



Pools and volume groups

SANtricity software

NetApp
November 03, 2025

Table of Contents

Pools and volume groups	1
Learn about pools and volume groups in SANtricity System Manager	1
What are pools and volume groups?	1
How do you create pools?	1
How do you create volume groups?	1
Related information	1
Concepts	2
How pools and volume groups work in SANtricity software	2
Learn about capacity terminology in SANtricity software	4
Learn about pools and volume groups in SANtricity software	6
Learn about automatic versus manual pool creation in SANtricity System Manager	9
Configure storage	9
Create a pool automatically in SANtricity System Manager	9
Create a pool manually in SANtricity System Manager	11
Create a volume group in SANtricity System Manager	13
Add capacity to a pool or volume group in SANtricity System Manager	18
Manage storage	21
Check volume redundancy in SANtricity System Manager	21
Delete a pool or volume group in SANtricity System Manager	23
Consolidate free capacity for a volume group in SANtricity System Manager	24
Export/Import volume groups in SANtricity CLI	25
Turn on locator lights in a pool, volume group, or SSD Cache in SANtricity System Manager	25
Remove capacity from a pool or SSD Cache in SANtricity System Manager	26
Modify pool and group settings	27
Change the configuration settings for a pool in SANtricity System Manager	27
Change the configuration settings for a volume group in SANtricity System Manager	29
Enable or disable resource provisioning on existing volume groups and pools in SANtricity System Manager	32
Enable or disable resource provisioning for new volume groups or pools in SANtricity System Manager	33
Enable security for a pool or volume group in SANtricity System Manager	33
Manage SSD cache	34
How SSD Cache works in SANtricity software	34
Learn about SSD cache restrictions in SANtricity software	36
Create an SSD Cache in SANtricity System Manager	36
Change the SSD Cache settings in SANtricity System Manager	38
View SSD Cache statistics in SANtricity System Manager	39
Manage reserved capacity	41
How reserved capacity works in SANtricity software	41
Increase reserved capacity in SANtricity System Manager	42
Decrease reserved capacity in SANtricity System Manager	43
Change the reserved capacity settings for a snapshot group in SANtricity System Manager	44
Change the reserved capacity settings for a snapshot volume in SANtricity System Manager	46

Change the reserved capacity settings for a consistency group member volume in SANtricity System Manager	46
Change the reserved capacity settings for a mirrored pair volume in SANtricity System Manager	47
Cancel a pending snapshot image in SANtricity System Manager	48
Delete a snapshot group in SANtricity System Manager	49
Pools and volume groups FAQ for SANtricity System Manager	49
What is a volume group?	49
What is a pool?	49
What is reserved capacity?	50
What is FDE/FIPS security?	50
What is redundancy check?	50
What are the differences between pools and volume groups?	50
Why would I want to manually configure a pool?	51
Why are capacity alerts important?	51
Why can't I increase my preservation capacity?	51
Is there a limit on the number of drives I can remove from a pool?	52
What media types are supported for a drive?	52
Why are some drives not showing up?	52
How do I maintain shelf/drawer loss protection?	53
What is the optimal drive positioning for pools and volume groups?	53
What RAID level is best for my application?	53
What is Data Assurance?	56
What is secure-capable (Drive Security)?	56
What do I need to know about increasing reserved capacity?	56
Why can't I choose another amount to decrease by?	57
Why do I need reserved capacity for each member volume?	57
How do I view and interpret all SSD Cache statistics?	57
What is optimization capacity for pools?	60
What is optimization capacity for volume groups?	60
What is resource provisioning capable?	61
What do I need to know about the resource-provisioned volumes feature?	61

Pools and volume groups

Learn about pools and volume groups in SANtricity System Manager

You can create logical storage capacity from a subset of unassigned drives in your storage array. This logical capacity can take the form of either a pool or a volume group, depending on the needs of your environment.

What are pools and volume groups?

A *pool* is a set of logically grouped drives. A *volume group* is a container for volumes with shared characteristics. You can use either a pool or volume group to create volumes accessible to a host.

Learn more:

- [How pools and volume groups work](#)
- [Capacity terminology](#)
- [Decide whether to use a pool or a volume group](#)

How do you create pools?

