group that includes all of the drives in the storage array.

RAID 1 or RAID 10

- **Description**
 - Striping/mirror mode.
- **How it works**
 - RAID 1 uses disk mirroring to write data to two duplicate disks simultaneously.
 - RAID 10 uses drive striping to stripe data across a set of mirrored drive pairs.
- **Data protection features**
 - RAID 1 and RAID 10 offer high performance and the best data availability.
 - RAID 1 and RAID 10 use drive mirroring to make an exact copy from one drive to another drive.
 - If one of the drives in a drive pair fails, the storage array can instantly switch to the other drive without any loss of data or service.
 - A single drive failure causes associated volumes to become degraded. The mirror drive allows access to the data.
 - A drive-pair failure in a volume group causes all of the associated volumes to fail, and data loss could occur.

- **Drive number requirements**

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

RAID 5

- **Description**

- High I/O mode.

- **How it works**

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

- **Data protection features**

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

- **Drive number requirements**

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

RAID 6

- **Description**

- High I/O mode.

- **How it works**

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

- **Data protection features**

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

- **Drive number requirements**

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



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

RAID levels and data protection

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

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

What is Data Assurance?

Data Assurance (DA) checks for and corrects errors that might occur as data is communicated between a host and a storage array. DA capabilities are presented at the pool and volume group level in System Manager.

The Data Assurance (DA) feature increases data integrity across the entire storage system. DA enables the storage array to check for errors that might occur when data is moved between the hosts and the drives. When this feature is enabled, the storage array appends error-checking codes (also known as cyclic redundancy checks or CRCs) to each block of data in the volume. After a data block is moved, the storage array uses these CRC codes to determine if any errors occurred during transmission. Potentially corrupted data is neither written to disk nor returned to the host.

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

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

What is secure-capable (Drive Security)?

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

What do I need to know about increasing reserved capacity?

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

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

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

- The volume in the pool or volume group must have an Optimal status and must not be in any state of

modification.

- Free capacity must exist in the pool or volume group that you want to use to increase capacity.
- You cannot increase reserved capacity for a snapshot volume that is read-only. Only snapshot volumes that are read-write require reserved capacity.

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

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

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

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

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

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

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

Why would I change this percentage?

Reserved capacity is typically 40 percent of the base volume for snapshot operations and 20 percent of the base volume for asynchronous mirroring operations. Usually this capacity is sufficient. The capacity needed varies, depending on the frequency and size of I/O writes to the base volume and how long you intend to use the storage object's copy service operation.

In general, choose a larger percentage for reserved capacity if one or both of these conditions exist:

- If the lifespan of a particular storage object's copy service operation will be very long.
- If a large percentage of data blocks will change on the base volume due to heavy I/O activity. Use historical performance data or other operating system utilities to help you determine typical I/O activity to the base volume.

Why do I need reserved capacity for each member volume?

Each member volume in a snapshot consistency group must have its own reserved capacity to save any modifications made by the host application to the base volume without affecting the referenced consistency group snapshot image. Reserved capacity

provides the host application with write access to a copy of the data contained in the member volume that is designated as read-write.

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

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

Why do I see more than one reserved capacity candidate?

You see more than one reserved capacity candidate when System Manager detects more than one volume in a pool or volume group that meets the capacity percentage amount you selected for the storage object.

You can choose to refresh the list of recommended candidates by changing the percentage of physical drive space that you want to reserve on the base volume for copy service operations. System Manager displays the best reserved capacity candidates based on your selection.

How do I view and interpret all SSD Cache statistics?

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

Nominal statistics

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

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

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

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

Detailed statistics

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

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

Detailed statistics	Description
Read Blocks	The number of blocks in host reads.

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

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

Volumes

Concepts

Volumes in the storage array

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

Volumes are the only data layer visible to hosts. In a SAN environment, volumes are mapped to logical unit

numbers (LUNs), which are visible to hosts. LUNs hold the user data that is accessible using one or more of the host access protocols supported by the storage array, including FC, iSCSI, and SAS.

Volume types you can create from pools and volume groups

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

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



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

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

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

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

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

Characteristics of volumes

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

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

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

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

name descriptive.

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

Volume terminology

Learn how the volume terms apply to your storage array.

All volume types

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

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

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

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

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

Specific to thin volumes




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

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

Workflow for creating volumes

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



Data integrity and data security for volumes

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

Data Assurance

The Data Assurance (DA) feature increases data integrity across the entire storage system. DA enables the storage array to check for errors that might occur when data is moved between the hosts and the drives. When this feature is enabled, the storage array appends error-checking codes (also known as cyclic redundancy checks or CRCs) to each block of data in the volume. After a data block is moved, the storage array uses these CRC codes to determine if any errors occurred during transmission. Potentially corrupted data is neither written to disk nor returned to the host.

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

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

Drive Security

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

How Drive Security works at the drive level

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

How Drive Security works at the volume level

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

How to implement Drive Security

To implement Drive Security, you perform the following steps.

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

With the Drive Security feature, you create a security key that is shared between the secure-enabled drives

and controllers in a storage array. Whenever power to the drives is turned off and on, the secure-enabled drives change to a Security Locked state until the controller applies the security key.

SSD Cache and volumes

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

Volumes

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

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

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

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



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

Drive Security feature

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

Application-specific workloads

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

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

Types of workloads

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

- **Application-specific.** When you are creating volumes using an application-specific workload, the system may recommend an optimized volume configuration to minimize contention between application workload I/O and other traffic from your application instance. Volume characteristics like I/O type, segment size, controller ownership, and read and write cache are automatically recommended and optimized for workloads that are created for the following application types.
 - Microsoft® SQL Server™
 - Microsoft® Exchange Server™
 - Video Surveillance applications
 - VMware ESXi™ (for volumes to be used with Virtual Machine File System) You can review the recommended volume configuration and edit, add, or delete the system-recommended volumes and characteristics using the **Add/Edit Volumes** dialog box.
- **Other** (or applications without specific volume creation support). Other workloads use a volume configuration that you must manually specify when you want to create a workload that is not associated with a specific application, or if System Manager does not have built-in optimization for the application you intend to use on the storage array. You must manually specify the volume configuration using the **Add/Edit Volumes** dialog box.

Application and workload views

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

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

Actions you can perform on volumes

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

Increase capacity

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

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

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

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

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

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

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

Delete

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

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

Copy

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

Initialize

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

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

Redistribute

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

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

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

Change volume ownership

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

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

Mirroring and volume ownership

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

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

Copy volume and volume ownership

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

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

Change cache settings

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

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

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

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

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

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

battery is replaced.

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

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

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

Change media scan settings

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

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

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

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

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

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

Capacity for volumes

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

Using capacity to create and expand volumes

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

- When you create a volume from unassigned capacity, you can create a pool or volume group and the volume at the same time.

- When you create a volume from free capacity, you are creating an additional volume on an already existing pool or volume group.

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

Capacity types for thick volumes and thin volumes

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

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

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

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

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



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

Capacity limits for thick volumes

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

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

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

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

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

Capacity limits for thin volumes

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

In addition to reported capacity and allocated capacity, thin volumes also contain Written capacity. Written

capacity is the amount of capacity that has been written from the reserved capacity allocated for thin volumes.

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

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

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

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

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

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

Thin volume monitoring

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

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

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

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

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



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

Comparison between thick volumes and thin volumes

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

means that the capacity is allocated as the data is being written to the volume.

When to use a thick or thin volume

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



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

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

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

Thin volume restrictions

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

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

Copy Volume function

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

Data copying for greater access

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

- Move data to larger drives.

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

Change a thin volume to a thick volume

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



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

Backup data

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

Restore snapshot volume data to the base volume

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

Source and target volumes

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

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

Types of Copy Volume operations

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

volume. An online operation uses a snapshot volume as the source and copies its data to a target volume.

To ensure data integrity, all I/O activity to the target volume is suspended during either type of Copy Volume operation. This suspension occurs because the state of data on the target volume is inconsistent until the procedure is complete.

The offline and online Copy Volume operations are described below.

Offline Copy Volume operation

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

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

Online Copy Volume operation

The online Copy Volume relationship is between a snapshot volume and a target volume. You can initiate a Copy Volume operation while the source volume is online and available for data writes. This function is

achieved by creating a snapshot of the volume and using the snapshot as the actual source volume for the copy.

When you initiate a Copy Volume operation for a source volume, System Manager creates a snapshot image of the base volume and a copy relationship between the snapshot image of the base volume and a target volume. Using the snapshot image as the source volume allows the storage array to continue to write to the source volume while the copy is in progress.

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

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

What you need to know about online copy operations

Reserved capacity volume

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

How to

Create storage

Create workloads

You can create workloads for any type of application.

About this task

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

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

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

Keep these guidelines in mind:

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

Steps

1. Select **Storage > Volumes**.

2. Select **Create › Workload**.

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

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

After you finish

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

Create volumes

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

About this task

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

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

- [Step 1: Select host](#)
- [Step 2: Select workload](#)
- [Step 3: Add or edit volumes](#)
- [Step 4: Review configuration](#)



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

Step 1: Select host

You select a specific host or host cluster to assign it to a volume. This assignment grants a host or a host cluster access to one or more volumes for I/O operations. You can choose to assign a host later, if needed.

Before you begin

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

About this task

Keep these guidelines in mind when you assign volumes:

- A host's operating system can have specific limits on how many volumes the host can access. Keep this limitation in mind when you create volumes for use by a particular host.
- You can define one assignment for each volume in the storage array.
- Assigned volumes are shared between controllers in the storage array.
- The same logical unit number (LUN) cannot be used twice by a host or a host cluster to access a volume. You must use a unique LUN.



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

Steps

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

The **Create Volumes** dialog box appears.

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

The Select Workload dialog box appears.

Step 2: Select workload

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

About this task

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

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

Steps

1. Do one of the following:
 - Select the **Create volumes for an existing workload** option to create volumes for an existing workload.
 - Select the **Create a new workload** option to define a new workload for a supported application or for

"Other" applications.

- From the drop-down list, select the name of the application you want to create the new workload for.

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

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

2. Click **Next**.

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

Step 3: Add or edit volumes

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

Before you begin

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

Selecting a DA capable pool or volume group

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

DA capabilities are presented at the pool and volume group level in System Manager. DA protection checks for and corrects errors that might occur as data is communicated between a host and a storage array. Selecting a DA-capable pool or volume group for the new volume ensures that any errors are detected and corrected.

If any of the host connections on the controllers in your storage array do not support DA, the associated hosts cannot access data on DA-enabled volumes. DA is not supported by iSCSI over TCP/IP, or by the SRP over InfiniBand.

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

Selecting a secure-capable pool or volume group

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

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

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

About this task

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

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

Steps

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

Field Details

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

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

Field	Description
Secure-capable	<p>Yes appears next to "Secure-capable" only if the drives in the pool or volume group are secure-capable.</p> <p>Drive Security prevents unauthorized access to the data on a drive that is physically removed from the storage array. This option is available only when the Drive Security feature has been enabled, and a security key is set up for the storage array.</p> <p>A pool or volume group can contain both secure-capable and non-secure-capable drives, but all drives must be secure-capable to use their encryption capabilities.</p>
DA	<p>Yes appears next to "DA" only if the drives in the pool or volume group support Data Assurance (DA).</p> <p>DA increases data integrity across the entire storage system. DA enables the storage array to check for errors that might occur when data is moved between the hosts and the drives. Using DA for the new volume ensures that any errors are detected.</p>

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

Field Details

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

Field	Description
Segment Size	<p data-bbox="867 155 1448 289">Shows the setting for segment sizing, which only appears for volumes in a volume group. You can change the segment size to optimize performance.</p> <p data-bbox="867 323 1448 667">Allowed segment size transitions — System Manager determines the segment size transitions that are allowed. Segment sizes that are inappropriate transitions from the current segment size are unavailable on the drop-down list. Allowed transitions usually are double or half of the current segment size. For example, if the current volume segment size is 32 KiB, a new volume segment size of either 16 KiB or 64 KiB is allowed.</p> <p data-bbox="867 701 1448 1045">SSD Cache-enabled volumes — You can specify a 4-KiB segment size for SSD Cache-enabled volumes. Make sure you select the 4-KiB segment size only for SSD Cache-enabled volumes that handle small-block I/O operations (for example, 16 KiB I/O block sizes or smaller). Performance might be impacted if you select 4 KiB as the segment size for SSD Cache-enabled volumes that handle large block sequential operations.</p> <p data-bbox="867 1079 1448 1213">Amount of time to change segment size — The amount of time to change a volume's segment size depends on these variables:</p> <ul data-bbox="889 1247 1448 1591" style="list-style-type: none"> • The I/O load from the host • The modification priority of the volume • The number of drives in the volume group • The number of drive channels • The processing power of the storage array controllers When you change the segment size for a volume, I/O performance is affected, but your data remains available.

Field	Description
Secure-capable	<p>Yes appears next to "Secure-capable" only if the drives in the pool or volume group are secure-capable.</p> <p>Drive security prevents unauthorized access to the data on a drive that is physically removed from the storage array. This option is available only when the drive security feature has been enabled, and a security key is set up for the storage array.</p> <p>A pool or volume group can contain both secure-capable and non-secure-capable drives, but all drives must be secure-capable to use their encryption capabilities.</p>
DA	<p>Yes appears next to "DA" only if the drives in the pool or volume group support Data Assurance (DA).</p> <p>DA increases data integrity across the entire storage system. DA enables the storage array to check for errors that might occur when data is moved between the hosts and the drives. Using DA for the new volume ensures that any errors are detected.</p>

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

Step 4: Review configuration

You can review a summary of the volumes you intend to create and make any necessary changes.

Steps

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

Results

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

After you finish

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

names.

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

Manage volumes

Increase capacity of a volume

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

Before you begin

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

About this task

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



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

Steps

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

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

3. Select **Yes** to continue.

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

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

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

Results

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

performance.

After you finish

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

Change settings for a volume

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

Before you begin

The volume you want to change is in Optimal status.

Steps


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

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

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

Field Details

Setting	Description
Name	Displays the name of the volume. Change the name of a volume when the current name is no longer meaningful or applicable.
Capacities	<p>Displays the reported and allocated capacity for the selected volume.</p> <p>Reported capacity and allocated capacity are the same for thick volumes, but are different for thin volumes. For a thick volume, the physically allocated space is equal to the space that is reported to the host. For a thin volume, reported capacity is the capacity that is reported to the hosts, whereas allocated capacity is the amount of drive space that is currently allocated for writing data.</p>
Pool / Volume group	Displays the name and RAID level of the pool or volume group. Indicates whether the pool or volume group is secure-capable and secure-enabled.

Setting	Description
Host	<p data-bbox="841 157 1453 359">Displays the volume assignment. You assign a volume to a host or host cluster so it can be accessed for I/O operations. This assignment grants a host or host cluster access to a particular volume or to a number of volumes in a storage array.</p> <ul data-bbox="867 394 1453 751" style="list-style-type: none"> <li data-bbox="867 394 1453 499">• Assigned to — Identifies the host or host cluster that has access to the selected volume. <li data-bbox="867 512 1453 751">• LUN — A logical unit number (LUN) is the number assigned to the address space that a host uses to access a volume. The volume is presented to the host as capacity in the form of a LUN. Each host has its own LUN address space. Therefore, the same LUN can be used by different hosts to access different volumes. <div data-bbox="922 787 1453 1276">  <p data-bbox="1036 793 1414 1270">For NVMe interfaces, this column displays Namespace ID. A namespace is NVM storage that is formatted for block access. It is analogous to a logical unit in SCSI, which relates to a volume in the storage array. The namespace ID is the NVMe controller's unique identifier for the namespace, and can be set to a value between 1 and 255. It is analogous to a logical unit number (LUN) in SCSI.</p> </div>
Identifiers	<p data-bbox="841 1354 1425 1381">Displays the identifiers for the selected volume.</p> <ul data-bbox="867 1417 1453 1648" style="list-style-type: none"> <li data-bbox="867 1417 1453 1480">• World-wide identifier (WWID) — A unique hexadecimal identifier for the volume. <li data-bbox="867 1495 1453 1558">• Extended unique identifier (EUI) — An EUI-64 identifier for the volume. <li data-bbox="867 1572 1453 1648">• Subsystem identifier (SSID) — The storage array subsystem identifier of a volume.

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

Field Details

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

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

Setting	Description
Modification priority	<p>Shows the setting for modification priority, which only appears for volumes in a volume group.</p> <p>The modification priority defines how much processing time is allocated for volume modification operations relative to system performance. You can increase the volume modification priority, although this might affect system performance.</p> <p>Move the slider bars to select a priority level.</p> <p>Modification priority rates — The lowest priority rate benefits system performance, but the modification operation takes longer. The highest priority rate benefits the modification operation, but system performance might be compromised.</p>
Caching	Shows the caching setting, which you can change to impact the overall I/O performance of a volume.
SSD Cache	<p>Shows the SSD Cache setting, which you can enable on compatible volumes as a way to improve read-only performance. Volumes are compatible if they share the same Drive Security and Data Assurance capabilities.</p> <p>The SSD Cache feature uses a single or multiple Solid State Disks (SSDs) to implement a read cache. Application performance is improved because of the faster read times for SSDs. Because the read cache is in the storage array, caching is shared across all applications using the storage array. Simply select the volume that you want to cache, and then caching is automatic and dynamic.</p>

5. Click **Save**.

Result

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

After you finish

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

Initialize volumes

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

Before you begin

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



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

About this task

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

The following types of volumes *cannot* be initialized:

- Base volume of a snapshot volume
- Primary volume in a mirror relationship
- Secondary volume in a mirror relationship
- Source volume in a volume copy
- Target volume in a volume copy
- Volume that already has an initialization in progress

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

Steps

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

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

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

Results

System Manager performs the following actions:

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

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

Redistribute volumes

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

Before you begin

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

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

About this task

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

Steps

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

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

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

Results

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

Change controller ownership of a volume

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

Before you begin

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

About this task

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

Steps

1. Select **Storage > Volumes**.

2. Select any volume, and then select **More › Change ownership**.

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

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

Results

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

Change cache settings for a volume

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

About this task

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

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

This action might occur when a new battery is charging, when a controller has been removed, or if a mismatch in cache sizes has been detected by the controller. After the condition has cleared, the cache properties selected in the dialog box become active. If the selected cache properties do not become active, contact technical support.

