



Volumes and workloads

SANtricity 11.7

NetApp
February 12, 2024

Table of Contents

- Volumes and workloads 1
 - Volumes and workloads overview 1
 - Concepts 1
 - Configure storage 17
 - Manage volumes 28
 - Manage settings 34
 - Use copy services 43
 - FAQs 49

Volumes and workloads

Volumes and workloads overview

You can create a volume as a container in which applications, databases, and file systems store data. When creating a volume, you also select a workload to customize the storage array configuration for a specific application.

What are volumes and workloads?

A *volume* is the logical component created with specific capacity for the host to access. Although a volume might consist of more than one drive, a volume appears as one logical component to the host. Once a volume is defined, you can add it to a workload. A *workload* is a storage object that supports an application, such as SQL Server or Exchange, which you can use to optimize storage for that application.

Learn more:

- [How volumes work](#)
- [How workloads work](#)
- [Volume terminology](#)
- [How capacity is allocated for volumes](#)
- [Actions you can perform on volumes](#)

How do you create volumes and workloads?

First, you create a workload. Go to **Storage > Volumes** and open a wizard that guides you through the steps. Next, you create a volume from the capacity available in a pool or a volume group, and then assign the workload you created.

Learn more:

- [Workflow for creating volumes](#)
- [Create workloads](#)
- [Create volumes](#)
- [Add volumes to workload](#)

Related information

Learn more about concepts related to volumes:

- [Data integrity and data security for volumes](#)
- [SSD Cache and volumes](#)
- [Thin volume monitoring](#)

Concepts

How volumes work

Volumes are data containers that manage and organize the storage space on your storage array.

You create volumes from the storage capacity available on your storage array and make it easy to organize and use your system's resources. This concept is similar to using folders/directories on a computer to organize files for easy and quick access.

Volumes are the only data layer visible to hosts. In a SAN environment, volumes are mapped to logical unit numbers (LUNs), which are visible to hosts. LUNs hold the user data that is accessible using one or more of the host access protocols supported by the storage array, including FC, iSCSI, and SAS.

Volume types you can create from pools and volume groups

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

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



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

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

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

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

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



The EF600 and EF300 storage systems do not support thin provisioning.

Characteristics of volumes

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

- **Segment size** — A segment is the amount of data in kilobytes (KiB) that is stored on a drive before the storage array moves to the next drive in the stripe (RAID group). The segment size is equal to or less than the capacity of the volume group. The segment size is fixed and cannot be changed for pools.
- **Capacity** — You create a volume from the free capacity available in either a pool or volume group. Before you create a volume, the pool or volume group must already exist, and it must have enough free capacity to create the volume.
- **Controller ownership** — All storage arrays can have either one or two controllers. On a single-controller array, a volume's workload is managed by a single controller. On a dual-controller array, a volume will have a preferred controller (A or B) that "owns" the volume. In a dual-controller configuration, volume ownership

is automatically adjusted using the Automatic Load Balancing feature to correct any load balance issues when workloads shift across the controllers. Automatic load balancing provides automated I/O workload balancing and ensures that incoming I/O traffic from the hosts is dynamically managed and balanced across both controllers.

- **Volume assignment** — You can give hosts access to a volume either when you create the volume or at a later time. All host access is managed through a logical unit number (LUN). Hosts detect LUNs that are, in turn, assigned to volumes. If you are assigning a volume to multiple hosts, use clustering software to make sure that the volume is available to all of the hosts.

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

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

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

How workloads work

When creating a volume, you select a workload to customize the storage array configuration for a specific application.

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

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

Types of workloads

You can create two types of workloads: Application-specific and Other.

- **Application-specific.** When you are creating volumes using an application-specific workload, the system may recommend an optimized volume configuration to minimize contention between application workload I/O and other traffic from your application instance. Volume characteristics like I/O type, segment size, controller ownership, and read and write cache are automatically recommended and optimized for workloads that are created for the following application types.
 - Microsoft® SQL Server™
 - Microsoft® Exchange Server™
 - Video Surveillance applications
 - VMware ESXi™ (for volumes to be used with Virtual Machine File System)

You can review the recommended volume configuration and edit, add, or delete the system-

recommended volumes and characteristics using the Add/Edit Volumes dialog box.