You can allow System Manager to create pools automatically when it detects unassigned capacity in a storage array. Alternatively, when automatic creation cannot determine the best configuration, you can create pools manually from **Storage > Pools & Volume Groups**.

Learn more:

- [Automatic versus manual pool creation](#)
- [Create pool automatically](#)
- [Create pool manually](#)
- [Add capacity to a pool or volume group](#)

How do you create volume groups?

You can create volume groups from **Storage > Pools & Volume Groups**.

Learn more:

- [Create a volume group](#)
- [Add capacity to a pool or volume group](#)

Related information

Learn more about concepts related to pools and volume groups:

- [How reserved capacity works](#)

- How SSD Cache works

Concepts


How pools and volume groups work in SANtricity software

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. System Manager automatically sets the RAID level 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.

Learn about capacity terminology in SANtricity software

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.
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.

Storage object	Description
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	Allocated capacity is the physical capacity allocated from the drives in a pool or volume group. You use allocated capacity to create volumes and for copy services operations.
Free capacity	Free capacity is the capacity available in a pool or volume group that has not yet been allocated to volume creation or copy services operations and storage objects.
Pool or volume group capacity	Pool, volume, or volume group capacity is the capacity in a storage array that has been assigned to a pool or volume group. This capacity is used to create volumes and service the various capacity needs of copy services operations and storage objects.
Pool unusable capacity	Pool unusable capacity is the space in a pool that cannot be used due to mismatched drive sizes.
Preservation capacity	Preservation capacity is the amount of capacity (number of drives) that is reserved in a pool to support potential drive failures.
Reported capacity	Reported capacity is the capacity that is reported to the host and can be accessed by the host.
Reserved capacity	Reserved capacity is the physical allocated capacity that is used for any copy service operation and storage object. It is not directly readable by the host.
SSD Cache	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.

Learn about pools and volume groups in SANtricity software

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 in System Manager.	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.

Learn about automatic versus manual pool creation in SANtricity System Manager

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.

Configure storage

Create a pool automatically in SANtricity System Manager

Pool creation is initiated automatically when SANtricity System Manager detects unassigned drives in the storage array. You can use automatic pool creation to easily

configure all unassigned drives in the storage array into one pool and to add drives into existing pools.

Before you begin

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

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

About this task

Keep in mind the following:

- When you add drives to a storage array, System Manager automatically detects the drives and prompts you to create a single pool or multiple pools based on the drive type and the current configuration.
- If pools were previously defined, System Manager automatically prompts you with the option of adding the compatible drives to an existing pool. When new drives are added to an existing pool, System Manager automatically redistributes the data across the new capacity, which now includes the new drives that you added.
- When configuring an EF600 or EF300 storage array, make sure each controller has access to an equal number of drives in the first 12 slots and an equal number of drives in the last 12 slots. This configuration helps the controllers use both drive-side PCIe buses more effectively.

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.

System Manager performs the following tasks:

- Creates a single pool if there are a sufficient number of drives with the same drive type (HDD or SSD) and have similar capacity.
- Creates multiple pools if the unassigned capacity consists of different drive types.
- Adds 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.
- Adds 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.

For EF600 and EF300 storage arrays, settings are also displayed for resource provisioning and volume block sizes.

5. Click **Accept**.

Create a pool manually in SANtricity System Manager

You can create a pool manually (from a set of candidates) if the Pool Auto Configuration feature does not provide a pool that meets your needs.

A pool provides the logical storage capacity necessary from which you can create individual volumes that can then be used to host your applications.

Before you begin

- You must have a minimum of 11 drives with the same drive type (HDD or SSD).
- Shelf loss protection requires that the drives comprising the pool are located in at least six different drive shelves and there are no more than two drives in a single drive shelf.
- Drawer loss protection requires that the drives comprising the pool are located in at least five different drawers and the pool includes an equal number of drive shelves from each drawer.
- When configuring an EF600 or EF300 storage array, make sure each controller has access to an equal number of drives in the first 12 slots and an equal number of drives in the last 12 slots. This configuration helps the controllers use both drive-side PCIe buses more effectively. Currently System Manager allows for drive selection under the Advanced feature when creating a volume group. For pool creation, it is recommended to use all drives in the storage array.