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


Steps

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

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


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

Field Details

Cache setting	Description
Read Caching	The read cache is a buffer that stores data that has been read from the drives. The data for a read operation might already be in the cache from a previous operation, which eliminates the need to access the drives. The data stays in the read cache until it is flushed.
Write Caching	<div><p>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.</p><div><p>Cache is automatically flushed after the Write caching is disabled for a volume.</p></div></div>

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

Field Details

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

Cache setting	Description
Write Caching with Mirroring	<p>Write caching with mirroring occurs when the data written to the cache memory of one controller is also written to the cache memory of the other controller. Therefore, if one controller fails, the other can complete all outstanding write operations. Write cache mirroring is available only if write caching is enabled and two controllers are present. Write caching with mirroring is the default setting at volume creation.</p> <p>This setting is available only if you enabled write caching. This setting is not available for thin volumes.</p>

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

Change media scan settings for a volume

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

Before you begin

Understand the following:

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

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

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

About this task

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

You can perform the following actions:

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

Steps

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

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

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

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

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

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

6. Click **Save**.

Results

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

Delete volume

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

Before you begin

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

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

About this task

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

- The volume is initializing.
- The volume is reconstructing.
- The volume is part of a volume group that contains a drive that is undergoing a copyback operation.

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



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



When a volume exceeds a given size (currently 64TB) the delete is being performed in background and the freed space may not be immediately available.

Steps

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

The **Delete Volumes** dialog box appears.

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

Results

System Manager performs the following actions:

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

Manage applications and workloads

Add to workload

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

About this task

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

Steps

1. Select **Storage > Volumes**.
2. Select the **Applications & Workloads** tab.

The Applications & Workloads view appears.

3. Select **Add to Workload**.

The Select Workload dialog box appears.

4. Do one of the following actions:

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

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

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

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

The Select Volumes dialog box appears.

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

Change workload settings

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

Steps

1. Select **Storage > Volumes**.
2. Select the **Applications & Workloads** tab.

The **Applications & Workloads** view appears.

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

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

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

Work with copy services

Copy volume

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

Before you begin

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

About this task

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

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

More about online and offline copies

Online copy

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

Offline copy

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



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

Steps

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

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

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

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

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

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

More about copy priority rates

There are five copy priority rates:

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

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

The volume candidate table displays only the candidates that support the reserved capacity specified. Reserved capacity is the physical allocated capacity that is used for any copy service operation and storage object. It is not directly readable by the host.

Allocate the reserved capacity using the following guidelines:

- The default setting for reserved capacity is 40% of the capacity of the base volume, and usually this capacity is sufficient.
- Reserved capacity, however, varies depending on the number of changes to the original data. The longer a storage object is active, the larger the reserved capacity should be.

Results

System Manager copies all data from the source volume to the target volume. After the Copy Volume operation is complete, the target volume automatically becomes read-only to the hosts.

After you finish

Select **Home > View Operations in Progress** to view the progress of the Copy Volume operation. This operation can be lengthy and could affect system performance.

Take action on a Copy Volume operation

You can view a Copy Volume operation in progress and stop, change priority, re-copy, or clear a Copy Volume operation.


Steps

1. Select **Home > View Operations in Progress**.

The **Operations in Progress** dialog box appears.

2. Find the Copy Volume operation that you want to take action on, and then click the link in the **Actions** column to take one of the following actions.

Read all cautionary text provided in dialogs, particularly when stopping an operation.

Action	Description
Stop	<p>You can stop a Copy Volume operation while the operation has a status of In Progress, Pending, or Failed.</p> <p>When the Copy Volume is stopped, all of the mapped hosts have write access to the source volume. If data is written to the source volume, the data on the target volume no longer matches the data on the source volume.</p>
Change priority	<p>You can change the priority of a Copy Volume operation while the operation has a status of In Progress to select the rate at which a Copy Volume operation completes.</p>
Re-copy	<p>You can re-copy a volume when you have stopped a Copy Volume operation and want to start it again or when a Copy Volume operation has failed or halted. The Copy Volume operation starts over from the beginning.</p> <p>The re-copy action overwrites existing data on the target volume and fails all snapshot volumes associated with the target volume, if any exist.</p>
Clear	<p>You can remove the Copy Volume operation while the operation has a status of In Progress, Pending, or Failed.</p> <div><p>Be sure that you want to do this operation before selecting Clear. There is no confirmation dialog.</p></div>

Create asynchronous mirrored volume

You mirror a volume asynchronously to maintain data at the remote storage array to be a point-in-time consistent copy of data at the local storage array. You do this by creating a mirror consistency group to establish the mirroring relationship between the two storage arrays, and then selecting the primary volume and secondary volume that you want to use in the mirror.

Before you begin

- The following conditions must be set up:

- The Web Services Proxy service is running.
- SANtricity Unified Manager is running on your local host through an HTTPS connection.
- Each controller in both the primary array and secondary array must have an Ethernet management port configured and must be connected to your network.
- SANtricity Unified Manager is showing valid SSL certificates for the storage array. You can accept a self-signed certificate or install your own security certificate using Unified Manager and navigating to **Certificate > Certificate Management**.
- SANtricity System Manager is launched from a Unified Manager.
- You must have discovered the two storage arrays you want to mirror data between. Then, from Unified Manager, you select the primary volume's storage array and click **Launch** to open the browser-based SANtricity System Manager.
- You must know the password for the local and remote storage arrays.
- Your local and remote storage arrays must be connected through a Fibre Channel fabric or iSCSI interface.

About this task

The process to mirror a volume asynchronously is a multi-step procedure:

- [Step 1: Create a mirror consistency group or select an existing one](#)
- [Step 2: Select the primary volume](#)
- [Step 3: Select the secondary volume](#)

A volume can participate in only one mirror relationship.

Step 1: Create a mirror consistency group or select an existing one

You create a mirror consistency group or select an existing one to establish the mirroring relationship between the local storage array and the remote storage array.

About this task

The number of mirror consistency group relationships and mirrored pair relationships that you can create depends on the hardware in your storage array.

Steps

1. Do one of the following actions to access the asynchronous mirroring sequence:
 - Select **Storage > Asynchronous Mirroring > Create Mirrored pair**.
 - Select **Storage > Volumes > Copy Services > Mirror a volume asynchronously**.
2. Either select an existing mirror consistency group or create a new one.

To create new mirror consistency group, do the following:

- a. Enter a unique name that best describes the data on the volumes that will be mirrored between the two storage arrays (for example, R&D Data).
- b. Select the remote storage array on which you want to establish a mirror relationship with the local storage array.



If your remote storage array is password protected, the system prompts for a password.

- c. Choose whether you want to resynchronize the mirrored pairs on the remote storage array either manually or automatically.
 - **Manual** — You must explicitly update the secondary point-in-time image using the Manual Resynchronization menu option. Select this option to manually start resynchronization for all asynchronous mirrored pairs within the asynchronous mirror group.
 - **Automatic** — Using the drop-down, specify the time from the beginning of the previous update to the beginning of the next update. To change the automatic synchronization interval from the default of every 10 minutes, edit the interval value, which is defined in minutes.
- d. Click **Create**.

System Manager creates the mirror consistency group on the local storage array first and then creates the mirror consistency group on the remote storage array.



If System Manager successfully creates the mirror consistency group on the local storage array, but fails to create it on the remote storage array, it automatically deletes the mirror consistency group from the local storage array. If an error occurs while System Manager is attempting to delete the mirror consistency group, you must manually delete it.

3. Select **Next** and go to [Step 2: Select the primary volume](#).

Step 2: Select the primary volume

You must select the primary volume that you want to use in the mirror relationship and allocate its reserved capacity. Any volumes added to the mirror consistency group on the local storage array will hold the primary role in the mirror relationship.

Steps

1. Select an existing volume that you want to use as the primary volume in the mirror, and then click **Next** to allocate the reserved capacity.
2. Allocate the reserved capacity for the primary volume you selected. Do one of the following actions:
 - **Accept the default settings** — Use this recommended option to allocate the reserved capacity for the primary volume with the default settings.
 - **Allocate your own reserved capacity settings to meet your data storage needs related to asynchronous mirroring** — Allocate the reserved capacity using the following guidelines.
 - The default setting for reserved capacity is 20% of the capacity of the base volume, and usually this capacity is sufficient.
 - The capacity needed varies, depending on the frequency and size of I/O writes to the primary volume and how long you need to keep the capacity.
 - In general, choose a larger capacity for reserved capacity if one or both of these conditions exist:
 - You intend to keep the mirrored pair for a long period of time.
 - A large percentage of data blocks will change on the primary volume due to heavy I/O activity. Use historical performance data or other operating system utilities to help you determine typical I/O activity to the primary volume.
3. Select **Next** and go to [Step 3: Select the secondary volume](#).

Step 3: Select the secondary volume

You must select the secondary volume that you want to use in the mirror relationship and allocate its reserved capacity. Any volumes added to the mirror consistency group on the remote storage array will hold the secondary role in the mirror relationship.

About this task

When you select a secondary volume on the remote storage array, the system displays a list of all the eligible volumes for that mirrored pair. Any volumes that are not eligible to be used do not display in that list.

Steps

1. Select an existing volume that you want to use as the secondary volume in the mirrored pair, and then click **Next** to allocate the reserved capacity.
2. Allocate the reserved capacity for the secondary volume you selected. Do one of the following actions:
 - **Accept the default settings** — Use this recommended option to allocate the reserved capacity for the secondary volume with the default settings.
 - **Allocate your own reserved capacity settings to meet your data storage needs related to asynchronous mirroring** — Allocate the reserved capacity using the following guidelines.
 - The default setting for reserved capacity is 20% of the capacity of the base volume, and usually this capacity is sufficient.
 - The capacity needed varies, depending on the frequency and size of I/O writes to the primary volume and how long you need to keep the capacity.
 - In general, choose a larger capacity for reserved capacity if one or both of these conditions exist:
 - You intend to keep the mirrored pair for a long period of time.
 - A large percentage of data blocks will change on the primary volume due to heavy I/O activity. Use historical performance data or other operating system utilities to help you determine typical I/O activity to the primary volume.
3. Select **Finish** to complete the asynchronous mirroring sequence.

Results

System Manager performs the following actions:

- Begins initial synchronization between the local storage array and the remote storage array.
- If the volume being mirrored is a thin volume, only the provisioned blocks (allocated capacity rather than reported capacity) are transferred to the secondary volume during the initial synchronization. This reduces the amount of data that must be transferred to complete the initial synchronization.
- Creates the reserved capacity for the mirrored pair on the local storage array and on the remote storage array.

Create synchronous mirrored volume

You mirror a volume synchronously to replicate data in real-time between storage arrays, so your information is protected from both system and site failures. You do this by selecting the primary volume and the secondary volume that you want to use in the synchronous mirroring relationship between a local storage array and a remote storage array.

Before you begin

- Because the Synchronous Mirroring feature requires the management of multiple storage arrays, you must have the browser-based SANtricity Unified Manager installed, and have discovered the two storage arrays you want to mirror data between. Then, from Unified Manager, you select the primary volume's storage array and click **Launch** to open the browser-based SANtricity System Manager.
- You must have two storage arrays.
- Each storage array must have two controllers.
- The primary and secondary volumes' storage arrays can run different OS versions. The minimum version supported is 7.84.
- You must know the password for the local and remote storage arrays.
- Your local and remote storage arrays must be connected through a Fibre Channel fabric.
- You must have created both the primary and secondary volumes that you want to use in the synchronous mirror relationship.

About this task

The process to mirror a volume synchronously is a multi-step procedure:

- [Step 1: Select the primary volume](#)
- [Step 2: Select the secondary volume](#)
- [Step 3: Select synchronization settings](#)

A volume can participate in only one mirror relationship.

Step 1: Select the primary volume

You must select the primary volume that you want to use in the synchronous mirror relationship. This volume holds the primary role in the mirror relationship.

Before you begin

- You must have created the primary volume that you want to use in the synchronous mirror relationship.
- The primary volume must be a standard volume. It cannot be a thin volume or a snapshot volume.

Steps

1. Do one of the following actions to access the synchronous mirroring sequence:
 - Select **Storage > Synchronous Mirroring > Mirror volume**.
 - Select **Storage > Volumes > Copy Services > Mirror a volume synchronously**. The **Create Synchronous Mirrored Pair** dialog appears.
2. Select an existing volume that you want to use as the primary volume in the mirror.



If a volume was selected in the Volumes tile and it is eligible to be mirrored, it will be selected by default.

3. Select **Next** and go to [Step 2: Select the secondary volume](#).

Step 2: Select the secondary volume

You must select the secondary volume that you want to use in the mirror relationship.

This volume will hold the secondary role in the mirror relationship.

Before you begin

- You must have created the secondary volume that you want to use in the synchronous mirror relationship.
- The secondary volume must be a standard volume. It cannot be a thin volume or a snapshot volume.
- The secondary volume must be at least as large as the primary volume.

About this task

When you select a secondary volume on the remote storage array, the system displays a list of all the eligible volumes for that mirrored pair. Any volumes that are not eligible to be used do not display in that list.

The volumes that appear in this dialog are sorted by capacity, starting with the volume nearest to the capacity of the primary volume capacity. Volumes with identical capacity are sorted alphabetically.

Steps

1. Select the remote storage array on which you want to establish a mirror relationship with the local storage array.



If your remote storage array is password protected, the system prompts for a password.

- Storage arrays are listed by their storage array name. If you have not named a storage array, it will be listed as "unnamed."
- If the storage array you want to use is not in the list, add it using the Enterprise Management Window (EMW) of SANtricity Storage Manager. Select **Edit > Add Storage Array**.

2. Select an existing volume that you want to use as the secondary volume in the mirror.



If a secondary volume is chosen with a capacity that is larger than the primary volume, the usable capacity is restricted to the size of the primary volume.

3. Click **Next** and go to [Step 3: Select synchronization settings](#).

Step 3: Select synchronization settings

You must set the priority at which the controller owner of the primary volume resynchronizes data with the secondary volume after a communication interruption. You must also select the resynchronization policy, either manual or automatic.

Steps

1. Use the slider bar to set the synchronization priority.

The synchronization priority determines how much of the system resources are used to complete initial synchronization and the resynchronization operation after a communication interruption as compared to service I/O requests.

The priority set on this dialog applies to both the primary volume and the secondary volume. You can modify the rate on the primary volume at a later time by selecting **Storage > Synchronous Mirroring > More > Edit Settings**.

More about synchronization rates

There are five synchronization priority rates:

- Lowest
- Low
- Medium
- High
- Highest If the synchronization priority is set to the lowest rate, I/O activity is prioritized, and the resynchronization operation takes longer. If the synchronization priority is set to the highest rate, the resynchronization operation is prioritized, but I/O activity for the storage array might be affected.

2. Choose whether you want to resynchronize the mirrored pairs on the remote storage array either manually or automatically.
 - **Manual** (the recommended option) — Select this option to require synchronization to be manually resumed after communication is restored to a mirrored pair. This option provides the best opportunity for recovering data.
 - **Automatic** — Select this option to start resynchronization automatically after communication is restored to a mirrored pair. To manually resume synchronization go to **Storage > Synchronous Mirroring**, highlight the mirrored pair in the table, and select **Resume** under **More**.
3. Click **Finish** to complete the synchronous mirroring sequence.

Results

System Manager performs the following actions:

- Activates the Synchronous Mirroring feature.
- Begins initial synchronization between the local storage array and the remote storage array.
- Sets the synchronization priority and resynchronization policy.

After you finish

Select **Home > View Operations in Progress** to view the progress of the synchronous mirroring operation. This operation can be lengthy and could affect system performance.

Create snapshot image

You can manually create a snapshot image from a base volume or snapshot consistency group. This is also called an *instant snapshot* or *instant image*.

Before you begin

- The base volume must be optimal.
- The drive must be optimal.
- The snapshot group cannot be designated as “reserved.”
- The reserved capacity volume must have the same Data Assurance (DA) settings as the associated base volume for the snapshot group.

Steps

1. Do one of the following actions to create a snapshot image:
 - Select **Storage › Volumes**. Select the object (base volume or snapshot consistency group), and then select **Copy Services › Create instant snapshot**.
 - Select **Storage › Snapshots**. Select the **Snapshot Images** tab, and then select **Create › Instant snapshot image**. The **Create Snapshot Image** dialog box appears. Select the object (base volume or snapshot consistency group), and then click **Next**. If a previous snapshot image was created for the volume or snapshot consistency group, then the system creates the instant snapshot immediately. Otherwise, if this is the first time a snapshot image is created for the volume or snapshot consistency group, the **Confirm Create Snapshot Image** dialog box appears.
2. Click **Create** to accept the notification that reserved capacity is needed and to proceed to the **Reserve Capacity** step.

The **Reserve Capacity** dialog box appears.

3. Use the spinner box to adjust the capacity percentage, and then click **Next** to accept the candidate volume highlighted in the table.

The **Edit Settings** dialog box appears.

4. Select the settings for the snapshot image as appropriate, and confirm that you want to perform the operation.

Field Details

Setting	Description
Snapshot image settings	
Snapshot image limit	Keep the check box selected if you want snapshot images automatically deleted after the specified limit; use the spinner box to change the limit. If you clear this check box, snapshot image creation stops after 32 images.
Reserved capacity settings	
Alert me when...	<p>Use the spinner box to adjust the percentage point at which the system sends an alert notification when the reserved capacity for a snapshot group is nearing full.</p> <p>When the reserved capacity for the snapshot group exceeds the specified threshold, use the advance notice to increase reserved capacity or to delete unnecessary objects before the remaining space runs out.</p>
Policy for full reserved capacity	<p>Choose one of the following policies:</p> <ul style="list-style-type: none">• Purge oldest snapshot image: The system automatically purges the oldest snapshot image in the snapshot group, which releases the snapshot image reserved capacity for reuse within the group.• Reject writes to base volume: When the reserved capacity reaches its maximum defined percentage, the system rejects any I/O write request to the base volume that triggered the reserved capacity access.

Results

- System Manager displays the new snapshot image in the Snapshot Images table. The table lists the new image by timestamp and associated base volume or snapshot consistency group.
- Snapshot creation might remain in a Pending state because of the following conditions:
 - The base volume that contains this snapshot image is a member of an asynchronous mirror group.
 - The base volume is currently in a synchronization operation. The snapshot image creation completes as soon as the synchronization operation is complete.

Schedule snapshot images

You create a snapshot schedule to enable recovery in case of a problem with the base volume and to perform scheduled backups. Snapshots of base volumes or snapshot consistency groups can be created on a daily, weekly, or monthly schedule, at any time of day.