- **Other** (or applications without specific volume creation support). Other workloads use a volume configuration that you must manually specify when you want to create a workload that is not associated with a specific application, or if the system does not have built-in optimization for the application you intend to use on the storage array. You must manually specify the volume configuration using the Add/Edit Volumes dialog box.

Application and workload views

To view applications and workloads, launch SANtricity System Manager. From that interface, you can view information associated with an application-specific workload in a couple of different ways:

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

Volume terminology

Learn how the volume terms apply to your storage array.

All volume types

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

Term	Description
Dynamic cache read prefetch	<p>Dynamic cache read prefetch allows the controller to copy additional sequential data blocks into the cache while it is reading data blocks from a drive to the cache. This caching increases the chance that future requests for data can be filled from the cache. Dynamic cache read prefetch is important for multimedia applications that use sequential I/O. The rate and amount of data that is prefetched into cache is self-adjusting based on the rate and request size of the host reads. Random access does not cause data to be prefetched into cache. This feature does not apply when read caching is disabled.</p> <p>For a thin volume, dynamic cache read prefetch is always disabled and cannot be changed.</p>
Free capacity area	<p>A free capacity area is the free capacity that can result from deleting a volume or from not using all available free capacity during volume creation. When you create a volume in a volume group that has one or more free capacity areas, the volume's capacity is limited to the largest free capacity area in that volume group. For example, if a volume group has a total of 15 GiB free capacity, and the largest free capacity area is 10 GiB, the largest volume you can create is 10 GiB.</p> <p>By consolidating free capacity, you can create additional volumes from the maximum amount of free capacity in a volume group.</p>
Host	A host is a server that sends I/O to a volume on a storage array.
Host cluster	A host cluster is a group of hosts. You create a host cluster to make it easy to assign the same volumes to multiple hosts.
Hot spare drive	Hot spare drives are supported only with volume groups. A hot spare drive contains no data and acts as a standby in case a drive fails in RAID 1, RAID 3, RAID 5, or RAID 6 volumes contained in a volume group. The hot spare drive adds another level of redundancy to your storage array.
LUN	<p>A logical unit number (LUN) is the number assigned to the address space that a host uses to access a volume. The volume is presented to the host as capacity in the form of a LUN.</p> <p>Each host has its own LUN address space. Therefore, the same LUN can be used by different hosts to access different volumes.</p>
Media scan	A media scan provides a way of detecting drive media errors before they are found during a normal read from or write to the drives. A media scan is performed as a background operation and scans all data and redundancy information in defined user volumes.
Namespace	A namespace is NVM storage that is formatted for block access. It is analogous to a logical unit in SCSI, which relates to a volume in the storage array.

Term	Description
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.)
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.
Read cache	The read cache is a buffer that stores data that has been read from the drives. The data for a read operation might already be in the cache from a previous operation, which eliminates the need to access the drives. The data stays in the read cache until it is flushed.
Reported capacity	Reported capacity is the capacity that is reported to the host and can be accessed by the host. Reported capacity and allocated capacity are the same for thick volumes, but are different for thin volumes. For a thick volume, the physically allocated space is equal to the space that is reported to the host. For a thin volume, reported capacity is the capacity that is reported to the hosts, whereas allocated capacity is the amount of drive space that is currently allocated for writing data.
Segment size	A segment is the amount of data in kilobytes (KiB) that is stored on a drive before the storage array moves to the next drive in the stripe (RAID group). The segment size is equal to or less than the capacity of the volume group. The segment size is fixed and cannot be changed for pools.
Striping	Striping is way of storing data on the storage array. Striping splits the flow of data into blocks of a certain size (called "block size") and then writes these blocks across the drives one by one. This way of data storage is used to distribute and store data across multiple physical drives. Striping is synonymous with RAID 0 and spreads the data across all the drives in a RAID group without parity.
Volume	A volume is a container in which applications, databases, and file systems store data. It is the logical component created for the host to access storage on the storage array.
Volume assignment	Volume assignment is how host LUNs are assigned to a volume.
Volume name	A volume name is a string of characters assigned to the volume when it is created. You can either accept the default name or provide a more descriptive name indicating the type of data stored in the volume.
Volume group	A volume group is a container for volumes with shared characteristics. A volume group has a defined capacity and RAID level. You can use a volume group to create one or more volumes accessible to a host. (You create volumes from either a volume group or a pool.)