Steps

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

The Create Pool dialog box appears.

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

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

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

Characteristic	Use
Free Capacity	Shows the free capacity of the pool candidate in GiB. Select a pool candidate with the capacity for your application's storage needs. Preservation (spare) capacity is also distributed throughout the pool and is not part of the free capacity amount.

Characteristic	Use
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>
Drive Block Size (EF300 and EF600 only)	<p>Shows the block size (sector size) that the drives in the pool can write. Values may include:</p> <ul style="list-style-type: none"> • 512 — 512-byte sector size. • 4K — 4,096-byte sector size.
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 or Yes - FIPS (Mixed). "Mixed" indicates a mixture of 140-2 and 140-3 level drives. If you use a mixture of these levels, be aware that the pool will then operate at the lower level of security (140-2). • 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> <p> The only way to remove Drive Security after it is enabled is to delete the pool and erase the drives.</p>
DA Capable	<p>Indicates if Data Assurance (DA) is available for this pool candidate. DA checks for and corrects errors that might occur as data is transferred through the controllers down to the drives.</p> <p>DA is enabled if all drives are DA-capable. DA may be disabled after the volume is created by selecting Storage > Volumes > View/Edit Settings > Advanced > Permanently disable data assurance. If DA is disabled on a volume, it cannot be re-enabled.</p>

Characteristic	Use
Resource Provisioning Capable (EF300 and EF600 only)	Shows if Resource Provisioning is available for this pool candidate. Resource Provisioning is a feature available in the EF300 and EF600 storage arrays, which allows volumes to be put in use immediately with no background initialization process.
Shelf Loss Protection	<p>Shows if shelf loss protection is available.</p> <p>Shelf loss protection guarantees accessibility to the data on the volumes in a pool if a total loss of communication occurs with a single drive shelf.</p>
Drawer Loss Protection	<p>Shows if drawer loss protection is available, which is provided only if you are using a drive shelf that contains drawers.</p> <p>Drawer loss protection guarantees accessibility to the data on the volumes in a pool if a total loss of communication occurs with a single drawer in a drive shelf.</p>
Volume Block Sizes Supported (EF300 and EF600 only)	<p>Shows the block sizes that can be created for the volumes in the pool:</p> <ul style="list-style-type: none"> • 512n — 512 bytes native. • 512e — 512 bytes emulated. • 4K — 4,096 bytes.

Create a volume group in SANtricity System Manager

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.

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.

Before you begin

Review these guidelines before you create a volume group:

- You need at least one unassigned drive.
- Limits exist on the number of drives you can have in a single volume group. These limits vary according to the RAID level.
- 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.
- If you have an EF600 or EF300 storage array, and you plan to create a volume group manually, make sure each controller has access to an equal number of drives in the first 12 slots and an equal number of drives in the last 12 slots. This configuration helps the controllers use both drive-side PCIe buses more effectively. Currently System Manager allows for drive selection under the Advanced feature when creating a volume group.
- 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.
Drive Block Size (EF300 and EF600 only)	<p>Shows the block size (sector size) that the drives in the group can write. Values may include:</p> <ul style="list-style-type: none"> • 512 — 512-byte sector size. • 4K — 4,096-byte sector size.
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 or Yes - FIPS (Mixed). "Mixed" indicates a mixture of 140-2 and 140-3 level drives. If you use a mixture of these levels, be aware that the volume group will then operate at the lower level of security (140-2). • 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> <p> The only way to remove Drive Security after it is enabled is to delete the volume group and erase the drives.</p>
DA Capable	<p>Indicates if Data Assurance (DA) is available for this group. Data Assurance (DA) checks for and corrects errors that might occur as data is transferred through the controllers down to the drives.</p> <p>If you want to use DA, select a volume group that is DA capable. (For DA-capable drives, DA is automatically enabled on volumes created in the pool.)</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>

Characteristic	Use
Resource Provisioning Capable (EF300 and EF600 only)	Shows if Resource Provisioning is available for this group. Resource Provisioning is a feature available in the EF300 and EF600 storage arrays, which allows volumes to be put in use immediately with no background initialization process.
Shelf Loss Protection	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.
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 volume group if a total loss of communication occurs with a single drawer in a drive shelf.
Volume Block Sizes Supported (EF300 and EF600 only)	<p>Shows the block sizes that can be created for the volumes in the group:</p> <ul style="list-style-type: none"> • 512n — 512 bytes native. • 512e — 512 bytes emulated. • 4K — 4,096 bytes.