Before you begin

The base volume must be Optimal.

About this task

This task describes how to create a snapshot schedule for an existing snapshot consistency group or base volume.



You also can create a snapshot schedule at the same time you create a snapshot image of a base volume or snapshot consistency group.

Steps

1. Do one of the following actions to create a snapshot schedule:

- Select **Storage > Volumes**.

Select the object (volume or snapshot consistency group) for this snapshot schedule, and then select **Copy Services > Create snapshot schedule**.

- Select **Storage > Snapshots**.

Select the **Schedules** tab, and then click **Create**.

2. Select the object (volume or snapshot consistency group) for this snapshot schedule, and then click **Next**.

The **Create Snapshot Schedule** dialog box appears.

3. Do one of the following actions:

- **Use a previously defined schedule from another snapshot object.**

Make sure advanced options are displayed. Click **Show more options**. Click **Import Schedule**, select the object with the schedule you want to import, and then click **Import**.

- **Modify the basic or advanced options.**

In the upper right of the dialog box, click **Show more options** to display all options, and then refer to the following table.

Field Details

Field	Description
Basic settings	
Select days	Select individual days of the week for snapshot images.
Start time	From the drop-down list, select a new start time for the daily snapshots (selections are provided in half-hour increments). The start time defaults to one half-hour ahead of the current time.
Time zone	From the drop-down list, select your array's time zone.
Advanced settings	
Day / month	<p>Choose one of the following options:</p> <ul style="list-style-type: none"> • Daily / Weekly — Select individual days for synchronization snapshots. You also can select the Select all days check box in the upper right if you want a daily schedule. • Monthly / Yearly — Select individual months for synchronization snapshots. In the On day(s) field, enter the days of the month for synchronizations to occur. Valid entries are 1 through 31 and Last. You can separate multiple days with a comma or semi-colon. Use a hyphen for inclusive dates. For example: 1,3,4,10-15,Last. You also can select the Select all months check box in the upper right if you want a monthly schedule.
Start time	From the drop-down list, select a new start time for the daily snapshots (selections are provided in half-hour increments). The start time defaults to one half-hour ahead of the current time.
Time zone	From the drop-down list, select your array's time zone.
Snapshots per day / Time between snapshots	Select the number of snapshot images to create per day. If you select more than one, also select the time between snapshot images. For multiple snapshot images, be sure that you have adequate reserved capacity.

Field	Description
Create snapshot image right now?	Select this check box to create an instant image in addition to the automatic images you are scheduling.
Start/End date or No end date	Enter the start date for synchronizations to begin. Also enter an end date or select No end date .

4. Do one of the following actions:

- If the object is a snapshot consistency group, click **Create** to accept the settings and create the schedule.
- If the object is a volume, click **Next** to allocate reserved capacity for the snapshot images.

The volume candidate table displays only the candidates that support the reserved capacity specified. Reserved capacity is the physical allocated capacity that is used for any copy service operation and storage object. It is not directly readable by the host.

5. Use the spinner box to allocate the reserved capacity for the snapshot images. Do one of the following actions:

- **Accept the default settings.**

Use this recommended option to allocate the reserved capacity for the snapshot images with the default settings.

- **Allocate your own reserved capacity settings to meet your data storage needs.**

If you change the default reserved capacity setting, click **Refresh Candidates** to refresh the candidate list for the reserved capacity you specified.

Allocate the reserved capacity using the following guidelines:

- The default setting for reserved capacity is 40% of the capacity of the base volume. Usually this capacity is sufficient.
- The capacity needed varies, depending on the frequency and size of I/O writes to the volumes and the quantity and duration of snapshot image collection.

6. Click **Next**.

The Edit Settings dialog box appears.

7. Edit the settings for the snapshot schedule as needed, and then click **Finish**.

Field Details

Setting	Description
Snapshot image limit	
Enable automatic deletion of snapshot images when...	Keep the check box selected if you want snapshot images automatically deleted after the specified limit; use the spinner box to change the limit. If you clear this check box, snapshot image creation stops after 32 images.
Reserved capacity settings	
Alert me when...	<p>Use the spinner box to adjust the percentage point at which the system sends an alert notification when the reserved capacity for a schedule is nearing full.</p> <p>When the reserved capacity for the schedule exceeds the specified threshold, use the advance notice to increase reserved capacity or to delete unnecessary objects before the remaining space runs out.</p>
Policy for full reserved capacity	<p>Choose one of the following policies:</p> <ul style="list-style-type: none">• Purge oldest snapshot image — The system automatically purges the oldest snapshot image, which releases the snapshot image reserved capacity for reuse within the snapshot group.• Reject writes to base volume — When the reserved capacity reaches its maximum defined percentage, the system rejects any I/O write request to the base volume that triggered the reserved capacity access.

Change allocated capacity limit for a thin volume

For thin volumes capable of allocating space on demand, you can change the limit that restricts the allocated capacity to which a thin volume can automatically expand. You also can change the percentage point at which an alert (warning threshold exceeded) is sent to the Notifications area on the Home page when a thin volume is near the allocated capacity limit. You can choose to enable or disable this alert notification.



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

About this task

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

When changing a thin volume's allocated capacity limit and warning threshold, you must take into account the space to be consumed by both the volume's user data and copy services data.

Steps

1. Select **Storage > Volumes**.

2. Select the **Thin Volume Monitoring** tab.

The **Thin Volume Monitoring** view appears.

3. Select the thin volume that you want to change, and then select **Change Limit**.

The **Change Limit** dialog box appears. The allocated capacity limit and warning threshold setting for the thin volume you selected appear in this dialog box.

4. Change the allocated capacity limit and warning threshold as needed.

Field Details

Setting	Description
Change allocated capacity limit to...	The threshold at which writes fail, preventing the thin volume from consuming additional resources. This threshold is a percentage of the volume's reported capacity size.
Alert me when... (warning threshold)	<p>Select the check box if you want the system to generate an alert when a thin volume is near the allocated capacity limit. The alert is sent to the Notifications area on the Home page. This threshold is a percentage of the volume's reported capacity size.</p> <p>Clear the check box to disable the warning threshold alert notification.</p>

5. Click **Save**.

FAQs

What is a volume?

A volume is a container in which applications, databases, and file systems store data. It is the logical component created for the host to access storage on the storage array.

A volume is created from the capacity available in a pool or a volume group. A volume has a defined capacity.

Although a volume might consist of more than one drive, a volume appears as one logical component to the host.

Why am I seeing a capacity over-allocation error when I have enough free capacity in a volume group to create volumes?

The selected volume group might have one or more free capacity areas. A free capacity area is the free capacity that can result from deleting a volume or from not using all available free capacity during volume creation.

When you create a volume in a volume group that has one or more free capacity areas, the volume's capacity is limited to the largest free capacity area in that volume group. For example, if a volume group has a total of 15 GiB free capacity, and the largest free capacity area is 10 GiB, the largest volume you can create is 10 GiB.

If a volume group has free capacity areas, the volume group graph contains a link indicating the number of existing free capacity areas. Select the link to display a pop-over that indicates the capacity of each area.

By consolidating free capacity, you can create additional volumes from the maximum amount of free capacity in a volume group. You can consolidate the existing free capacity on a selected volume group using one of the following methods:

- When at least one free capacity area is detected for a volume group, the **Consolidate free capacity** recommendation appears on the **Home** page in the Notification area. Click the **Consolidate free capacity** link to launch the dialog box.
- You can also select **Pools & Volume Groups > Uncommon Tasks > Consolidate volume group free capacity** to launch the dialog box.

If you want to use a specific free capacity area rather than the largest free capacity area, use the Command Line Interface (CLI).

How does my selected workload impact volume creation?

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

- **Application-specific.** When you are creating volumes using an application-specific workload, the system may recommend an optimized volume configuration to minimize contention between application workload I/O and other traffic from your application instance. Volume characteristics like I/O type, segment size, controller ownership, and read and write cache are automatically recommended and optimized for workloads that are created for the following application types.
 - Microsoft® SQL Server™
 - Microsoft® Exchange Server™
 - Video surveillance applications
 - VMware ESXi™ (for volumes to be used with Virtual Machine File System) You can review the

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

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

Why aren't these volumes associated with a workload?

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

Why can't I delete the selected workload?

This workload consists of a group of volumes that were created using the command line interface (CLI) or migrated (imported/exported) from a different storage array. As a result, the volumes in this workload are not associated with an application-specific workload, so the workload cannot be deleted.

How do application-specific workloads help me manage my storage array?

An application is software such as SQL Server or Exchange. You define one or more workloads to support each application. For some applications, System Manager will automatically recommend a volume configuration that optimizes storage. Characteristics such as I/O type, segment size, controller ownership, and read and write cache are included in the volume configuration.

The volume characteristics of your application-specific workload dictate how the workload interacts with the components of your storage array and helps determine the performance of your environment under a given configuration.

How does providing this information help create storage?

The workload information is used to optimize the volume characteristics such as I/O type, segment size, and read/write cache for the workload selected. These optimized characteristics dictate how your workload interacts with the storage array components.

Based on the workload information you provide, System Manager creates the appropriate volumes and places them on the available pools or volume groups that currently exist on the system. The system creates the volumes and optimizes their characteristics based on the current best practices for the workload you selected.

Before you finish creating volumes for a given workload, you can review the recommended volume configuration and edit, add, or delete the system-recommended volumes and characteristics using the **Add/Edit Volumes** dialog box.

Refer to your application-specific documentation for best practice information.

What do I need to do to recognize the expanded capacity?

If you increase the capacity for a volume, the host might not immediately recognize the increase in volume capacity.

Most operating systems recognize the expanded volume capacity and automatically expand after the volume expansion is initiated. However, some might not. If your OS does not automatically recognize the expanded volume capacity, you might need to perform a disk rescan or reboot.

After you have expanded the volume capacity, you must manually increase the file system size to match. How you do this depends on the file system you are using.

Refer to your host operating system documentation for additional details.

Why don't I see all my pools and/or volume groups?

Any pool or volume group to which you cannot move the volume does not display in the list.

Pools or volume groups are not eligible for any of the following reasons:

- The Data Assurance (DA) capabilities of a pool or volume group pool do not match.
- A pool or volume group is in a non-optimal state.
- The capacity of a pool or volume group is too small.

What is segment size?

A segment is the amount of data in kilobytes (KiB) that is stored on a drive before the storage array moves to the next drive in the stripe (RAID group). Segment size applies only to volume groups, not pools.

Segment size is defined by the number of data blocks it contains. For example:

- 64 KiB segment = 128 data blocks
- 512 KiB segment = 1024 data blocks

When determining segment size, you must know what type of data you will store in a volume. If an application typically uses small, random reads and writes (IOPS), a smaller segment size typically works better. Alternatively, if the application has large, sequential reads and writes (throughput), a large segment size is generally better.

Whether an application uses small random reads and writes, or large sequential reads and writes, the storage array performs better if the segment size is larger than the typical data block chunk size. This normally makes it easier and faster for the drives to access the data, which is important for better storage array performance.

In an environment where IOPS performance is important

In an I/O operations per second (IOPS) environment, the storage array performs better if you use a segment size that is larger than the typical data block size ("chunk") that is read/written to a drive. This ensures that each chunk is written to a single drive.

In an environment where throughput is important

In a throughput environment, the segment size should be an even fraction of the total drives for data and the typical data chunk size (I/O size). This spreads the data as a single stripe across the drives in the volume group leading to faster reads and writes.

For example, in a 5-drive RAID 5 volume group (4+1), if the typical read/write “chunk” size is 2 MiB, a segment size of 512 KiB (an even fraction [1/4] of the total chunk size) would be the best choice for the application’s volume segment size because it ensures that each read/write is written as a single stripe of the volume group drives.

What is preferred controller ownership?

Preferred controller ownership defines the controller that is designated to be the owning, or primary, controller of the volume.

Controller ownership is very important and should be planned carefully. Controllers should be balanced as closely as possible for total I/Os.

For example, if one controller reads primarily large, sequential data blocks and the other controller has small data blocks with frequent reads and writes, the loads are very different. Knowing which volumes contain what type of data allows you to balance I/O transfers equally over both controllers.

What is Automatic Load Balancing?

The Automatic Load Balancing feature provides automated I/O balancing and ensures that incoming I/O traffic from the hosts is dynamically managed and balanced across both controllers.

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

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

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

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



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

Hosts

Concepts

Host terminology


Learn how the host terms apply to your storage array.

Component	Definition
Host	A host is a server that sends I/O to a volume on a storage array.
Host name	The host name should equate to the system name of the host.
Host cluster	A host cluster is a group of hosts. You create a host cluster to make it easy to assign the same volumes to multiple hosts.
Host interface protocol	A host interface protocol is the connection (such as Fibre Channel, iSCSI, etc.) between the controllers and the hosts.
HBA or Network Interface Card (NIC)	A host bus adapter (HBA) is a board that resides in a host and contains one or more host ports.
Host port	A host port is a port on a host bus adapter (HBA) that provides the physical connection to a controller and is used for I/O operations.
Host port identifier	<p>A host port identifier is a unique world-wide name associated with each host port on a host bus adapter (HBA).</p> <ul style="list-style-type: none">• Internet Small Computer System Interface (iSCSI) host port identifiers must have between 1 and 233 characters. iSCSI host port identifiers display in standard IQN format (e.g., <code>iqn.xxx.com.xxx:8b3ad</code>).• Non-iSCSI host port identifiers such as Fibre Channel and Serial Attached SCSI (SAS) display as colon-delimited after every two characters (e.g., <code>xx:yy:zz</code>). Fibre Channel host port identifiers must have 16 characters.
Host operating system type	The host operating system type is a configuration setting that defines how the controllers in the storage array react to I/O depending on the operating system (or variant) of the host. This is also sometimes called <i>host type</i> for short.

Component	Definition
Controller host port	A controller host port is a port on the controller that provides the physical connection to a host and is used for I/O operations.
LUN	<p>A logical unit number (LUN) is the number assigned to the address space that a host uses to access a volume. The volume is presented to the host as capacity in the form of a LUN.</p> <p>Each host has its own LUN address space. Therefore, the same LUN can be used by different hosts to access different volumes.</p>

Workflow for host creation and volume assignment

The following figure illustrates how to configure host access.



Automatic versus manual host creation

Creating a host is one of the steps required to let the storage array know which hosts are attached to it and to allow I/O access to the volumes. You can create a host automatically or manually.

Automatic creation

Automatic host creation is initiated by the Host Context Agent (HCA). The HCA is a utility that you must install on each host attached to the storage array. Each host that has the HCA installed pushes its configuration information to the storage array controllers through the I/O path. Based on the host information, the controllers automatically create the host and the associated host ports and set the host type. If needed, you can make any additional changes to the host configuration using System Manager.

After the HCA performs its automatic detection, the host automatically appears in the Hosts page with the

following attributes:

- The host name derived from the system name of the host.
- The host identifier ports that are associated with the host.
- The Host Operating System Type of the host.

Hosts are created as stand-alone hosts; the HCA does not automatically create or add to host clusters.

Manual creation

You might want to manually create a host for one of the following reasons:

1. You chose not to install the HCA utility on your hosts.
2. You want to ensure that the host port identifiers that were detected by the storage array controllers are associated correctly with the hosts.

During manual host creation, you associate host port identifiers by selecting them from a list or manually entering them. After you create a host, you can assign volumes to it or add it to a host cluster if you plan to share access to volumes.

How volumes are assigned to hosts and host clusters

For a host or host cluster to send I/O to a volume, you must assign the volume to the host or host cluster.

You can select a host or host cluster when you create a volume or you can assign a volume to a host or host cluster later. A host cluster is a group of hosts. You create a host cluster to make it easy to assign the same volumes to multiple hosts.

Assigning volumes to hosts is flexible, allowing you to meet your particular storage needs.

- **Stand-alone host, not part of a host cluster** — You can assign a volume to an individual host. The volume can be accessed only by the one host.
- **Host cluster** — You can assign a volume to a host cluster. The volume can be accessed by all the hosts in the host cluster.
- **Host within a host cluster** — You can assign a volume to an individual host that is part of a host cluster. Even though the host is part of a host cluster, the volume can be accessed only by the individual host and not by any other hosts in the host cluster.

When volumes are created, logical unit numbers (LUNs) are assigned automatically. The LUN serves as the "address" between the host and the controller during I/O operations. You can change LUNs after the volume is created.

Access volumes

An access volume is a factory-configured volume on the storage array that is used for communication with the storage array and the host through the host I/O connection. The access volume requires a Logical Unit Number (LUN).

The access volume is used in two instances:

- **Automatic host creation** — The access volume is used by the Host Context Agent (HCA) utility to push

host information (name, ports, host type) to System Manager for automatic host creation.

- **In-band management** — The access volume is used for an in-band connection to manage the storage array. This can only be done if you are managing the storage array with the command line interface (CLI).

An access volume is automatically created the first time you assign a volume to a host. For example, if you assign Volume_1 and Volume_2 to a host, when you view results of that assignment, you see three volumes (Volume_1, Volume_2, and Access).

If you are not automatically creating hosts or managing a storage array in-band with the CLI, you do not need the access volume, and you can free up the LUN by deleting the access volume. This action removes the volume-to-LUN assignment as well as any in-band management connections to the host.

Maximum number of LUNs

The storage array has a maximum number of logical unit numbers (LUNs) that can be used for each host.

The maximum number depends on the operating system of the host. The storage array tracks the number of LUNs used. If you try to assign a volume to a host that exceeds the maximum number of LUNs, the host cannot access the volume.

How to

Configure host access

Create host automatically

You can allow the Host Context Agent (HCA) to automatically detect the hosts, and then verify that the information is correct. Creating a host is one of the steps required to let the storage array know which hosts are attached to it and to allow I/O access to the volumes.

Before you begin

The Host Context Agent (HCA) is installed and running on every host connected to the storage array. Hosts with the HCA installed and connected to the storage array are created automatically. To install the HCA, install SANtricity Storage Manager on the host and select the Host option. The HCA is not available on all supported operating systems. If it is not available, you must create the host manually.

Steps

1. Select **Storage > Hosts**.

The table lists the automatically-created hosts.

2. Verify that the information provided by the HCA is correct (name, host type, host port identifiers).

If you need to change any of the information, select the host, and then click **View/Edit Settings**.

3. (Optional) If you want the automatically-created host to be in a cluster, create a host cluster and add the host or hosts.

Results

After a host is created automatically, the system displays the following items in the Hosts tile table:

- The host name derived from the system name of the host.
- The host identifier ports that are associated with the host.
- The Host Operating System Type of the host.