Term	Description
Workload	A workload is a storage object that supports an application. You can define one or more workloads, or instances, per application. For some applications, the system configures the workload to contain volumes with similar underlying volume characteristics. These volume characteristics are optimized based on the type of application the workload supports. For example, if you create a workload that supports a Microsoft SQL Server application and then subsequently create volumes for that workload, the underlying volume characteristics are optimized to support Microsoft SQL Server.
Write cache	The write cache is a buffer that stores data from the host that has not yet been written to the drives. The data stays in the write cache until it is written to the drives. Write caching can increase I/O performance.
Write caching with mirroring	Write caching with mirroring occurs when the data written to the cache memory of one controller is also written to the cache memory of the other controller. Therefore, if one controller fails, the other can complete all outstanding write operations. Write cache mirroring is available only if write caching is enabled and two controllers are present. Write caching with mirroring is the default setting at volume creation.
Write caching without batteries	The write caching without batteries setting lets write caching continue even when the batteries are missing, failed, discharged completely, or not fully charged. Choosing write caching without batteries is not typically recommended, because data might be lost if power is lost. Typically, write caching is turned off temporarily by the controller until the batteries are charged or a failed battery is replaced.

Specific to thin volumes



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




Thin volumes are not available on the EF600 or EF300 storage system.

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

Workflow for creating volumes

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



Data integrity and data security for volumes

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

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).

Drive Security

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

How Drive Security works at the drive level

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

How Drive Security works at the volume level

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

How to implement Drive Security

To implement Drive Security, you perform the following steps.

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

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

SSD Cache and volumes

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



This feature is not available on the EF600 or EF300 storage system.

Volumes

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

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

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

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



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

Drive Security feature

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

Actions you can perform on volumes

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

Increase capacity

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

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

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

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

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

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

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

Delete

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

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

Copy

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

Initialize

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

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

Redistribute

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

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

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

Change volume ownership

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

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

Mirroring and volume ownership

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

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

Copy volume and volume ownership

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

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

Change cache settings

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

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

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

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

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

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

caching without batteries is not typically recommended, because data might be lost if power is lost. Typically, write caching is turned off temporarily by the controller until the batteries are charged or a failed battery is replaced.

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

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

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

Change media scan settings

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

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

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

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

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

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

How capacity is allocated for volumes

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

Using capacity to create and expand volumes

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

- When you create a volume from unassigned capacity, you can create a pool or volume group and the volume at the same time.
- When you create a volume from free capacity, you are creating an additional volume on an already existing pool or volume group.

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

Capacity types for thick volumes and thin volumes

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

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

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

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

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



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

Capacity limits for thick volumes

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

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

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

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

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

Capacity limits for thin volumes

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

allowing for better storage utilization.

In addition to reported capacity and allocated capacity, thin volumes also contain Written capacity. Written capacity is the amount of capacity that has been written from the reserved capacity allocated for thin volumes.

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

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

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

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

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

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

Thin volume monitoring

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

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

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

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

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



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



Thin volumes are not available on the EF600 or EF300 storage system.