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) <p>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.</p>
Drive Block Size (EF300 and EF600 only)	<p>Shows the block size (sector size) that the drives in the group can write. Values may include:</p> <ul style="list-style-type: none"> • 512 — 512-byte sector size. • 4K — 4,096-byte sector size.
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 or Yes - FIPS (Mixed). "Mixed" indicates a mixture of 140-2 and 140-3 level drives. If you use a mixture of these levels, be aware that the volume group will then operate at the lower level of security (140-2). • You can create a volume group comprised of drives that might or might not be secure-capable or are a mix of security levels. If the drives in the volume group include drives that are not secure-capable, you cannot make the volume group secure.
DA Capable	<p>Indicates if Data Assurance (DA) is available for this group. Data Assurance (DA) checks for and corrects errors that might occur as data is communicated through the controllers down to the drives.</p> <p>If you want to use DA, select a volume group that is DA capable. (For DA-capable drives, DA is automatically enabled on volumes created in the pool.)</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>
Volume Block Sizes Supported (EF300 and EF600 only)	<p>Shows the block sizes that can be created for the volumes in the group:</p> <ul style="list-style-type: none"> • 512n — 512 bytes native. • 512e — 512 bytes emulated. • 4K — 4,096 bytes.
Resource Provisioning Capable (EF300 and EF600 only)	<p>Shows if Resource Provisioning is available for this group. Resource Provisioning is a feature available in the EF300 and EF600 storage arrays, which allows volumes to be put in use immediately with no background initialization process.</p>

Add capacity to a pool or volume group in SANtricity System Manager

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.
- The maximum number of volumes allowed in a volume group is 256.
- The maximum number of volumes allowed in a pool depends on the storage system model:
 - 2,048 volumes (EF600 and E5700 series)
 - 1,024 volumes (EF300)
 - 512 volumes (E4000 and E2800 series)
- 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. 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 system.)



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 if the drive is secure-capable.</p> <ul style="list-style-type: none"> To protect your pool or volume group with the Drive Security feature, all the drives must be secure-capable. It is possible to create a pool or volume group with a mix of secure-capable and non-secure-capable drives, but the Drive Security feature cannot be enabled. A pool or volume group with all secure-capable drives cannot accept a non-secure-capable drive for sparing or expansion, even if the encryption capability is not in use. Drives that are reported as secure-capable can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives. A FIPS drive can be level 140-2 or 140-3, with level 140-3 as the higher level of security. If you select a mixture of 140-2 and 140-3 level drives, the pool or volume group will then operate at the lower level of security (140-2).

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.
DULBE capable	Indicates whether the drive has the option for Deallocated or Unwritten Logical Block Error (DULBE). DULBE is an option on NVMe drives that allows the EF300 or EF600 storage array to support resource-provisioned volumes.

4. Click **Add**.

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

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



* Currently, the confirmation dialog box does not display when adding drives to a pool with shelf loss protection or drawer loss protection.

1. 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 storage

Check volume redundancy in SANtricity System Manager

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.



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 a pool or volume group in SANtricity System Manager

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

Before you begin

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

Steps

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

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

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

Results

System Manager performs the following actions:

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

Consolidate free capacity for a volume group in SANtricity System Manager

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**.

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 in SANtricity CLI

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.

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

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 in SANtricity System Manager

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.

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

Modify pool and group settings

Change the configuration settings for a pool in SANtricity System Manager

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

About this task

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



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

Steps

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

The Pool Setting dialog box appears.