Create host manually

For hosts that cannot be automatically discovered, you can manually create a host. Creating a host is one of the steps required to let the storage array know which hosts are attached to it and to allow I/O access to the volumes.

About this task

Keep these guidelines in mind when you create a host:

- You must define the host identifier ports that are associated with the host.
- Make sure that you provide the same name as the host's assigned system name.
- This operation does not succeed if the name you choose is already in use.
- The length of the name cannot exceed 30 characters.

Steps

1. Select **Storage › Hosts**.
2. Click **Create › Host**.

The Create Host dialog box appears.

3. Select the settings for the host as appropriate.

Field details

Setting	Description
Name	Type a name for the new host.
Host operating system type	Select the operating system that is running on the new host from the drop-down list.
Host interface type	(Optional) If you have more than one type of host interface supported on your storage array, select the host interface type that you want to use.
Host ports	<p>Do one of the following:</p> <ul style="list-style-type: none">• Select I/O Interface <p>Generally, the host ports should have logged in and be available from the drop-down list. You can select the host port identifiers from the list.</p> <ul style="list-style-type: none">• Manual add <p>If a host port identifier is not displayed in the list, it means that the host port has not logged in. An HBA utility or the iSCSI initiator utility may be used to find the host port identifiers and associate them with the host.</p> <p>You can manually enter the host port identifiers or copy/paste them from the utility (one at a time) into the Host ports field.</p> <p>You must select one host port identifier at a time to associate it with the host, but you can continue to select as many identifiers that are associated with the host. Each identifier is displayed in the Host ports field. If necessary, you also can remove an identifier by selecting the X next to it.</p>

Setting	Description
CHAP initiator	<p>(Optional) If you selected or manually entered a host port with an iSCSI IQN, and if you want to require a host that tries to access the storage array to authenticate using Challenge Handshake Authentication Protocol (CHAP), select the CHAP initiator checkbox. For each iSCSI host port you selected or manually entered, do the following:</p> <ul style="list-style-type: none"> • Enter the same CHAP secret that was set on each iSCSI host initiator for CHAP authentication. If you are using mutual CHAP authentication (two-way authentication that enables a host to validate itself to the storage array and for a storage array to validate itself to the host), you also must set the CHAP secret for the storage array at initial setup or by changing settings. • Leave the field blank if you do not require host authentication. Currently, the only iSCSI authentication method used by System Manager is CHAP.

4. Click **Create**.

Results

After the host is successfully created, the system creates a default name for each host port configured for the host (user label).

The default alias is <Hostname_Port Number>. For example, the default alias for the first port created for host IPT is IPT_1.

Create host cluster

You create a host cluster when two or more hosts require I/O access to the same volumes.

About this task

Keep these guidelines in mind when you create a host cluster:

- This operation does not start unless there are two or more hosts available to create the cluster.
- Hosts in host clusters can have different operating systems (heterogeneous).
- To create a Data Assurance (DA)-enabled volume, the host connection you are planning to use must support DA.

If any of the host connections on the controllers in your storage array do not support DA, the associated hosts cannot access data on DA-enabled volumes. DA is **not** supported by iSCSI over TCP/IP, or by the SRP over InfiniBand.

- This operation does not succeed if the name you choose is already in use.
- The length of the name cannot exceed 30 characters.

Steps

1. Select **Storage › Hosts**.
2. Select **Create › Host Cluster**.

The Create Host Cluster dialog box appears.

3. Select the settings for the host cluster as appropriate.

Field details

Setting	Description
Name	Type the name for the new host cluster.
Hosts	Select two or more hosts from the drop-down list. Only those hosts that are not already part of a host cluster appear in the list.

4. Click **Create**.

If the selected hosts are attached to interface types that have different Data Assurance (DA) capabilities, a dialog appears with the message that DA will be unavailable on the host cluster. This unavailability prevents DA-enabled volumes from being added to the host cluster. Select **Yes** to continue or **No** to cancel.

DA increases data integrity across the entire storage system. DA enables the storage array to check for errors that might occur when data is moved between the hosts and the drives. Using DA for the new volume ensures that any errors are detected.

Results

The new host cluster appears in the table with the assigned hosts in the rows beneath.

Assign volumes

You must assign a volume to a host or a host cluster so it can be used for I/O operations. This assignment grants a host or host cluster access to one or more volumes in a storage array.

Before you begin

Keep these guidelines in mind when you assign volumes:

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

Assigning a volume fails under these conditions:

- All volumes are assigned.
- The volume is already assigned to another host or host cluster.

The ability to assign a volume is unavailable under these conditions:

- No valid hosts or host clusters exist.
- No host port identifiers have been defined for the host.
- All volume assignments have been defined.

About this task

All unassigned volumes are displayed, but functions for hosts with or without Data Assurance (DA) apply as follows:

- For a DA-capable host, you can select volumes that are either DA-enabled or not DA-enabled.
- For a host that is not DA-capable, if you select a volume that is DA-enabled, a warning states that the system must automatically turn off DA on the volume before assigning the volume to the host.

Steps

1. Select **Storage › Hosts**.
2. Select the host or host cluster to which you want to assign volumes, and then click **Assign Volumes**.

A dialog box appears that lists all the volumes that can be assigned. You can sort any of the columns or type something in the **Filter** box to make it easier to find particular volumes.

3. Select the check box next to each volume that you want to assign or select the check box in the table header to select all volumes.
4. Click **Assign** to complete the operation.

Results

After successfully assigning a volume or volumes to a host or a host cluster, the system performs the following actions:

- The assigned volume receives the next available LUN number. The host uses the LUN number to access the volume.
- The user-supplied volume name appears in volume listings associated to the host. If applicable, the factory-configured access volume also appears in volume listings associated to the host.

Manage hosts and host clusters

Change the settings for a host

You can change the name, host operating system type, and associated host clusters for a host.

Steps

1. Select **Storage › Hosts**.
2. Select the host that you want to edit, and then click **View/Edit Settings**.

A dialog box appears that shows the current host settings.

3. If it is not already selected, click the **Properties** tab.
4. Change the settings as appropriate.

Field Details

Setting	Description
Name	You can change the user-supplied name of the host. Specifying a name for the host is required.
Associated host cluster	<p>You can choose one of the following options:</p> <ul style="list-style-type: none">• None — The host remains a standalone host. If the host was associated to a host cluster, the system removes the host from the cluster.• <Host Cluster> — The system associates the host to the selected cluster.
Host operating system type	You can change the type of operating system running on the host you defined.

5. Click **Save**.

Change the settings for a host cluster

You can change the host cluster name, or add or remove hosts in a host cluster.

Steps

1. Select **Storage > Hosts**.
2. Select the host cluster you want to edit, and then click **View/Edit Settings**.

A dialog box appears that shows the current host cluster settings.

3. Change the settings for the host cluster as appropriate.

Field Details

Setting	Description
Name	You can specify the user-supplied name of the host cluster. Specifying a name for a cluster is required.
Associated Hosts	<p>To add a host, click the Associated Hosts box, and then select a host name from the drop-down list. You cannot manually enter a host name.</p> <p>To delete a host, click the X next to the host name.</p>

4. Click **Save**.

Unassign volumes

Unassign volumes from hosts or host clusters if you no longer need I/O access to that volume from the host or host cluster.

About this task

Keep these guidelines in mind when you unassign a volume:

- If you are removing the last assigned volume from a host cluster, and the host cluster also has hosts with specific assigned volumes, make sure that you remove or move those assignments before removing the last assignment for the host cluster.
- If a host cluster, a host, or a host port is assigned to a volume that is registered to the operating system, you must clear this registration before you can remove these nodes.

Steps

1. Select **Storage > Hosts**.
2. Select the host or host cluster that you want to edit, and then click **Unassign Volumes**.

A dialog box appears that shows all the volumes that are currently assigned.

3. Select the check box next to each volume that you want to unassign or select the check box in the table header to select all volumes.
4. Click **Unassign**.

Results

- The volumes that were unassigned are available for a new assignment.
- Until the changes are configured on the host, the volume is still recognized by the host operating system.

Change host port identifiers for a host

Change the host port identifiers when you want to change the user label on a host port identifier, add a new host port identifier to the host, or delete a host port identifier from the

host.

About this task

When changing host port identifiers, keep the following guidelines in mind:

- **Add** — When you add a host port, you are associating the host port identifier to the host you created to connect to your storage array. You can manually enter port information using a host bus adapter (HBA) utility.
- **Edit** — You can edit the host ports to move (associate) a host port to a different host. You might have moved the host bus adapter or iSCSI initiator to a different host, so you must move (associate) the host port to the new host.
- **Delete** — You can delete host ports to remove (unassociate) host ports from a host.

Steps

1. Select **Storage > Hosts**.
2. Select the host to which the ports will be associated, and then click **View/Edit Settings**.


If you want to add ports to a host in a host cluster, expand the host cluster and select the desired host. You cannot add ports at the host cluster level.

A dialog box appears that shows the current host settings.

3. Click the **Host Ports** tab.

The dialog box shows the current host port identifiers.

4. Change the host port identifier settings as appropriate.

Setting	Description
Host Port	<p>You can choose one of the following options:</p> <ul style="list-style-type: none"> • Add — Use Add to associate a new host port identifier to the host. The length of the host port identifier name is determined by the host interface technology. Fibre Channel and Infiniband host port identifier names must have 16 characters. iSCSI host port identifier names have a maximum of 223 characters. The port must be unique. A port number that has already been configured is not allowed. • Delete — Use Delete to remove (unassociate) a host port identifier. The Delete option does not physically remove the host port. This option removes the association between the host port and the host. Unless you remove the host bus adapter or the iSCSI initiator, the host port is still recognized by the controller. <div>  <p>If you delete a host port identifier, it is no longer associated with this host. Also, the host loses access to any of its assigned volumes through this host port identifier.</p> </div>
Label	<p>To change the port label name, click the Edit icon (pencil). The port label name must be unique. A label name that has already been configured is not allowed.</p>
CHAP Secret	<p>Appears only for iSCSI hosts. You can set or change the CHAP secret for the initiators (iSCSI hosts).</p> <p>System Manager uses the Challenge Handshake Authentication Protocol (CHAP) method, which validates the identity of targets and initiators during the initial link. Authentication is based on a shared security key called a CHAP secret.</p>

5. Click **Save**.

Delete host or host cluster

You can delete a host or host cluster.

About this task

Keep these guidelines in mind when you delete a host or a host cluster:

- Any specific volume assignments are deleted, and the associated volumes are available for a new assignment.
- If the host is part of a host cluster that has its own specific assignments, the host cluster is unaffected. However, if the host is part of a host cluster that does not have any other assignments, the host cluster and any other associated hosts or host port identifiers inherit any default assignments.
- Any host port identifiers that were associated with the host become undefined.

Steps

1. Select **Storage > Hosts**.
2. Select the host or host cluster that you want to delete, and then click **Delete**.

The confirmation dialog box appears.

3. Confirm that you want to perform the operation, and then click **Delete**.

Results

If you deleted a host, the system performs the following actions:

- Deletes the host and, if applicable, removes it from the host cluster.
- Removes access to any assigned volumes.
- Returns the associated volumes to an unassigned state.
- Returns any host port identifiers associated with the host to an unassociated state.

If you deleted a host cluster, the system performs the following actions:

- Deletes the host cluster and its associated hosts (if any).
- Removes access to any assigned volumes.
- Returns the associated volumes to an unassigned state.
- Returns any host port identifiers associated with the hosts to an unassociated state.

FAQs

What are hosts and host clusters?

A host is a server that sends I/O to a volume on a storage array. A host cluster is a group of hosts. You create a host cluster to make it easy to assign the same volumes to multiple hosts.

You define a host separately. It can either be an independent entity or be added to a host cluster. You can assign volumes to an individual host, or a host can be part of a host cluster that shares access to one or more volumes with other hosts in the host cluster.

The host cluster is a logical entity that you create in SANtricity System Manager. You must add hosts to the host cluster before you can assign volumes.

Why would I need to create a host cluster?

You need to create a host cluster if you want to have two or more hosts share access to the same set of volumes. Normally, the individual hosts have clustering software installed

on them to coordinate volume access.

How do I know which host operating system type is correct?

The Host Operating System Type field contains the operating system of the host. You can select the recommended host type from the drop-down list or allow the Host Context Agent (HCA) to configure the host and appropriate host operating system type.

Host Operating System type	Operating System (OS) and multipath driver
AIX MPIO	The Advanced Interactive Executive (AIX) OS and the native MPIO driver
AVT_4M	Silicon Graphics, Inc. (SGI) proprietary multipath driver; refer to the SGI installation documentation for more information
Factory Default	This is reserved for the initial startup of the storage array and should be changed to match the host operating system and multipath driver being used for the particular host
HP-UX	The HP-UX OS with native multipath driver
Linux (ATTO)	The Linux OS and the ATTO Technology, Inc. driver (must use ATTO FC HBAs)
Linux (DM-MP)	The Linux OS and the native DM-MP driver
Linux (Pathmanager)	The Linux OS and the SGI proprietary multipath driver; refer to the SGI installation documentation for more information
Mac OS	The Mac OS and the ATTO Technology, Inc. driver
ONTAP	FlexArray
Solaris (version 11 or later)	The Solaris 11 or later OS and the native MPxIO driver
Solaris (version 10 or earlier)	The Solaris 10 or earlier OS and the native MPxIO driver
SVC	IBM SAN Volume Controller
VMware	The ESXi OS

Host Operating System type	Operating System (OS) and multipath driver
Windows or Windows Clustered	The Windows Server OS and Windows MPIO with a DSM driver
Windows (ATTO)	The Windows OS and the ATTO Technology, Inc. driver

After the HCA is installed and the storage is attached to the host, the HCA sends the host topology to the storage controllers through the I/O path. Based on the host topology, the storage controllers automatically define the host and the associated host ports, and then set the host type.



If the HCA does not select the recommended host type, you must manually set the host type in System Manager.

What are HBAs and adapter ports?

A host bus adapter (HBA) is a board that resides in a host and contains one or more host ports. A host port is a port on a host bus adapter (HBA) that provides the physical connection to a controller and is used for I/O operations.

The adapter ports on the HBA are called host ports. Most HBAs have either one or two host ports. The HBA has a unique World Wide Identifier (WWID), and each HBA host port has a unique WWID. The host port identifiers are used to associate the appropriate HBA with the physical host when you are either manually creating the host through SANtricity System Manager or automatically creating the host using the host context agent.

How do I match the host ports to a host?

If you are manually creating a host, you first must use the appropriate host bus adapter (HBA) utility available on the host to determine the host port identifiers associated with each HBA installed in the host.

When you have this information, select the host port identifiers that have logged into the storage array from the list provided in the Create Host dialog of System Manager.



Make sure you select the appropriate host port identifiers for the host you are creating. If you associate the wrong host port identifiers, you might cause unintended access from another host to this data.

If you are automatically creating hosts using the host context agent (HCA) installed on each host, the HCA should automatically associate the host port identifiers with each host and configure them appropriately.

How do I create CHAP secrets?

If you set up Challenge Handshake Authentication Protocol (CHAP) authentication on any iSCSI host connected to the storage array, you must re-enter that initiator CHAP secret for each iSCSI host. To do this, you can use System Manager either as part of the Create Host operation or through the View/Edit Settings option.

If you are using CHAP mutual authentication, you also must define a target CHAP secret for the storage array in the Settings page and re-enter that target CHAP secret on each iSCSI host.

What is the default cluster?

The default cluster is a system-defined entity that allows any unassociated host bus adapter (HBA) host port identifier that has logged into the storage array to gain access to any volumes assigned to the default cluster. An unassociated host port identifier is a host port that while physically installed in a host and logged into the storage array is not logically associated with a particular host.



If you want your hosts to have specific access to certain volumes in the storage array, you must *not* use the default cluster. Instead, you must associate the host port identifiers with their corresponding hosts. This can be done either manually using System Manager during the Create Host operation or automatically using the host context agent (HCA) installed on each host. Then, you assign volumes either to an individual host or to a host cluster.

You should *only* use the default cluster in special situations where your external storage environment is conducive to allowing all the hosts and all the logged-in host port identifiers connected to the storage array have access to all of the volumes (all-access mode) without specifically making the hosts known to the storage array or System Manager.

Initially, you can assign volumes only to the default cluster through the command line interface (CLI). However, after you assign at least one volume to the default cluster, this entity (called Default Cluster) is displayed in System Manager, and you can then use System Manager to manage this entity.

Performance

Concepts

Performance overview

The Performance page provides graphs and tables of data that enable you to assess the storage array's performance in several key areas.

Performance functions allow you to accomplish these tasks:

- View performance data in near real-time to help you determine whether a storage array is experiencing problems.
- Export performance data to construct a historical view of a storage array and identify when a problem started or what caused a problem.
- Select the objects, performance metrics, and time frame you want to view.
- Compare metrics.

You can view performance data in three formats:

- **Real-time graphical** — Plots performance data on a graph in near real-time.
- **Near real-time tabular** — Lists performance data in a table in near real-time.
- **Exported CSV file** — Allows you to save tabular performance data in a file of comma-separated values for

further viewing and analysis.

Characteristics of performance data formats

Type of performance monitoring	Sampling interval	Length of time displayed	Maximum number of objects displayed	Ability to save data
Real-time graphical, live Real-time graphical, historical	10 sec (live) 5 min (historical) Data points shown depend on selected time frame	Default time frame is 1 hour. Choices: <ul style="list-style-type: none">• 5 minutes• 1 hour• 8 hours• 1 day• 7 days• 30 days	5	No
Near real-time tabular (table view)	10 sec -1 hr	Most current value	Unlimited	Yes
Comma-separated values (CSV) file	Depends on selected time frame	Depends on selected time frame	Unlimited	Yes

Guidelines for viewing performance data

- Performance data collection is always on. There is no option to turn it off.
- Each time the sampling interval elapses, the storage array is queried and the data is updated.
- For graphical data, the 5-minute time frame supports 10-second updating averaged over 5 minutes. All other time frames are updated every 5 minutes, averaged over the selected time frame.
- Performance data in the graphical views is updated in real time. Performance data in the table view is updated in near real time.
- If a monitored object changes during the time data is collected, the object might not have a complete set of data points spanning the selected time frame. For example, volume sets can change as volumes are created, deleted, assigned, or unassigned; or drives can be added, removed, or failed.

Performance terminology

Learn how the performance terms apply to your storage array.