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

Field details

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

Setting	Description
Preservation capacity ("Optimization capacity" for the EF600 or EF300)	<p>Preservation capacity — You can define the number of drives to determine the capacity that is reserved on the pool to support potential drive failures. When a drive failure occurs, the preservation capacity is used to hold the reconstructed data. Pools use preservation capacity during the data reconstruction process instead of hot spare drives, which are used in volume groups.</p> <p>Use the spinner controls to adjust the number of drives. Based on the number of drives, the preservation capacity in the pool appears next to the spinner box.</p> <p>Keep the following information in mind about preservation capacity.</p> <ul style="list-style-type: none"> Because preservation capacity is subtracted from the total free capacity of a pool, the amount of capacity that you reserve affects how much free capacity is available to create volumes. If you specify 0 for the preservation capacity, all of the free capacity on the pool is used for volume creation. If you decrease the preservation capacity, you increase the capacity that can be used for pool volumes. <p>Additional optimization capacity (EF600 and EF300 arrays only) — When a pool is created, a recommended optimization capacity is generated that provides a balance of available capacity versus performance and drive wear life. You can adjust this balance by moving the slider to the right for better performance and drive wear life at the expense of increased available capacity, or by moving it to the left for increased available capacity at the expense of better performance and drive wear life.</p> <p>SSD drives will have longer life and better maximum write performance when a portion of their capacity is unallocated. For drives associated with a pool, unallocated capacity is comprised of a pool's preservation capacity, the free capacity (capacity not used by volumes), and a portion of the usable capacity set aside as additional optimization capacity. The additional optimization capacity ensures a minimum level of optimization capacity by reducing the usable capacity, and as such, is not available for volume creation.</p>

4. Click **Save**.

Change the configuration settings for a volume group in SANtricity System Manager

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

Before you begin

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

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

Steps

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

The Volume Group Settings dialog box appears.

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

Field details

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

4. Click **Save**.

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

Results

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

Enable or disable resource provisioning on existing volume groups and pools in SANtricity System Manager

For any DULBE-capable drives, you can enable or disable resource provisioning on existing volumes in a pool or volume group.

Resource provisioning is a feature available in the EF300 and EF600 storage arrays, which allows volumes to be put in use immediately with no background initialization process. All drive blocks assigned to the volume are deallocated (unmapped), which can improve SSD wear life and increase maximum write performance.

By default, resource provisioning is enabled on systems where the drives support DULBE. There is no need to enable resource provisioning unless you have previously disabled it.

Before you begin

- You must have an EF300 or EF600 storage array.



DULBE is not supported on EF300C or EF600C storage arrays at this time.

- You must have SSD volume groups or pools, where all the drives support the NVMe Deallocated or Unwritten Logical Block Error Enable (DULBE) error recovery capability. Otherwise, the resource provisioning option is not available.

About this task

When you enable resource provisioning for existing volume groups and pools, all volumes in the selected volume group or pool are changed to allow the blocks to be deallocated. This process might involve a background operation to ensure consistent allocation at the unmap granularity. This operation does not unmap any space. Once the background operation completes, the operating system needs to unmap any unused blocks to create free space.

When you disable resource provisioning for existing volume groups or pools, a background operation rewrites all the logical blocks in every volume. Existing data remains intact. The writes will map or provision the blocks on the drives associated with the volume group or pool.



For new volume groups and pools, you can enable or disable resource provisioning from **Settings** > **System** > **Additional Settings** > **Enable/Disable Resource-Provisioned Volumes**.

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**, and then either **Enable resource provisioning** or **Disable resource provisioning**.
4. In the dialog box, confirm the operation.



If you re-enabled DULBE — After the background operation completes, you might need to reboot the host so it detects the DULBE configuration changes, and then remount all the filesystems.

Enable or disable resource provisioning for new volume groups or pools in SANtricity System Manager

If you previously disabled the default feature for resource provisioning, you can re-enable it for any new SSD volume groups or pools that you create. You can also disable the setting again.

Resource provisioning is a feature available in the EF300 and EF600 storage arrays, which allows volumes to be put in use immediately with no background initialization process. All drive blocks assigned to the volume are deallocated (unmapped), which can improve SSD wear life and increase maximum write performance.



By default, resource provisioning is enabled on systems where the drives support DULBE.

Before you begin

- You must have an EF300 or EF600 storage array.
- You must have SSD volume groups or pools, where all the drives support the NVMe Deallocated or Unwritten Logical Block Error Enable (DULBE) error recovery capability.



DULBE is not supported on EF300C or EF600C storage arrays at this time.

About this task

When you re-enable resource provisioning for new volume groups or pools, only newly created volume groups and pools are affected. Any existing volume groups and pools with resource provisioning enabled will remain unchanged.

Steps

1. Select **Settings > System**.
2. Scroll down to **Additional Settings**, and then click **Enable/Disable Resource-Provisioned Volumes**.

The setting description indicates whether resource provisioning is currently enabled or disabled.

3. In the dialog box, confirm the operation.

Results

Enabling or disabling resource provisioning affects only new SSD pools or volume groups that you create. Existing pools or volume groups remain unchanged.

Enable security for a pool or volume group in SANtricity System Manager

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**.

Manage SSD cache

How SSD Cache works in SANtricity software

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 if 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 it is beneficial to place the data in cache to improve overall system performance.

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

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

Improved performance

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

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

SSD Cache and the Drive Security feature

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

Implement SSD Cache

To implement SSD Cache, do the following:

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



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

Learn about SSD cache restrictions in SANtricity software

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

Restrictions

- Any volume assigned to use a controller's SSD Cache is not eligible for an automatic load balance transfer.
- Currently, only one SSD Cache is supported per storage array.
- The maximum usable SSD Cache capacity on a storage array is 10 TB.
- 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.

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.

Create an SSD Cache in SANtricity System Manager

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

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**.

Change the SSD Cache settings in SANtricity System Manager

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

Steps

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

The SSD Cache Settings dialog box appears.

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

Field details

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

4. Click **Save**.

View SSD Cache statistics in SANtricity System Manager

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

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.

Manage reserved capacity

How reserved capacity works in SANtricity software

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.

Increase reserved capacity in SANtricity System Manager

You can increase reserved capacity, which is the physically allocated capacity used for any copy service operation on a storage object.

For snapshot operations, it is typically 40 percent of the base volume; for asynchronous mirroring operations, it is typically 20 percent of the base volume. Typically, you increase reserved capacity when you receive a warning that the storage object's reserved capacity is becoming full.

Before you begin

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

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

About this task

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

- Snapshot group
- Snapshot volume

- Consistency group member volume
- Mirrored pair volume

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



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

Steps

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

The Increase Reserved Capacity dialog box appears.

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

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

5. Click **Increase**.

Results

System Manager performs the following actions:

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

Decrease reserved capacity in SANtricity System Manager

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 in SANtricity System Manager

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

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

Steps

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

The Snapshot Group Settings dialog box appears.

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

Field details

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

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

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

Change the reserved capacity settings for a snapshot volume in SANtricity System Manager

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 in SANtricity System Manager

You can change the settings for a consistency group member volume to adjust the percentage point at which SANtricity 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> <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>
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 in SANtricity System Manager

You can change the settings for a mirrored pair volume to adjust the percentage point at

which SANtricity 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> <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>

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

Cancel a pending snapshot image in SANtricity System Manager

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 a snapshot group in SANtricity System Manager

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.

Pools and volume groups FAQ for SANtricity System Manager

This FAQ can help if you're just looking for a quick answer to a question.

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 SANtricity 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, the system 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 from **Pools & Volume Groups**. 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?

SANtricity 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 the following:
 - Hard Disk Drives (HDDs) with Solid State Disks (SSDs)
 - NVMe with SAS drives
 - Drives with 512-byte and 4KiB volume block sizes
- If a pool or volume group contains all secure-capable drives, non-secure-capable drives are not listed.
- If a pool or volume group contains all Federal Information Processing Standards (FIPS) drives, non-FIPS drives are not listed.
- If a pool or volume group contains all Data Assurance (DA)-capable drives and there is at least one DA-enabled volume in the pool or volume group, a drive that is not DA capable is not eligible, so it cannot be added to that pool or volume group. However, if there is no DA-enabled volume in the pool or volume group, a drive that is not DA capable can be added to that pool or volume group. If you decide to mix these drives, keep in mind that you cannot create any DA-enabled volumes.



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

How do I maintain shelf/drawer loss protection?

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

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

When creating pools and volume groups, make sure to balance the drive selection between the upper and lower drive slots.

For the EF600 and EF300 controllers, drive slots 0-11 are connected to one PCI bridge, while slots 12-23 are connected to a different PCI bridge. For optimal performance, you should balance the drive selection to include a roughly equal number of drives from the upper and lower slots. This positioning ensures that your volumes do not hit a bandwidth limit sooner than necessary.