Term	Description
Application	An application is a software program, such as SQL or Exchange.

Term	Description
CPU	CPU is short for "central processing unit." CPU indicates the percentage of the storage array's processing capacity being used.
Host	A host is a server that sends I/O to a volume on a storage array.
IOPS	IOPS stands for input/output operations per second.
Latency	Latency is the time interval between a request, such as for a read or write command, and the response from the host or the storage array.
LUN	<p>A logical unit number (LUN) is the number assigned to the address space that a host uses to access a volume. The volume is presented to the host as capacity in the form of a LUN.</p> <p>Each host has its own LUN address space. Therefore, the same LUN can be used by different hosts to access different volumes.</p>
MiB	MiB is an abbreviation for mebibyte (mega binary byte). One MiB is 220, or 1,048,576 bytes. Compare with MB, which signifies a base 10 value. One MB equals 1,024 bytes.
Object	<p>An object is any logical or physical storage component.</p> <p>Logical objects include volume groups, pools, and volumes. Physical objects include the storage array, array controllers, hosts, and drives.</p>
Pool	A pool is a set of drives that is logically grouped. You can use a pool to create one or more volumes accessible to a host. (You create volumes from either a pool or a volume group.)
Read	Read is short for "read operation," which occurs when the host requests data from the storage array.

Term	Description
Volume	<p>A volume is a container in which applications, databases, and file systems store data. It is the logical component created for the host to access storage on the storage array.</p> <p>A volume is created from the capacity available in a pool or a volume group. A volume has a defined capacity. Although a volume might consist of more than one drive, a volume appears as one logical component to the host.</p>
Volume name	A volume name is a string of characters assigned to the volume when it is created. You can either accept the default name or provide a more descriptive name indicating the type of data stored in the volume.
Volume group	A volume group is a container for volumes with shared characteristics. A volume group has a defined capacity and RAID level. You can use a volume group to create one or more volumes accessible to a host. (You create volumes from either a volume group or a pool.)
Workload	A workload is a storage object that supports an application. You can define one or more workloads, or instances, per application. For some applications, System Manager configures the workload to contain volumes with similar underlying volume characteristics. These volume characteristics are optimized based on the type of application the workload supports. For example, if you create a workload that supports a Microsoft SQL Server application and then subsequently create volumes for that workload, the underlying volume characteristics are optimized to support Microsoft SQL Server.
Write	Write is short for "write operation," when data is sent from the host to the array for storage.

How tos

View graphical performance data

You can view graphical performance data for logical objects, physical objects, applications, and workloads.

About this task

The performance graphs show historical data as well as live data currently being captured. A vertical line on the graph, labeled **Live updating**, distinguishes historical data from live data.

Home page view

The **Home** page contains a graph showing storage array level performance. You can select limited metrics from this view, or you can click **View Performance Details** to select all the available metrics.

Detailed view

The graphs available from the detailed performance view are arranged under three tabs:

- **Logical View** — Displays performance data for logical objects grouped by volume groups and pools. Logical objects include volume groups, pools, and volumes.
- **Physical View** — Displays performance data for the controller, host channels, drive channels, and drives.
- **Applications & Workloads View** — Displays a list of logical objects (volumes) grouped by the application types and workloads you have defined.

Steps

1. Select **Home**.
2. To select an array-level view, click the IOPS, MiB/s, or CPU button.
3. To see more details, click **View Performance Details**.
4. Select **Logical View** tab, **Physical View** tab, or **Applications & Workloads View** tab.

Depending on the object type, different graphs appear in each tab.

View tabs	Performance data displayed for each object type
Logical View	<ul style="list-style-type: none">• Storage array: IOPS, MiB/s• Pools: Latency, IOPS, MiB/s• Volume groups: Latency, IOPS, MiB/s• Volumes: Latency, IOPS, MiB/s
Physical View	<ul style="list-style-type: none">• Controllers: IOPS, MiB/s, CPU, Headroom• Host channels: Latency, IOPS, MiB/s, Headroom• Drive channels: Latency, IOPS, MiB/s• Drives: Latency, IOPS, MiB/s
Applications & Workloads View	<ul style="list-style-type: none">• Storage array: IOPS, MiB/s• Applications: Latency, IOPS, MiB/s• Workloads: Latency, IOPS, MiB/s• Volumes: Latency, IOPS, MiB/s


5. Use the options to view the objects and information you need.

Options

Options for viewing objects	Description
Expand a drawer to see the list of objects.	<p><i>Navigation drawers</i> contain storage objects, such as pools, volume groups, and drives.</p> <p>Click the drawer to view the list of objects in the drawer.</p>
Select objects to view.	Select the check box to the left of each object to choose the performance data you want to view.
Use Filter to find object names or partial names.	In the Filter box, enter the name or a partial name of objects to list just those objects in the drawer.
Click Refresh Graphs after selecting objects.	After selecting objects from the drawers, select Refresh Graphs to view graphical data for the items you have selected.
Hide or show graph	Select the graph title to hide or show the graph.

6. As needed, use the additional options for viewing performance data.

Additional options

Option	Description
Time frame	<p>Select the length of time you want to view (5 minutes, 1 hour, 8 hours, 1 day, 7 days, or 30 days). The default is 1 hour.</p> <div><p>Loading performance data for a 30-day time frame can take several minutes. Do not navigate away from the web page, refresh the web page, or close the browser while data is loading.</p></div>
Data point details	<p>Hover the cursor over the graph to see metrics for a particular data point.</p>
Scroll bar	<p>Use the scroll bar below the graph to view an earlier or later time span.</p>
Zoom bar	<p>Below the graph, drag the zoom bar handles to zoom out on a time span. The wider the zoom bar, the less granular the details of the graph.</p> <p>To reset the graph, select one of the time frame options.</p>
Drag and drop	<p>On the graph, drag the cursor from one point in time to another to zoom in on a time span.</p> <p>To reset the graph, select one of the time frame options.</p>

View and save tabular performance data

You can view and save performance graphs data in tabular format. This allows you to filter the data you want displayed.

Steps

1. From any performance data graph, click **Launch table view**.

A table appears that lists all the performance data for the selected objects.

2. Use the object selection pull-down and the filter as needed.
3. Click the Show/Hide Columns button to select the columns you want to include in the table.

You can click each check box to select or deselect an item.

4. Select **Export** at the bottom of the screen to save the tabular view to a file of comma-separated values (CSV).

The **Export Table** dialog box appears, indicating the number of rows to be exported and the file format of the export (comma-separated values, or CSV format).

5. Click **Export** to proceed with the download, or click **Cancel**.

Depending on your browser settings, the file is either saved, or you are prompted to choose a name and location for the file.

The default file name format is `performanceStatistics-yyyy-mm-dd_hh-mm-ss.csv`, which includes the date and time when the file was exported.

Interpret performance data

Performance data can guide you in tuning the performance of your storage array.

When interpreting Performance data, keep in mind that several factors affect the performance of your storage array. The following table describes the main areas to consider.

Performance data	Implications for performance tuning
Latency (milliseconds, or ms)	<p data-bbox="821 155 1341 186">Monitor the I/O activity of a specific object.</p> <p data-bbox="821 222 1398 254">Potentially identify objects that are bottlenecks:</p> <ul data-bbox="846 289 1479 1199" style="list-style-type: none"> <li data-bbox="846 289 1479 422">• If a volume group is shared among several volumes, the individual volumes might need their own volume groups to improve the sequential performance of the drives and decrease latency. <li data-bbox="846 443 1479 575">• With pools, larger latencies are introduced and uneven workloads might exist between drives, making the latency values less meaningful and, in general, higher. <li data-bbox="846 596 1479 695">• Drive type and speed influence latency. With random I/O, faster spinning drives spend less time moving to and from different locations on the disk. <li data-bbox="846 716 1479 848">• Too few drives result in more queued commands and a greater period of time for the drive to process the command, increasing the general latency of the system. <li data-bbox="846 869 1479 932">• Larger I/Os have greater latency due to the additional time involved with transferring data. <li data-bbox="846 953 1479 1085">• Higher latency might indicate that the I/O pattern is random in nature. Drives with random I/O will have greater latency than those with sequential streams. <li data-bbox="846 1106 1479 1199">• A disparity in latency among drives or volumes of a common volume group could indicate a slow drive.

Performance data	Implications for performance tuning
IOPS	<p>Factors that affect input/output operations per second (IOPS or IOs/sec) include these items:</p> <ul style="list-style-type: none"> • Access pattern (random or sequential) • I/O size • RAID level • Cache block size • Whether read caching is enabled • Whether write caching is enabled • Dynamic cache read prefetch • Segment size • The number of drives in the volume groups or storage array <p>The higher the cache hit rate, the higher I/O rates will be. Higher write I/O rates are experienced with write caching enabled compared to disabled. In deciding whether to enable write caching for an individual volume, look at the current IOPS and the maximum IOPS. You should see higher rates for sequential I/O patterns than for random I/O patterns. Regardless of your I/O pattern, enable write caching to maximize the I/O rate and to shorten the application response time.</p> <p>You can see performance improvements caused by changing the segment size in the IOPS statistics for a volume. Experiment to determine the optimal segment size, or use the file system size or database block size.</p>
MiB/s	<p>Transfer or throughput rates are determined by the application I/O size and the I/O rate. Generally, small application I/O requests result in a lower transfer rate but provide a faster I/O rate and shorter response time. With larger application I/O requests, higher throughput rates are possible.</p> <p>Understanding your typical application I/O patterns can help you determine the maximum I/O transfer rates for a specific storage array.</p>

Performance data	Implications for performance tuning
CPU	<p>This value is a percentage of processing capacity that is being used.</p> <p>You might notice a disparity in the CPU usage of the same types of objects. For example, the CPU usage of one controller is heavy or is increasing over time while that of the other controller is lighter or more stable. In this case, you might want to change the controller ownership of one or more volumes to the controller with the lower CPU percentage.</p> <p>You might want to monitor CPU across the storage array. If CPU continues to increase over time while application performance decreases, you might need to add storage arrays. By adding storage arrays to your enterprise, you can continue to meet application needs at an acceptable performance level.</p>
Headroom	<p>Headroom refers to the remaining performance capability of the controllers, the controller host channels, and the controller drive channels. This value is expressed as a percentage and represents the gap between the maximum possible performance these objects are able to deliver and the current performance levels.</p> <ul style="list-style-type: none"> • For the controllers, headroom is a percentage of maximum possible IOPS. • For the channels, headroom is a percentage of maximum throughput, or MiB/s. Read throughput, write throughput, and bidirectional throughput are included in the calculation.

FAQs

How do performance statistics for individual volumes relate to the total?

The statistics for pools and volume groups are calculated by aggregating all volumes, including reserved capacity volumes.

Reserved capacity is used internally by the storage system to support thin volumes, snapshots, and asynchronous mirroring, and are not visible to I/O hosts. As a result, the pool, controller, and storage array statistics may not add up to be the sum of the viewable volumes.

However, for application and workload statistics, only the visible volumes are aggregated.

Why does data display as zero in the graphs and table?

When a zero is displayed for a data point in the graphs and table, it means there is no I/O activity for the object for that point in time. This situation could occur because the host is

not initiating I/O to that object, or it could be a problem with the object itself.

The historical data for the object is still available for viewing. The graphs and table will show non-zero data once I/O activity begins occurring for the object.

The following table lists the most common reasons why a data point value might be zero for any given object.

Array-level object type	Reason data displays as zero
Volume	<ul style="list-style-type: none">• Volume had no host assignment.
Volume group	<ul style="list-style-type: none">• Volume group is being imported.• Volume group does not contain a volume that is assigned to a host, and volume group does not contain any reserved capacity.
Drive	<ul style="list-style-type: none">• Drive has failed.• Drive has been removed.• Drive is in an unknown state.
Controller	<ul style="list-style-type: none">• Controller is offline.• Controller has failed.• Controller has been removed.• Controller is in an unknown state.
Storage array	<ul style="list-style-type: none">• Storage array does not contain volumes.

What does the Latency graph show?

The **Latency** graph provides latency statistics, in milliseconds (ms), for volumes, volume groups, pools, applications, and workloads. This graph appears in the Logical View, Physical View, and Applications & Workloads View tabs.

Latency refers to any delay that occurs as data is read or written. Hover your cursor over a point on the graph to view the following values, in milliseconds (ms), for that point in time:

- Read time.
- Write time.
- Average I/O size.

What does the IOPS graph show?

The **IOPS** graph displays statistics for input/output operations per second. On the **Home** page, this graph displays statistics for the storage array. In the Logical View, Physical View, and Applications & Workloads View tabs of the **Performance** tile, this graph displays statistics for the storage array, volumes, volume groups, pools, applications, and

workloads.

IOPS is an abbreviation for *input/output (I/O) operations per second*. Hover your cursor over a point on the graph to view the following values for that point in time:

- Number of read operations.
- Number of write operations.
- Total read and write operations combined.

What does the MiB/s graph show?

The **MiB/s** graph displays transfer speed statistics in mebibytes per second. On the **Home** page, this graph displays statistics for the storage array. In the Logical View, Physical View, and Applications & Workloads View tabs of the **Performance** tile, this graph displays statistics for the storage array, volumes, volume groups, pools, applications, and workloads.

MiB/s is an abbreviation for *mebibytes per second*, or 1,048,576 bytes per second. Hover your cursor over a point on the graph to view the following values for that point in time:

- The amount of data read.
- The amount of data written.
- The combined total amount of data read and written.

What does the CPU graph show?

The CPU graph displays processing capacity statistics for each controller (controller A and controller B). CPU is an abbreviation for *central processing unit*. On the **Home** page, this graph displays statistics for the storage array. On the Physical View tab of the **Performance** tile, this graph displays statistics for the storage array and drives.

The CPU graph shows the percentage of CPU processing capacity being used against operations on the array. Even when no external I/O is occurring, the CPU utilization percentage can be non-zero because the storage operating system might be doing background operations and monitoring. Hover your cursor over a point on the graph to view a percentage of processing capability being used at that point in time.

What does the Headroom graph show?

The Headroom graph is related to remaining performance capability for the storage array controllers. This graph is visible on the **Home** page and on the Physical View tab of the **Performance** tile.

The Headroom graph shows the remaining performance capability of the physical objects in the storage system. Hover your cursor over a point on the graph to view the percentages of IOPS and MiB/s capability remaining for controller A and for controller B.

Snapshots

Concepts


Overview of snapshot storage

A snapshot image is a logical, read-only copy of volume content, captured at a particular point in time. You can use snapshots to protect against data loss.

Snapshot images also are useful for test environments. By creating a virtual copy of data, you can test data using the snapshot without altering the actual volume itself. In addition, hosts do not have write access to snapshot images, so your snapshots are always a secure backup resource.

The Snapshots feature uses copy-on-write technology to store snapshot images and use allocated reserved capacity. As snapshots are created, the Snapshots feature stores image data as follows:

- When a snapshot image is created, it exactly matches the base volume. The Snapshots feature uses copy-on-write technology. After the snapshot is taken, the first write to any block or set of blocks on the base volume causes the original data to be copied to the reserved capacity before writing the new data to the base volume.
- Subsequent snapshots include only changed data blocks. Before data is overwritten on the base volume, the Snapshots feature uses its copy-on-write technology to save the required images of the affected sectors to the snapshot reserved capacity.



¹ Base volume (physical disk capacity); ² Snapshots (logical disk capacity); ³ Reserved capacity (physical disk capacity)

- The reserved capacity stores original data blocks for portions of the base volume that have changed after the snapshot was taken and includes an index for tracking changes. Generally, the size of the reserved capacity defaults to 40 percent of the base volume. (If you need more reserved capacity, you can increase reserved capacity.)
- Snapshot images are stored in a specific order, based on their timestamp. Only the oldest snapshot image of a base volume is available for manual deletion.

To restore data to a base volume, you can use either a snapshot volume or snapshot image:

- **Snapshot volume** — If you need to retrieve deleted files, create a snapshot volume from a known good snapshot image, and then assign it to the host.
- **Snapshot image** — If you need to restore a base volume to a specific point-in-time, use a previous snapshot image to roll back data to the base volume.

Requirements and guidelines for snapshots

When creating and using snapshots, review the following requirements and guidelines.

Snapshot images and snapshot groups

- Each snapshot image is associated with exactly one snapshot group.
- A snapshot group is created the first time you create a scheduled or instant snapshot image for an associated object. This creates reserved capacity.

You can view snapshot groups from the **Pools & Volume Groups** page.

- Scheduled snapshot images do not occur when the storage array is offline or powered off.
- If you delete a snapshot group that has a snapshot schedule, the snapshot schedule is also deleted.
- If you have a snapshot volume that you no longer need, you can reuse it, along with any associated reserved capacity, instead of deleting it. This creates a different snapshot volume of the same base volume. You can re-associate the snapshot volume or snapshot consistency group snapshot volume with the same snapshot image or a different snapshot image, as long as the snapshot image is in the same base volume.

Snapshot consistency group

- A snapshot consistency group contains one snapshot group for each volume that is a member of the snapshot consistency group.
- You can associate a snapshot consistency group with only one schedule.
- If you delete a snapshot consistency group that has a snapshot schedule, the snapshot schedule is also deleted.
- You cannot individually manage a snapshot group that is associated with a snapshot consistency group. Instead, you must perform the manage operations (create snapshot image, delete snapshot image or snapshot group, and rollback snapshot image) at the snapshot consistency group level.

Base volume

- A snapshot volume must have the same Data Assurance (DA) and security settings as the associated base volume.
- You cannot create a snapshot volume of a failed base volume.
- If the base volume resides on a volume group, the member volumes for any associated snapshot consistency group can reside on either a pool or volume group.
- If a base volume resides on a pool, all member volumes for any associated snapshot consistency group must reside on the same pool as the base volume.

Reserved capacity

- Reserved capacity is associated with only one base volume.
- Using a schedule can result in a large number of snapshot images. Make sure you have sufficient reserved capacity for scheduled snapshots.
- The reserved capacity volume for a snapshot consistency group must have the same Data Assurance (DA) and security settings as its associated base volume for the snapshot consistency group member volume.

Pending snapshot images

Snapshot image creation might remain in a Pending state in the following conditions:

- The base volume that contains this snapshot image is a member of an asynchronous mirror group.
- The base volume is currently in a synchronization operation. The snapshot image creation completes as soon as the synchronization operation is complete.

Maximum number of snapshot images

- If a volume is a member of a snapshot consistency group, System Manager creates a snapshot group for that member volume. This snapshot group counts towards the maximum allowable number of snapshot groups per base volume.
- If you attempt to create a snapshot image on a snapshot group or snapshot consistency group, but the associated group has reached its maximum number of snapshot images, you have two options:
 - Enable automatic deletion for the snapshot group or snapshot consistency group.
 - Manually delete one or more snapshot images from the snapshot group or snapshot consistency group and retry the operation.

Auto-deletion

If the snapshot group or snapshot consistency group is enabled for automatic deletion, System Manager deletes the oldest snapshot image when the system creates a new one for the group.

Rollback operation


- You cannot perform the following actions when a rollback operation is in progress:
 - Delete the snapshot image that is being used for the rollback.
 - Create a new snapshot image for a base volume that is participating in a rollback operation.
 - Change the associated snapshot group's Repository-Full Policy.
- You cannot start a rollback operation when any of these operations are in progress:
 - Capacity expansion (adding capacity to a pool or volume group)
 - Volume expansion (increasing the capacity of a volume)
 - RAID level change for a volume group
 - Segment size change for a volume
- You cannot start a rollback operation if the base volume is participating in a volume copy.
- You cannot start a rollback operation if the base volume is a secondary volume in a remote mirror.
- A rollback operation fails if any of the used capacity in the associated snapshot repository volume has unreadable sectors.

Base volumes, reserved capacity, and snapshot groups

A *base volume* is the volume used as the source for a snapshot image. A base volume can be either a thick volume or a thin volume and can reside in either a pool or volume group.

To take snapshots of the base volume, you can create an instant image at any time, or you can automate the process by defining a regular schedule for snapshots.

The following figure shows the relationship between snapshot objects and the base volume.



¹ Base volume; ² Snapshot objects in the group (images and reserved capacity); ³ Reserved capacity for the snapshot group.

Reserved capacity and snapshot groups

System Manager organizes snapshot images into *snapshot groups*. When System Manager establishes the snapshot group, it automatically creates associated *reserved capacity* to hold the snapshot images for the group and to keep track of subsequent changes to additional snapshots.

If the base volume resides in a volume group, the reserved capacity can be located in either a pool or volume group. If the base volume resides in a pool, the reserved capacity must be located in the same pool as the base volume.

Snapshot groups require no user action, but you can adjust reserved capacity on a snapshot group at any time. Additionally, you might be prompted to create reserved capacity when the following conditions are met:

- Any time you take a snapshot of a base volume that does not yet have a snapshot group, System Manager automatically creates a snapshot group. This action also creates reserved capacity for the base volume that is used to store subsequent snapshot images.
- Any time you create a snapshot schedule for a base volume, System Manager automatically creates a snapshot group.

Auto-deletion

When working with snapshots, use the default option to have auto-deletion turned on. Auto-deletion automatically deletes the oldest snapshot image when the snapshot group reaches the snapshot group limit of 32 images. If you turn off auto-deletion, then snapshot group limits are eventually exceeded, and you must take manual actions to configure snapshot group settings and manage reserved capacity.

Snapshot schedules and snapshot consistency groups

Use schedules for collection of snapshot images, and use snapshot consistency groups to manage multiple base volumes.

To easily manage snapshot operations for base volumes, you can use the following features:

- **Snapshot schedule** — Automate snapshots for a single base volume.
- **Snapshot consistency group** — Manage multiple base volumes as one entity.


Snapshot schedule

If you want to automatically take snapshots for a base volume, you can create a schedule. For example, you can define a schedule that takes snapshot images every Saturday at midnight, on the first of every month, or on any dates and times you decide. After the maximum of 32 snapshots is reached for a single schedule, you can suspend scheduled snapshots, create more reserved capacity, or you can delete snapshots. Snapshots can be deleted manually or by automating the deletion process. After a snapshot image is deleted, additional reserved capacity is available for reuse.

Snapshot consistency group

You create a snapshot consistency group when you want to make sure snapshot images are taken on multiple volumes at the same time. Snapshot image actions are performed on the snapshot consistency group as a whole. For example, you can schedule synchronized snapshots of all volumes with the same timestamp. Snapshot consistency groups are ideal for applications that span multiple volumes, such as database applications that store logs on one volume and the database files on another volume.

The volumes included in a snapshot consistency group are called member volumes. When you add a volume to a consistency group, System Manager automatically creates new reserved capacity that corresponds to that member volume. You can define a schedule to automatically create a snapshot image of each member volume.




¹ Reserved capacity; ² Member volume; ³ Consistency group snapshot images

Snapshot volumes

You can create a snapshot volume and assign it to a host if you want to read or write snapshot data. The snapshot volume shares the same characteristics as the base volume (RAID level, I/O characteristics, and so on).

When you create a snapshot volume, you can designate it as *read-only* or *read-write accessible*.

When you create read-only snapshot volumes, you do not need to add reserved capacity. When you create read-write snapshot volumes, you must add reserved capacity to provide write-access.



¹ Base volume; ² Host; ³ Read-only snapshot volume; ⁴ Read-write snapshot volume; ⁵ Reserved capacity

Snapshot rollback


A rollback operation returns a base volume to a previous state, determined by the selected snapshot.

For the rollback, you can select a snapshot image from either of the following sources:

- **Snapshot image rollback**, for a full restore of a base volume.
- **Snapshot consistency group rollback**, which can be used to roll back one or more volumes.

During the rollback, the Snapshots feature preserves all snapshot images in the group. It also allows the host to access the base volume during this process, if needed for I/O operations.

When a rollback is launched, a background process sweeps through the logical block addresses (LBAs) for the base volume, and then finds copy-on-write data in the rollback snapshot image to be restored. Because the base volume is host-accessible for reads and writes, and all previously written data is available immediately, the reserved capacity volume must be large enough to contain all changes while the rollback is processing. The data transfer continues as a background operation until the rollback completes.



¹ Base volume; ² Snapshot objects in a group; ³ Snapshot group reserved capacity

Snapshot terminology

Learn how the snapshot terms apply to your storage array.

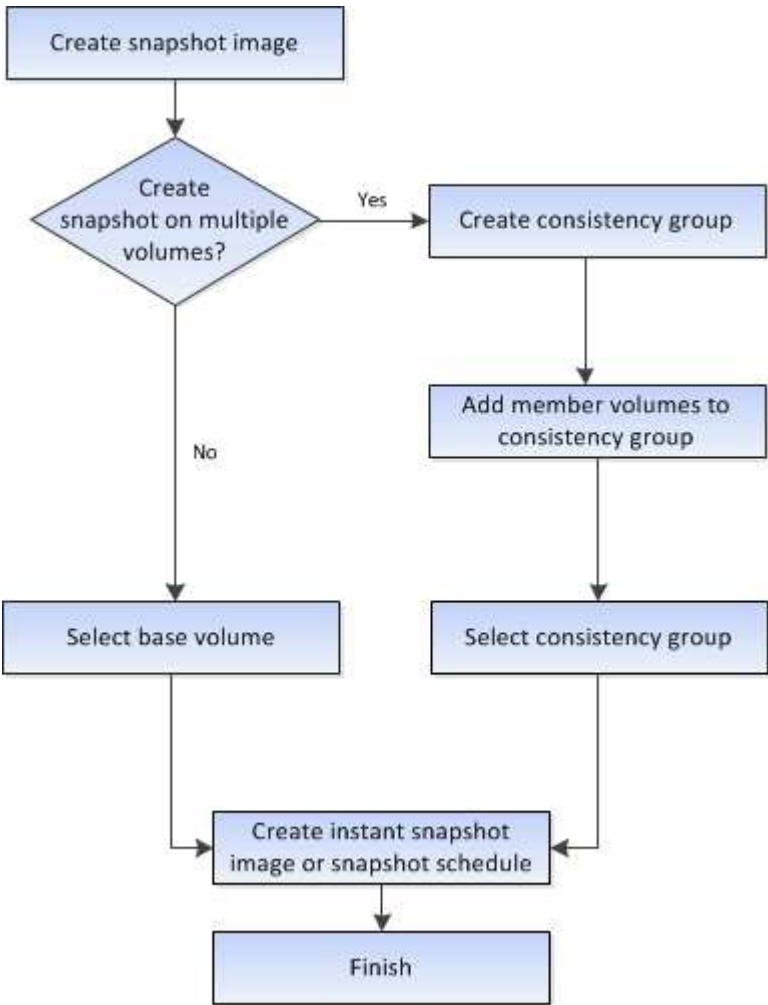
Term	Description
Snapshots feature	The Snapshots feature is used to create and manage images of volumes.

Term	Description
Snapshot image	A snapshot image is a logical copy of volume data, captured at a particular point-in-time. Like a restore point, snapshot images allow you to roll back to a known good data set. Although the host can access the snapshot image, it cannot directly read or write to it.
Base volume	A base volume is the source from which a snapshot image is created. It can be a thick or thin volume and is typically assigned to a host. The base volume can reside in either a volume group or disk pool.
Snapshot volume	A snapshot volume allows the host to access data in the snapshot image. The snapshot volume contains its own reserved capacity, which saves any modifications to the base volume without affecting the original snapshot image.
Snapshot group	A snapshot group is a collection of snapshot images from a single base volume.
Reserved capacity volume	A reserved capacity volume tracks which data blocks of the base volume are overwritten and the preserved content of those blocks.
Snapshot schedule	A snapshot schedule is a timetable for creating automated snapshot images. Through the schedule, you can control the frequency of image creations.
Snapshot consistency group	A snapshot consistency group is a collection of volumes that are treated as a single entity when a snapshot image is created. Each of these volumes has its own snapshot image, but all the images are created at the same point in time.
Snapshot consistency group member volume	Each volume that belongs to a snapshot consistency group is referred to as a member volume. When you add a volume to a snapshot consistency group, System Manager automatically creates a new snapshot group that corresponds to this member volume.
Rollback	A rollback is the process of returning data in a base volume to a previous point in time.
Reserved capacity	Reserved capacity is the physical allocated capacity that is used for any copy service operation and storage object. It is not directly readable by the host.


Workflow for creating snapshot images and snapshot volumes

In System Manager, you can create snapshot images and snapshot volumes by following these steps.

Workflow for creating snapshot images



Workflow for creating snapshot volumes



How tos

Create snapshots and snapshot objects

Create snapshot image

You can manually create a snapshot image from a base volume or snapshot consistency group. This is also called an *instant snapshot* or *instant image*.

Before you begin

- The base volume must be optimal.
- The drive must be optimal.
- The snapshot group cannot be designated as “reserved.”
- The reserved capacity volume must have the same Data Assurance (DA) settings as the associated base volume for the snapshot group.

Steps

1. Do one of the following actions to create a snapshot image:
 - Select **Storage > Volumes**. Select the object (base volume or snapshot consistency group), and then

select **Copy Services › Create instant snapshot**.

- Select **Storage › Snapshots**. Select the **Snapshot Images** tab, and then select **Create › Instant snapshot image**. The **Create Snapshot Image** dialog box appears. Select the object (base volume or snapshot consistency group), and then click **Next**. If a previous snapshot image was created for the volume or snapshot consistency group, then the system creates the instant snapshot immediately. Otherwise, if this is the first time a snapshot image is created for the volume or snapshot consistency group, the **Confirm Create Snapshot Image** dialog box appears.

2. Click **Create** to accept the notification that reserved capacity is needed and to proceed to the **Reserve Capacity** step.

The **Reserve Capacity** dialog box appears.

3. Use the spinner box to adjust the capacity percentage, and then click **Next** to accept the candidate volume highlighted in the table.

The **Edit Settings** dialog box appears.

4. Select the settings for the snapshot image as appropriate, and confirm that you want to perform the operation.

Field Details

Setting	Description
Snapshot image settings	
Snapshot image limit	Keep the check box selected if you want snapshot images automatically deleted after the specified limit; use the spinner box to change the limit. If you clear this check box, snapshot image creation stops after 32 images.
Reserved capacity settings	
Alert me when...	<p>Use the spinner box to adjust the percentage point at which the system sends an alert notification when the reserved capacity for a snapshot group is nearing full.</p> <p>When the reserved capacity for the snapshot group exceeds the specified threshold, use the advance notice to increase reserved capacity or to delete unnecessary objects before the remaining space runs out.</p>
Policy for full reserved capacity	<p>Choose one of the following policies:</p> <ul style="list-style-type: none">• Purge oldest snapshot image: The system automatically purges the oldest snapshot image in the snapshot group, which releases the snapshot image reserved capacity for reuse within the group.• Reject writes to base volume: When the reserved capacity reaches its maximum defined percentage, the system rejects any I/O write request to the base volume that triggered the reserved capacity access.

Results

- System Manager displays the new snapshot image in the Snapshot Images table. The table lists the new image by timestamp and associated base volume or snapshot consistency group.
- Snapshot creation might remain in a Pending state because of the following conditions:
 - The base volume that contains this snapshot image is a member of an asynchronous mirror group.
 - The base volume is currently in a synchronization operation. The snapshot image creation completes as soon as the synchronization operation is complete.

Schedule snapshot images

You create a snapshot schedule to enable recovery in case of a problem with the base volume and to perform scheduled backups. Snapshots of base volumes or snapshot consistency groups can be created on a daily, weekly, or monthly schedule, at any time of day.

Before you begin

The base volume must be Optimal.

About this task

This task describes how to create a snapshot schedule for an existing snapshot consistency group or base volume.



You also can create a snapshot schedule at the same time you create a snapshot image of a base volume or snapshot consistency group.

Steps

1. Do one of the following actions to create a snapshot schedule:

- Select **Storage > Volumes**.

Select the object (volume or snapshot consistency group) for this snapshot schedule, and then select **Copy Services > Create snapshot schedule**.

- Select **Storage > Snapshots**.

Select the **Schedules** tab, and then click **Create**.

2. Select the object (volume or snapshot consistency group) for this snapshot schedule, and then click **Next**.

The **Create Snapshot Schedule** dialog box appears.

3. Do one of the following actions:

- **Use a previously defined schedule from another snapshot object.**

Make sure advanced options are displayed. Click **Show more options**. Click **Import Schedule**, select the object with the schedule you want to import, and then click **Import**.

- **Modify the basic or advanced options.**

In the upper right of the dialog box, click **Show more options** to display all options, and then refer to the following table.

Field Details

Field	Description
Basic settings	
Select days	Select individual days of the week for snapshot images.
Start time	From the drop-down list, select a new start time for the daily snapshots (selections are provided in half-hour increments). The start time defaults to one half-hour ahead of the current time.
Time zone	From the drop-down list, select your array's time zone.
Advanced settings	
Day / month	<p>Choose one of the following options:</p> <ul style="list-style-type: none"> • Daily / Weekly — Select individual days for synchronization snapshots. You also can select the Select all days check box in the upper right if you want a daily schedule. • Monthly / Yearly — Select individual months for synchronization snapshots. In the On day(s) field, enter the days of the month for synchronizations to occur. Valid entries are 1 through 31 and Last. You can separate multiple days with a comma or semi-colon. Use a hyphen for inclusive dates. For example: 1,3,4,10-15,Last. You also can select the Select all months check box in the upper right if you want a monthly schedule.
Start time	From the drop-down list, select a new start time for the daily snapshots (selections are provided in half-hour increments). The start time defaults to one half-hour ahead of the current time.
Time zone	From the drop-down list, select your array's time zone.
Snapshots per day / Time between snapshots	Select the number of snapshot images to create per day. If you select more than one, also select the time between snapshot images. For multiple snapshot images, be sure that you have adequate reserved capacity.

Field	Description
Create snapshot image right now?	Select this check box to create an instant image in addition to the automatic images you are scheduling.
Start/End date or No end date	Enter the start date for synchronizations to begin. Also enter an end date or select No end date .

4. Do one of the following actions:

- If the object is a snapshot consistency group, click **Create** to accept the settings and create the schedule.
- If the object is a volume, click **Next** to allocate reserved capacity for the snapshot images.

The volume candidate table displays only the candidates that support the reserved capacity specified. Reserved capacity is the physical allocated capacity that is used for any copy service operation and storage object. It is not directly readable by the host.

5. Use the spinner box to allocate the reserved capacity for the snapshot images. Do one of the following actions:

- **Accept the default settings.**

Use this recommended option to allocate the reserved capacity for the snapshot images with the default settings.

- **Allocate your own reserved capacity settings to meet your data storage needs.**

If you change the default reserved capacity setting, click **Refresh Candidates** to refresh the candidate list for the reserved capacity you specified.

Allocate the reserved capacity using the following guidelines:

- The default setting for reserved capacity is 40% of the capacity of the base volume. Usually this capacity is sufficient.
- The capacity needed varies, depending on the frequency and size of I/O writes to the volumes and the quantity and duration of snapshot image collection.

6. Click **Next**.

The Edit Settings dialog box appears.

7. Edit the settings for the snapshot schedule as needed, and then click **Finish**.

Field Details

Setting	Description
Snapshot image limit	
Enable automatic deletion of snapshot images when...	Keep the check box selected if you want snapshot images automatically deleted after the specified limit; use the spinner box to change the limit. If you clear this check box, snapshot image creation stops after 32 images.
Reserved capacity settings	
Alert me when...	<p>Use the spinner box to adjust the percentage point at which the system sends an alert notification when the reserved capacity for a schedule is nearing full.</p> <p>When the reserved capacity for the schedule exceeds the specified threshold, use the advance notice to increase reserved capacity or to delete unnecessary objects before the remaining space runs out.</p>
Policy for full reserved capacity	<p>Choose one of the following policies:</p> <ul style="list-style-type: none">• Purge oldest snapshot image — The system automatically purges the oldest snapshot image, which releases the snapshot image reserved capacity for reuse within the snapshot group.• Reject writes to base volume — When the reserved capacity reaches its maximum defined percentage, the system rejects any I/O write request to the base volume that triggered the reserved capacity access.

Create snapshot consistency group

To ensure that you have consistent copies, you can create a set of multiple volumes called a *snapshot consistency group*. This group allows you to make snapshot images of all the volumes at the same time for consistency. Each volume that belongs to a snapshot consistency group is referred to as a *member volume*. When you add a volume to a snapshot consistency group, the system automatically creates a new snapshot group that corresponds to this member volume.

About this task

The snapshot consistency group creation sequence lets you select member volumes for the group and allocate capacity to the member volumes.

The process to create a snapshot consistency group is a multi-step procedure:

- [Step 1: Add members](#)
- [Step 2: Reserve capacity](#)
- [Step 3: Edit settings](#)

Step 1: Add members

You can select members to specify a collection of volumes that comprise the snapshot consistency group. Any actions you perform on the snapshot consistency group extend uniformly to selected member volumes.