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 will 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, go to **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. The user interface automatically configures pools as RAID 6.

RAID levels and data protection

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

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

What is Data Assurance?

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

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

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

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

What is secure-capable (Drive Security)?

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

What do I need to know about increasing reserved capacity?

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

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

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

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

For snapshot operations, reserved capacity is typically 40 percent of the base volume. For asynchronous

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

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

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

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

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

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

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

Why do I need reserved capacity for each member volume?

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

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

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

How do I view and interpret all SSD Cache statistics?

You can view nominal statistics and detailed statistics for SSD Cache. Nominal statistics are a subset of the detailed statistics.

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

Nominal statistics

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

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.
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 (%)	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.
Cache Utilization (%)	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.

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.
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.

Detailed statistics	Description
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.

What is optimization capacity for pools?

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

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

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



The Additional Optimization Capacity slider is only available for EF600 and EF300 storage systems.

What is optimization capacity for volume groups?

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

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

When a volume group is created, a recommended optimization capacity is generated that provides a balance of performance, drive wear life, and available capacity. The Additional Optimization Capacity slider in the

Volume Group Settings dialog allows adjustments to a volume group's optimization capacity. Adjusting the slider provides for better performance and drive wear life at the expense of available capacity, or additional available capacity at the expense of performance and drive wear life.



The Additional Optimization Capacity slider is only available for EF600 and EF300 storage systems.

What is resource provisioning capable?

Resource Provisioning is a feature available in the EF300 and EF600 storage arrays, which allows volumes to be put in use immediately with no background initialization process.

A resource-provisioned volume is a thick volume in an SSD volume group or pool, where drive capacity is allocated (assigned to the volume) when the volume is created, but the drive blocks are deallocated (unmapped). By comparison, in a traditional thick volume, all drive blocks are mapped or allocated during a background volume initialization operation in order to initialize the Data Assurance protection information fields and to make data and RAID parity consistent in each RAID stripe. With a resource provisioned volume, there is no time-bound background initialization. Instead, each RAID stripe is initialized upon the first write to a volume block in the stripe.

Resource-provisioned volumes are supported only on SSD volume groups and pools, where all drives in the group or pool support the NVMe Deallocated or Unwritten Logical Block Error Enable (DULBE) error recovery capability. When a resource-provisioned volume is created, all drive blocks assigned to the volume are deallocated (unmapped). In addition, hosts can deallocate logical blocks in the volume using the NVMe Dataset Management command or the SCSI Unmap command. Deallocating blocks can improve SSD wear life and increase maximum write performance. The improvement varies with each drive model and capacity.



DULBE is not supported on EF300C or EF600C storage arrays at this time.

What do I need to know about the resource-provisioned volumes feature?

Resource Provisioning is a feature available in the EF300 and EF600 storage arrays, which allows volumes to be put in use immediately with no background initialization process.

A resource-provisioned volume is a thick volume in an SSD volume group or pool, where drive capacity is allocated (assigned to the volume) when the volume is created, but the drive blocks are deallocated (unmapped). By comparison, in a traditional thick volume, all drive blocks are mapped or allocated during a background volume initialization operation in order to initialize the Data Assurance protection information fields and to make data and RAID parity consistent in each RAID stripe. With a resource provisioned volume, there is no time-bound background initialization. Instead, each RAID stripe is initialized upon the first write to a volume block in the stripe.

Resource-provisioned volumes are supported only on SSD volume groups and pools, where all drives in the group or pool support the NVMe Deallocated or Unwritten Logical Block Error Enable (DULBE) error recovery capability. When a resource-provisioned volume is created, all drive blocks assigned to the volume are deallocated (unmapped). In addition, hosts can deallocate logical blocks in the volume using the NVMe Dataset Management command or the SCSI Unmap command. Deallocating blocks can improve SSD wear life and increase maximum write performance. The improvement varies with each drive model and capacity.

Resource provisioning is enabled by default on systems where the drives support DULBE. You can disable that default setting from **Pools & Volume Groups**.



DULBE is not supported on EF300C or EF600C storage arrays at this time.

Copyright information

Copyright © 2025 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—with 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.