Before you begin

The member volumes must be Optimal.

Steps

1. Select **Storage** › **Snapshots**.
2. Click the **Snapshot Consistency Groups** tab.
3. Select **Create** › **Snapshot consistency group**.

The **Create Snapshot Consistency Group** dialog box appears.

4. Select the volume(s) to be added as member volumes to the snapshot consistency group.
5. Click **Next**, and go to [Step 2: Reserve capacity](#).

Step 2: Reserve capacity

You must associate reserved capacity to the snapshot consistency group. System Manager suggests the volumes and capacity based on the properties of the snapshot consistency group. You can accept the recommended reserved capacity configuration or customize the allocated storage.

About this task

On the **Reserve Capacity** dialog box, the volume candidate table displays only the candidates that support the reserved capacity specified. Reserved capacity is the physical allocated capacity that is used for any copy service operation and storage object. It is not directly readable by the host.

Steps

1. Use the spinner box to allocate the reserved capacity for the snapshot consistency group. Do one of the following actions:
 - **Accept the default settings.**

Use this recommended option to allocate the reserved capacity for each member volume with the default settings.
 - **Allocate your own reserved capacity settings to meet your data storage needs.**

Allocate the reserved capacity using the following guidelines.

- The default setting for reserved capacity is 40% of the capacity of the base volume. Usually this capacity is sufficient.
 - The capacity needed varies, depending on the frequency and size of I/O writes to the volumes and the quantity and duration of snapshot image collection.
2. (Optional) If you change the default reserved capacity setting, click **Refresh Candidates** to refresh the candidate list for the reserved capacity you specified.
 3. Click **Next**, and go to [Step 3: Edit settings](#).

Step 3: Edit settings

You can accept or choose automatic deletion settings and reserved capacity alert thresholds for the snapshot consistency group.

Steps

1. Accept or change the default settings for the snapshot consistency group as appropriate.

Field Details

Setting	Description
Snapshot consistency group settings	
Name	Specify the name for the snapshot consistency group.
Enable automatic deletion of snapshot images when...	Keep the check box selected if you want snapshot images automatically deleted after the specified limit; use the spinner box to change the limit. If you clear this check box, snapshot image creation stops after 32 images.
Reserved capacity settings	
Alert me when...	<p>Use the spinner box to adjust the percentage point at which the system sends an alert notification when the reserved capacity for a snapshot consistency group is nearing full.</p> <p>When the reserved capacity for the snapshot consistency group exceeds the specified threshold, use the advance notice to increase reserved capacity or to delete unnecessary objects before the remaining space runs out.</p>
Policy for full reserved capacity	<p>Choose one of the following policies:</p> <ul style="list-style-type: none">• Purge oldest snapshot image — The system automatically purges the oldest snapshot image in the snapshot consistency group, which releases the snapshot image reserved capacity for reuse within the group.• Reject writes to base volume — When the reserved capacity reaches its maximum defined percentage, the system rejects any I/O write request to the base volume that triggered the reserved capacity access.

2. After you are satisfied with your snapshot consistency group configuration, click **Finish**.

Create snapshot volume

You create a snapshot volume to provide host access to a snapshot image of a volume or snapshot consistency group. You can designate the snapshot volume as either read-only or read-write.

About this task

The snapshot volume creation sequence lets you create a snapshot volume from a snapshot image and provides options to allocate reserved capacity if the volume is read/write. A snapshot volume can be designated as one of the following:

- A read-only snapshot volume provides a host application with read access to a copy of the data contained in the snapshot image, but without the ability to modify the snapshot image. A read-only snapshot volume does not have associated reserved capacity.
- A read-write snapshot volume provides the host application with write access to a copy of the data contained in the snapshot image. It has its own reserved capacity that is used to save any subsequent modifications made by the host application to the base volume without affecting the referenced snapshot image.

The process to create a snapshot volume is a multi-step procedure:

- [Step 1: Review members](#)
- [Step 2: Assign to host](#)
- [Step 3: Reserve capacity](#)
- [Step 4: Edit settings](#)

Step 1: Review members

You can select either a snapshot image of a base volume or a snapshot consistency group. If you select a snapshot consistency group snapshot image, the member volumes of the snapshot consistency group appear for review.

Steps

1. Select **Storage** > **Snapshots**.
2. Select the **Snapshot Volumes** tab.
3. Select **Create**. The **Create Snapshot Volume** dialog box appears.
4. Select the snapshot image (volume or snapshot consistency group) you want to convert into a snapshot volume, and then click **Next**. Use a text entry in the **Filter** field to narrow down the list.

If the selection was for a snapshot consistency group snapshot image, the **Review Members** dialog box appears.

On the **Review Members** dialog box, review the list of volumes that are selected for conversion to snapshot volumes, and then click **Next**.

5. Go to [Step 2: Assign to host](#).

Step 2: Assign to host

You select a specific host or host cluster to assign it to the snapshot volume. This assignment grants a host or host cluster access to the snapshot volume. You can choose to assign a host later, if needed.

Before you begin

- Valid hosts or host clusters exist under the **Hosts** page.
- Host port identifiers must have been defined for the host.

- Before creating a DA-enabled volume, verify that your planned host connection supports the Data Assurance (DA) feature. If any of the host connections on the controllers in your storage array do not support DA, the associated hosts cannot access data on DA-enabled volumes.

About this task

When you assign volumes, keep these guidelines in mind:

- A host's operating system can have specific limits on how many volumes the host can access.
- You can define one host assignment for each snapshot volume in the storage array.
- Assigned volumes are shared between controllers in the storage array.
- The same logical unit number (LUN) cannot be used twice by a host or a host cluster to access a snapshot volume. You must use a unique LUN.



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

Steps

1. On the **Assign to Host** dialog box, select the host or host cluster that you want to assign to the new volume. If you want to create the volume without assigning a host, select **Assign later** from the drop-down list.
2. Select the access mode. Choose one of the following:
 - **Read/write** — This option provides the host with read/write access to the snapshot volume and requires reserved capacity.
 - **Read only** — This option provides the host with read-only access to the snapshot volume and does not require reserved capacity.
3. Click **Next**, and do one of the following:
 - If your snapshot volume is read/write, the **Review Capacity** dialog box appears. Go to [Step 3: Reserve capacity](#).
 - If your snapshot volume is read only, the **Edit Priority** dialog box appears. Go to [Step 4: Edit settings](#).

Step 3: Reserve capacity

You must associate reserved capacity to a read/write snapshot volume. System Manager suggests the volumes and capacity based on the properties of the base volume or snapshot consistency group. You can accept the recommended reserved capacity configuration or customize the allocated storage.

About this task

You can increase or decrease the reserved capacity for the snapshot volume as needed. If you find that the snapshot reserved capacity is larger than you need, you can reduce its size to free up space that is needed by other logical volumes.

Steps

1. Use the spinner box to allocate the reserved capacity for the snapshot volume.

The Volume Candidate table displays only the candidates that support the reserved capacity specified.

Do one of the following actions:

- **Accept the default settings.**

Use this recommended option to allocate the reserved capacity for the snapshot volume with the default settings.

- **Allocate your own reserved capacity settings to meet your data storage needs.**

If you change the default reserved capacity setting, click **Refresh Candidates** to refresh the candidate list for the reserved capacity you specified.

Allocate the reserved capacity using the following guidelines.

- The default setting for reserved capacity is 40% of the capacity of the base volume, and usually this capacity is sufficient.
 - The capacity needed varies, depending on the frequency and size of I/O writes to the volumes and the quantity and duration of snapshot image collection.
2. (Optional) If you are creating the snapshot volume for a snapshot consistency group, the option to **Change candidate** appears in the Reserved Capacity Candidates table. Click **Change candidate** to select an alternate reserved capacity candidate.
 3. Click **Next**, and go to [Step 4: Edit settings](#).

Step 4: Edit settings

You can change the settings for a snapshot volume such as its name, caching, reserved capacity alert thresholds, and so on.

About this task

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

Steps

1. Accept or change the settings for the snapshot volume as appropriate.

Field Details

Setting	Description
Snapshot volume settings	
Name	Specify the name for the snapshot volume.
Enable SSD Cache	Choose this option to enable read-only caching on SSDs.
Reserved capacity settings	
Alert me when...	Appears only for a read/write snapshot volume. Use the spinner box to adjust the percentage point at which the system sends an alert notification when the reserved capacity for a snapshot group is nearing full. When the reserved capacity for the snapshot group exceeds the specified threshold, use the advance notice to increase reserved capacity or to delete unnecessary objects before the remaining space runs out.

2. Review the snapshot volume configuration. Click **Back** to make any changes.
3. When you are satisfied with your snapshot volume configuration, click **Finish**.

Results

System Manager creates the snapshot volume in a normal state.

If the snapshot volume shows in a pending state, the base volume is a member of an asynchronous mirror group completing a synchronizing operation.

Manage snapshot schedules

Change the settings for a snapshot schedule

For a snapshot schedule, you can change automatic collection times or the frequency of collection.

About this task

You can import settings from an existing snapshot schedule, or you can modify settings as needed.

Because a snapshot schedule is associated to a snapshot group or snapshot consistency group, reserved capacity might be affected by changes to schedule settings.

Steps

1. Select **Storage > Snapshots**.
2. Click the **Schedules** tab.
3. Select the snapshot schedule that you want to change, and then click **Edit**.

The **Edit Snapshot Schedule** dialog box appears.

4. Do one of the following:
 - **Use a previously defined schedule from another snapshot object** — Click **Import Schedule**, select the object with the schedule you want to import, and then click **Import**.
 - **Edit the schedule settings** — Refer to Field Details below.

Field Details

Setting	Description
Day / month	<p>Choose one of the following options:</p> <ul style="list-style-type: none">• Daily / Weekly — Select individual days for synchronization snapshots. You also can select the Select all days check box in the upper right if you want a daily schedule.• Monthly / Yearly — Select individual months for synchronization snapshots. In the On day(s) field, enter the days of the month for synchronizations to occur. Valid entries are 1 through 31 and Last. You can separate multiple days with a comma or semi-colon. Use a hyphen for inclusive dates. For example: 1,3,4,10-15,Last. You also can select the Select all months check box in the upper right if you want a monthly schedule.
Start time	From the drop-down list, select a new start time for the daily snapshots. Selections are provided in half-hour increments. The start time defaults to one half-hour ahead of the current time.
Time zone	From the drop-down list, select your storage array's time zone.
Snapshots per day Time between snapshots	<p>Select the number of snapshot images to create per day.</p> <p>If you select more than one, also select the time between restore points. For multiple restore points, be sure that you have adequate reserved capacity.</p>
Start date End date No end date	Enter the start date for synchronizations to begin. Also enter an end date or select No end date .

5. Click **Save**.

Activate and suspend snapshot schedule

You can temporarily suspend scheduled collection of snapshot images when you need to conserve storage space. This method is more efficient than deleting and later re-creating

the snapshot schedule.

About this task

The state of the snapshot schedule stays suspended until you use the **Activate** option to resume scheduled snapshot activity.

Steps

1. Select **Storage > Snapshots**.
2. If it is not already displayed, click the **Schedules** tab.

The schedules are listed on the page.

3. Select an active snapshot schedule that you want to suspend, and then click **Activate / Suspend**.

The **State** column status changes to **Suspended**, and the snapshot schedule stops collection of all snapshot images.

4. To resume collecting snapshot images, select the suspended snapshot schedule that you want to resume, and then click **Activate / Suspend**.

The **State** column status changes to **Active**.

Delete snapshot schedule

If you no longer want to collect snapshot images, you can delete an existing snapshot schedule.

About this task

When you delete a snapshot schedule, the associated snapshot images are not deleted along with it. If you think the collection of snapshot images might be resumed at some point, you should suspend the snapshot schedule rather than delete it.

Steps

1. Select **Storage > Snapshots**.
2. Click the **Schedules** tab.
3. Select the snapshot schedule that you want to delete, and confirm the operation.

Results

The system removes all schedule attributes from the base volume or snapshot consistency group.

Manage snapshot images

View snapshot image settings

You can view the properties, status, reserved capacity, and associated objects assigned to each snapshot image.

About this task

Associated objects for a snapshot image include the base volume or snapshot consistency group for which this snapshot image is a restore point, the associated snapshot group, and any snapshot volumes created from the snapshot image. Use the snapshot settings to determine whether you want to copy or convert the snapshot

image.

Steps

1. Select **Storage › Snapshots**.
2. Click the **Snapshot Images** tab.
3. Select the snapshot image that you want to view, and then click **View Settings**.

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

4. View the settings for the snapshot image.

Start snapshot image rollback for a base volume

You can perform a rollback operation to change the content of a base volume to match the content that is saved in a snapshot image. The rollback operation does not change the content of the snapshot images that are associated with the base volume.

Before you begin

- Enough reserved capacity is available to start a rollback operation.
- The selected snapshot image is Optimal and the selected volume is Optimal.
- The selected volume does not have a rollback operation already in progress.

About this task

The rollback start sequence lets you start rollback on a snapshot image of a base volume while providing options to add storage capacity. You cannot start more than one rollback operation for a base volume at a time.



The host can immediately access the new rolled-back base volume, but the existing base volume does not allow the host read-write access after the rollback begins. You can create a snapshot of the base volume just before starting the rollback to preserve the pre-rollback base volume for recovery.

Steps

1. Select **Storage › Snapshots**.
2. Select the **Snapshot Images** tab.
3. Select the snapshot image, and then select **Rollback › Start**.

The **Confirm Start Rollback** dialog box appears.

4. **Optional:** Select the option to **Increase Capacity** if necessary.

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

- a. Use the spinner box to adjust the capacity percentage.

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

- b. Click **Increase**.

5. Confirm that you want to perform this operation, and then click **Rollback**.

Results

System Manager performs the following actions:

- Restores the volume with the content saved on the selected snapshot image.
- Makes the rolled-back volumes immediately available for host access. You do not need to wait for the rollback operation to complete.

After you finish

Select **Home** › **View Operations in Progress** to view the progress of the rollback operation.

If the rollback operation is not successful, the operation pauses. You can resume the paused operation and, if still unsuccessful, follow the Recovery Guru procedure to correct the problem or contact technical support.

Start snapshot image rollback for snapshot consistency group member volumes

You can perform a rollback operation to change the content of snapshot consistency group member volumes to match the content that is saved in a snapshot image. The rollback operation does not change the content of the snapshot images that are associated with the snapshot consistency group.

Before you begin

- Enough reserved capacity is available to start a rollback operation.
- The selected snapshot image is Optimal and the selected volume is Optimal.
- The selected volume does not have a rollback operation already in progress.

About this task

The rollback start sequence lets you start rollback on a snapshot image of a snapshot consistency group while providing options to add storage capacity. You cannot start more than one rollback operation for a snapshot consistency group at a time.



The host has immediate access to the new rolled-back volumes, but the existing member volumes no longer allow host read-write access after the rollback starts. You can create a snapshot image of the member volumes just before starting the rollback to preserve the pre-rollback base volumes for recovery purposes.

The process to start rollback of a snapshot image of a snapshot consistency group is a multi-step procedure:

- [Step 1: Select members](#)
- [Step 2: Review capacity](#)
- [Step 3: Edit priority](#)

Step 1: Select members

You must select the member volumes to be rolled back.

Steps

1. Select **Storage** › **Snapshots**.

2. Select the **Snapshot Images** tab.
3. Select the snapshot consistency group snapshot image, and then select **Rollback > Start**.

The **Start Rollback** dialog box appears.

4. Select the member volume or volumes.
5. Click **Next**, and do one of the following:
 - If any of the selected member volumes are associated with more than one reserved capacity object that stores snapshot images, the Review Capacity dialog box appears. Go to [Step 2: Review capacity](#).
 - If none of the selected member volumes are associated with more than one reserved capacity object that stores snapshot images, the Edit Priority dialog box appears. Go to [Step 3: Edit priority](#).

Step 2: Review capacity

If you selected member volumes associated to more than one reserved capacity object, such as a snapshot group and reserved capacity volume, you can review and increase reserved capacity for the rolled-back volume(s).

Steps

1. Next to any member volumes with very low (or zero) reserved capacity, click the **Increase capacity** link in the **Edit** column.

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

2. Use the spinner box to adjust the capacity percentage, and then click **Increase**.

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

3. Click **Next**, and go to [Step 3: Edit priority](#).

The Edit Priority dialog box appears.

Step 3: Edit priority

You can edit the priority of the rollback operation if needed.

About this task

The rollback priority determines how many system resources are dedicated to the rollback operation at the expense of system performance.

Steps

1. Use the slider to adjust the rollback priority as needed.
2. Confirm that you want to perform this operation, and then click **Finish**.

Results

System Manager performs the following actions:

- Restores the snapshot consistency group member volumes with the content saved on the selected

snapshot image.

- Makes the rolled-back volumes immediately available for host access. You do not need to wait for the rollback operation to complete.

After you finish

Select **Home** › **View Operations in Progress** to view the progress of the rollback operation.

If the rollback operation is not successful, the operation pauses. You can resume the paused operation and, if still unsuccessful, follow the Recovery Guru procedure to correct the problem or contact technical support.

Resume snapshot image rollback

If an error occurs during a snapshot image rollback operation, the operation is automatically paused. You can resume a rollback operation that is in a paused state.

Steps

1. Select **Storage** › **Snapshots**.
2. Click the **Snapshot Images** tab.
3. Highlight the paused rollback, and then select **Rollback** › **Resume**.

The operation resumes.

Results

System Manager performs the following actions:

- If the rollback operation resumes successfully, you can view the progress of the rollback operation in the **Operations in Progress** window.
- If the rollback operation is not successful, the operation pauses again. You can follow the Recovery Guru procedure to correct the problem or contact technical support.

Cancel snapshot image rollback

You can cancel an active rollback that is in progress (actively copying data), a pending rollback (in a pending queue awaiting resources to start), or a rollback that has been paused due to an error.

About this task

When you cancel a rollback operation that is in progress, the base volume reverts to an unusable state and appears as failed. Therefore, consider canceling a rollback operation only when recovery options exist for restoring the content of the base volume.



If the snapshot group on which the snapshot image resides has one or more snapshot images that have been automatically purged, the snapshot image used for the rollback operation might not be available for future rollbacks.

Steps

1. Select **Storage** › **Snapshots**.
2. Click the **Snapshot Images** tab.

3. Select the active or paused rollback, and then select **Rollback > Cancel**.

The Confirm Cancel Rollback dialog box appears.

4. Click **Yes** to confirm.

Results

System Manager stops the rollback operation. The base volume is usable but might have data that is inconsistent or not intact.

After you finish

After you cancel a rollback operation, you must take one of the following actions:

- Reinitialize the content of the base volume.
- Perform a new rollback operation to restore the base volume using either the same snapshot image that was used in the Cancel Rollback operation or a different snapshot image to perform the new rollback operation.

Delete snapshot image

You delete snapshot images to clean up the oldest snapshot image from a snapshot group or snapshot consistency group.

About this task

You can delete a single snapshot image, or you can delete snapshot images from snapshot consistency groups that have the same creation timestamp. You also can delete snapshot images from a snapshot group.

You cannot delete a snapshot image if it is not the oldest snapshot image for the associated base volume or snapshot consistency group.

Steps

1. Select **Storage > Snapshots**.
2. Click the **Snapshot Images** tab.
3. Select the snapshot image that you want to delete, and confirm that you want to perform the operation.

If you selected a snapshot image of a snapshot consistency group, select each member volume that you want to delete, and confirm that you want to perform the operation.

4. Click **Delete**.

Results

System Manager performs the following actions:

- Deletes the snapshot image from the storage array.
- Releases the reserved capacity for reuse within the snapshot group or snapshot consistency group.
- Disables all the associated snapshot volumes that exist for the deleted snapshot image.
- From a snapshot consistency group deletion, moves any member volume associated with the deleted snapshot image to a Stopped state.

Manage snapshot consistency groups

Add member volume to a snapshot consistency group

You can add a new member volume to an existing snapshot consistency group. When you add a new member volume, you also must reserve capacity for the member volume.

Before you begin

- The member volume must be Optimal.
- The snapshot consistency group must have less than the maximum number of allowable volumes (as defined by your configuration).
- Each reserved capacity volume must have the same Data Assurance (DA) and security settings as the associated member volume.

About this task

You can add standard volumes or thin volumes to the snapshot consistency group. The base volume can reside in either a pool or volume group.

Steps

1. Select **Storage > Snapshots**.
2. Select the **Snapshot Consistency Groups** tab.

The table appears and displays all the snapshot consistency groups associated with the storage array.

3. Select the snapshot consistency group you want to modify, and then click **Add Members**.

The Add Members dialog box appears.

4. Select the member volume(s) you want to add, and then click **Next**.

The Reserve Capacity step appears. The Volume Candidate table displays only the candidates that support the reserved capacity specified.

5. Use the spinner box to allocate the reserved capacity for the member volume. Do one of the following actions:

- **Accept the default settings.**

Use this recommended option to allocate the reserved capacity for the member volume with the default settings.

- **Allocate your own reserved capacity settings to meet your data storage needs.**

If you change the default reserved capacity setting, click **Refresh Candidates** to refresh the candidate list for the reserved capacity you specified.

Allocate the reserved capacity using the following guidelines.

- The default setting for reserved capacity is 40% of the capacity of the base volume, and usually this capacity is sufficient.
- The capacity needed varies, depending on the frequency and size of I/O writes to the volumes and the quantity and duration of snapshot image collection.

6. Click **Finish** to add the member volumes.

Remove a member volume from a snapshot consistency group

You can remove a member volume from an existing snapshot consistency group.

About this task

When you remove a member volume from a snapshot consistency group, System Manager automatically deletes the snapshot objects associated with that member volume.

Steps

1. Select **Storage > Snapshots**.
2. Click the **Snapshot Consistency Groups** tab.
3. Expand the snapshot consistency group you want to modify by selecting the plus (+) sign next to it.
4. Select the member volume that you want to remove, and then click **Remove**.
5. Confirm that you want to perform the operation, and then click **Remove**.

Results

System Manager performs the following actions:

- Deletes all snapshot images and snapshot volumes associated with the member volume.
- Deletes the snapshot group associated with the member volume.
- The member volume is not otherwise changed or deleted.

Change the settings for a snapshot consistency group

Change the settings for a snapshot consistency group when you want to change its name, automatic deletion settings, or the maximum number of allowed snapshot images.

Steps

1. Select **Storage > Snapshots**.
2. Click the **Snapshot Consistency Groups** tab.
3. Select the snapshot consistency group that you want to edit, and then click **View/Edit Settings**.

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

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

Field Details

Setting	Description
Snapshot consistency group settings	
Name	You can change the name for the snapshot consistency group.
Auto-deletion	Keep the check box selected if you want snapshot images automatically deleted after the specified limit; use the spinner box to change the limit. If you clear this check box, snapshot image creation stops after 32 images.
Snapshot image limit	You can change the maximum number of snapshot images allowed for a snapshot group.
Snapshot schedule	This field indicates whether a schedule is associated with the snapshot consistency group.
Associated objects	
Member volumes	You can view the quantity of member volumes associated with the snapshot consistency group.

5. Click **Save**.

Delete snapshot consistency group

You can delete snapshot consistency groups that are no longer needed.

Before you begin

Confirm that the images for all member volumes are no longer needed for backup or testing purposes.

About this task

This operation deletes all the snapshot images or schedules associated with the snapshot consistency group.

Steps

1. Select **Storage > Snapshots**.
2. Select the **Snapshot Consistency Groups** tab.
3. Select the snapshot consistency group that you want to delete, and then select **Uncommon Tasks > Delete**.

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

4. Confirm that you want to perform this operation, and then click **Delete**.

Results

System Manager performs the following actions:

- Deletes all existing snapshot images and snapshot volumes from the snapshot consistency group.
- Deletes all the associated snapshot images that exist for each member volume in the snapshot consistency group.
- Deletes all the associated snapshot volumes that exist for each member volume in the snapshot consistency group.
- Deletes all associated reserved capacity for each member volume in the snapshot consistency group (if selected).

Manage snapshot volumes

Convert snapshot volume to read-write mode

You can convert a read-only snapshot volume or a snapshot consistency group snapshot volume to read-write mode if needed. A snapshot volume that is converted to read-write accessible contains its own reserved capacity. This capacity is used to save any subsequent modifications made by the host application to the base volume without affecting the referenced snapshot image.

Steps

1. Select **Storage > Snapshots**.
2. Select the **Snapshot Volumes** tab.

The Snapshot Volumes table appears and displays all the snapshot volumes associated with the storage array.

3. Select the read-only snapshot volume you want to convert, and then click **Convert to Read/Write**.

The Convert to Read/Write dialog box appears with the **Reserve Capacity** step activated. The Volume Candidate table displays only the candidates that support the reserved capacity specified.

4. To allocate the reserved capacity for the read-write snapshot volume, do one of the following actions:
 - **Accept the default settings** — Use this recommended option to allocate the reserved capacity for the snapshot volume with the default settings.
 - **Allocate your own reserved capacity settings to meet your data storage needs** — Allocate the reserved capacity using the following guidelines.
 - The default setting for reserved capacity is 40% of the capacity of the base volume, and usually this capacity is sufficient.
 - The capacity needed varies, depending on the frequency and size of I/O writes to the volume.
5. Select **Next** to review or edit settings.

The **Edit Settings** dialog box appears.

6. Accept or specify the settings for the snapshot volume as appropriate, and then select **Finish** to convert the snapshot volume.

Field Details

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

Change the volume settings for a snapshot volume

You can change the settings for a snapshot volume or snapshot consistency group snapshot volume to rename it, enable or disable SSD caching, or change the host, host cluster, or logical unit number (LUN) assignment.

Steps

1. Select **Storage > Snapshots**.
2. Click the **Snapshot Volumes** tab.
3. Select the snapshot volume that you want to change, and then click **View/Edit Settings**.

The Snapshot Volume Settings dialog box appears.

4. View or edit the settings for the snapshot volume as appropriate.

Field Details

Setting	Description
Snapshot volume	
Name	You can change the name for the snapshot volume.
Assigned to	You can change the host or host cluster assignment for the snapshot volume.
LUN	You can change the LUN assignment for the snapshot volume.
SSD Cache	You can enable/disable read-only caching on solid state disks (SSDs).
Associated objects	
Snapshot image	You can view the snapshot images associated with the snapshot volume. A snapshot image is a logical copy of volume data, captured at a particular point-in-time. Like a restore point, snapshot images allow you to roll back to a known good data set. Although the host can access the snapshot image, it cannot directly read or write to it.
Base volume	You can view the base volume associated with the snapshot volume. A base volume is the source from which a snapshot image is created. It can be a thick or thin volume and is typically assigned to a host. The base volume can reside in either a volume group or disk pool.
Snapshot group	You can view the snapshot group associated with the snapshot volume. A snapshot group is a collection of snapshot images from a single base volume.

Copy snapshot volume

You can perform a Copy Volume process on a snapshot volume or a snapshot consistency group snapshot volume.

About this task

You can copy a snapshot volume to the target volume as performed in a normal Copy Volume operation.

However, snapshot volumes cannot remain online during the copy volume process.

Steps

1. Select **Storage › Snapshots**.
2. Select the **Snapshot Volumes** tab.

The Snapshot Volumes table appears and displays all the snapshot volumes associated with the storage array.

3. Select the snapshot volume that you want to copy, and then select **Copy Volume**.

The **Copy Volume** dialog box appears, prompting you to select a target.

4. Select the target volume to be used as the copy destination, and then click **Finish**.

Re-create snapshot volume

You can re-create a snapshot volume or a snapshot consistency group snapshot volume that you previously disabled. Re-creating a snapshot volume takes less time than creating a new one.

Before you begin

- The snapshot volume must be in either an Optimal or Disabled state.
- All member snapshot volumes must be in a Disabled state before you can re-create the snapshot consistency group snapshot volume.

About this task

You cannot re-create an individual member snapshot volume; you can re-create only the overall snapshot consistency group snapshot volume.



If the snapshot volume or snapshot consistency group snapshot volume is part of an online copy relationship, you cannot perform the re-create option on the volume.

Steps

1. Select **Storage › Snapshots**.
2. Select the **Snapshot Volumes** tab.

The Snapshot Volumes table appears and displays all the snapshot volumes associated with the storage array.

3. Select the snapshot volume that you want to re-create, and then select **Uncommon Tasks › Recreate**.

The **Recreate Snapshot Volume** dialog box appears.

4. Select one of the following options:

- **An existing snapshot image created from volume <name>**

Select this option to indicate an existing snapshot image from which to re-create the snapshot volume.

- **A new (instant) snapshot image of volume <name>**

Select this option to create a new snapshot image from which to re-create the snapshot volume.

5. Click **Recreate**.

Results

System Manager performs the following actions:

- Deletes all `write` data on any associated snapshot repository volume.
- Snapshot volume or snapshot consistency group snapshot volume parameters remain the same as the previously disabled volume parameters.
- Retains the original names for the snapshot volume or snapshot consistency group snapshot volume.

Disable snapshot volume

You can disable a snapshot volume or a snapshot volume in a snapshot consistency group when you no longer need it or want to temporarily stop using it.

About this task

Use the Disable option if one of these conditions applies:

- You are finished with the snapshot volume or snapshot consistency group snapshot volume for the time being.
- You intend to re-create the snapshot volume or snapshot consistency group snapshot volume (that is designated as read-write) at a later time and want to retain the associated reserved capacity so you do not need to create it again.
- You want to increase the storage array performance by stopping write activity to a read-write snapshot volume.

If the snapshot volume or snapshot consistency group snapshot volume is designated as read-write, this option also lets you stop any further write activity to its associated reserved capacity volume. If you decide to re-create the snapshot volume or snapshot consistency group snapshot volume, you must choose a snapshot image from the same base volume.



If the snapshot volume or snapshot consistency group snapshot volume is part of an online copy relationship, you cannot perform the Disable option on the volume.

Steps

1. Select **Storage > Snapshots**.
2. Select the **Snapshot Volumes** tab.

System Manager displays all the snapshot volumes associated with the storage array.

3. Select the snapshot volume that you want to disable, and then select **Uncommon Tasks > Disable**.
4. Confirm that you want to perform the operation, and then click **Disable**.

Results

- The snapshot volume remains associated with its base volume.
- The snapshot volume retains its World Wide Name (WWN).
- If read-write, the snapshot volume retains its associated reserved capacity.

- The snapshot volume retains any host assignments and access. However, read-write requests fail.
- The snapshot volume loses its association with its snapshot image.

Delete snapshot volume

You can delete a snapshot volume or a snapshot consistency group snapshot volume that is no longer needed for backup or software application testing purposes. You can also specify whether you want to delete the snapshot reserved capacity volume associated with a `read-write` snapshot volume or retain the snapshot reserved capacity volume as an unassigned volume.

About this task

Deleting a base volume automatically deletes any associated snapshot volume or consistency group snapshot volume. You cannot delete a snapshot volume that is in a volume copy with a status of **In Progress**.

Steps

1. Select **Storage > Snapshots**.
2. Select the **Snapshot Volumes** tab.

System Manager displays all the snapshot volumes associated with the storage array.

3. Select the snapshot volume that you want to delete, and then select **Uncommon Tasks > Delete**.
4. Confirm that you want to perform the operation, and then click **Delete**.

Results

System Manager performs the following actions:

- Deletes all member snapshot volumes (for a snapshot consistency group snapshot volume).
- Removes all associated host assignments.

FAQs

Why don't I see all my volumes, hosts, or host clusters?

Snapshot volumes with a DA-enabled base volume are not eligible to be assigned to a host that is not Data Assurance (DA) capable. You must disable DA on the base volume before a snapshot volume can be assigned to a host that is not DA capable.

Consider the following guidelines for the host to which you are assigning the snapshot volume:

- A host is not DA capable if it is connected to the storage array through an I/O interface that is not DA capable.
- A host cluster is not DA capable if it has at least one host member that is not DA capable.



You cannot disable DA on a volume that is associated with snapshots (consistency groups, snapshot groups, snapshot images, and snapshot volumes), volume copies, and mirrors. All associated reserved capacity and snapshot objects must be deleted before DA can be disabled on the base volume.

What is a snapshot image?

A snapshot image is a logical copy of volume content, captured at a particular point in time. Snapshot images use minimal storage space.

Snapshot image data is stored as follows:

- When a snapshot image is created, it exactly matches the base volume. After the snapshot is taken, when the first write request occurs for any block or set of blocks on the base volume, the original data is copied to the snapshot reserved capacity before the new data is written to the base volume.
- Subsequent snapshots include only data blocks that have changed since the first snapshot image was created. Each subsequent copy-on-write operation saves original data that is about to be overwritten on the base volume to the snapshot reserved capacity before the new data is written to the base volume.

Why use snapshot images?

You can use snapshots to protect against and allow recovery from accidental or malicious loss or corruption of data.

Select a base volume or a group of base volumes, called a snapshot consistency group, and then capture snapshot images in one or more of the following ways:

- You can create a snapshot image of a single base volume or a snapshot consistency group consisting of multiple base volumes.
- You can take snapshots manually or create a schedule for a base volume or snapshot consistency group to automatically capture periodic snapshot images.
- You can create a host-accessible snapshot volume of a snapshot image.
- You can perform a rollback operation to restore a snapshot image.

The system retains multiple snapshot images as restore points you can use to roll back to known good data sets at specific points in time. The ability to roll back provides protection against accidental deletion of data and data corruption.

What kinds of volumes can be used for snapshots?

Standard volumes and thin volumes are the only types of volumes that can be used to store snapshot images. Non-standard volumes cannot be used. The base volume can reside on either a pool or volume group.

Why would I create a snapshot consistency group?

You create a snapshot consistency group when you want to make sure snapshot images are taken on multiple volumes at the same time. For example, a database made up of multiple volumes that need to stay consistent for recovery purposes would require a snapshot consistency group to collect coordinated snapshots of all volumes and use them to restore the entire database.

The volumes included in a snapshot consistency group are called member volumes.

You can perform the following snapshot operations on a snapshot consistency group:

- Create a snapshot image of a snapshot consistency group to get simultaneous images of the member volumes.
- Create a schedule for the snapshot consistency group to automatically capture periodic simultaneous images of the member volumes.
- Create a host-accessible snapshot volume of a snapshot consistency group image.
- Perform a rollback operation for a snapshot consistency group.

What is a snapshot volume and when does it need reserved capacity?

A snapshot volume allows the host to access data in the snapshot image. The snapshot volume contains its own reserved capacity, which saves any modifications to the base volume without affecting the original snapshot image. Snapshot images are not read- or write-accessible to hosts. If you want to read or write to snapshot data, create a snapshot volume and assign it to a host.

You can create two types of snapshot volumes. The type of snapshot volume determines if it uses reserved capacity.

- **Read-only** — A snapshot volume that is created as read-only provides a host application with read access to a copy of the data contained in the snapshot image. A read-only snapshot volume does not use reserved capacity.
- **Read-write** — A snapshot volume that is created as read-write allows you to make changes to the snapshot volume without affecting the referenced snapshot image. A read-write snapshot volume uses reserved capacity to store these changes. You can convert a read-only snapshot volume to read-write at any time.

What is a snapshot group?

A snapshot group is a collection of point-in-time snapshot images of a single associated base volume.

System Manager organizes snapshot images into *snapshot groups*. Snapshot groups require no user action, but you can adjust reserved capacity on a snapshot group at any time. Additionally, you might be prompted to create reserved capacity when the following conditions are met:

- Any time you take a snapshot of a base volume that does not yet have a snapshot group, System Manager automatically creates a snapshot group. This creates reserved capacity for the base volume that is used to store subsequent snapshot images.
- Any time you create a snapshot schedule for a base volume, System Manager automatically creates a snapshot group.

Why would I disable a snapshot volume?

You disable a snapshot volume when you want to assign a different snapshot volume to the snapshot image. You can reserve the disabled snapshot volume for later use.

If you no longer need the snapshot volume or the consistency group snapshot volume and do not intend to re-create it at a later time, you should delete the volume instead of disabling it.

What is the Disabled state?

A snapshot volume in Disabled status is not currently assigned to a snapshot image. To enable the snapshot volume, you must use the re-create operation to assign a new snapshot image to the disabled snapshot volume.

The snapshot volume characteristics are defined by the snapshot image assigned to it. Read-write activity is suspended on a snapshot volume in Disabled status.

Why would I suspend a snapshot schedule?

When a schedule is suspended, the scheduled snapshot image creations do not occur. You can pause a snapshot schedule to conserve storage space, and then resume the scheduled snapshots at a later time.

If you do not need the snapshot schedule, you should delete the schedule instead of suspending it.

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.