Comparison between thick volumes and thin volumes

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



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

Volume type	Description
Thick volumes	<ul style="list-style-type: none">• Thick volumes are created from either a pool or volume group.• With thick volumes, a large amount of storage space is provided in advance in anticipation of future storage needs.• Thick volumes are created with the entire size of the volume pre-allocated on physical storage at the time the volume is created. This pre-allocation means that creating a 100 GiB volume actually consumes 100 GiB of allocated capacity on your drives. However, the space might remain unused, causing under-utilization of storage capacity.• When creating thick volumes, make sure not to over-allocate capacity for a single volume. Over-allocating capacity for a single volume can quickly consume all the physical storage in your system.• Keep in mind that storage capacity is also required for copy services (snapshot images, snapshot volumes, volume copies, and asynchronous mirroring), so do not allocate all of the capacity to thick volumes. Insufficient space can block writes to the pool or volume group. You receive a free-capacity alert threshold warning if this situation occurs.
Thin volumes	<ul style="list-style-type: none">• Thin volumes are created only from a pool, not from a volume group.• Thin volumes must be RAID 6.• Thin volumes are not available on the EF600 or EF300 storage system.• You must use the CLI to create thin volumes.• Unlike thick volumes, space required for the thin volume is not allocated during creation, but is supplied, on demand at a later time.• A thin volume lets you over-allocate its size. That is, you can assign a LUN size that is larger than the size of the volume. You can then expand the volume as needed (if necessary, adding drives in the process) without expanding the size of the LUN, and therefore without disconnecting users.• You can use thin provisioning block space reclamation (UNMAP) to reclaim blocks of a thin-provisioned volume on the storage array through a host-issued SCSI UNMAP command. A storage array that supports thin provisioning can re-purpose the reclaimed space to satisfy allocation requests for some other thin provisioned volume within the same storage array, which allows better reporting of disk space consumption and more efficient use of resources.

Thin volume restrictions

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

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

Configure storage

Create workloads

You can create workloads for any type of application.

About this task

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

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

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

Keep these guidelines in mind:

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

Steps

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

The Create Application Workload dialog box appears.

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

4. Click **Create**.

After you finish

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

Create volumes

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

About this task

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

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

Step 1: Select host for a volume

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

Before you begin

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

About this task

Keep these guidelines in mind when you assign volumes:

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



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

Steps

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

The Create Volumes dialog box appears.

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

The Select Workload dialog box appears.

Step 2: Select a workload for a volume

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

About this task

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

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

Steps

1. Do one of the following:
 - Select the **Create volumes for an existing workload** option to create volumes for an existing workload.
 - Select the **Create a new workload** option to define a new workload for a supported application or for "Other" applications.
 - From the drop-down list, select the name of the application you want to create the new workload for.

Select one of the "Other" entries if the application you intend to use on this storage array is not listed.
 - Enter a name for the workload you want to create.
2. Click **Next**.
3. If your workload is associated with a supported application type, enter the information requested; otherwise, go to [Step 3: Add or edit volumes](#).

Step 3: Add or edit volumes

System Manager may suggest a volume configuration based on the application or workload you selected. This volume configuration is optimized based on the type of application the workload supports. You can accept the

recommended volume configuration or you can edit it as needed. If you selected one of the "Other" applications, you must manually specify the volumes and characteristics you want to create.

Before you begin

- The pools or volume groups must have sufficient free capacity.
- 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 (E2800 series)
- To create a Data Assurance (DA)-enabled volume, the host connection you are planning to use must support DA.

Selecting a secure-capable pool or volume group

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

DA capabilities are presented at the pool and volume group level in System Manager. DA protection checks for and corrects errors that might occur as data is transferred through the controllers down to the drives. Selecting a DA-capable pool or volume group for the new volume ensures that any errors are detected and corrected.

If any of the host connections on the controllers in your storage array do not support DA, the associated hosts cannot access data on DA-enabled volumes.

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

Selecting a secure-capable pool or volume group

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

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

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

- To create a resource-provisioned volume, all drives must be NVMe drives with the Deallocated or Unwritten Logical Block Error (DULBE) option.

About this task

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

For some application-specific workloads, each eligible pool or volume group shows the proposed capacity

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

Steps

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

Field details

Field	Description
Volume Name	A volume is assigned a default name by System Manager during the volume creation sequence. You can either accept the default name or provide a more descriptive one indicating the type of data stored in the volume.
Reported Capacity	<p>Define the capacity of the new volume and the capacity units to use (MiB, GiB, or TiB). For Thick volumes, the minimum capacity is 1 MiB, and the maximum capacity is determined by the number and capacity of the drives in the pool or volume group.</p> <p>Keep in mind that storage capacity is also required for copy services (snapshot images, snapshot volumes, volume copies, and remote mirrors); therefore, do not allocate all of the capacity to standard volumes.</p> <p>Capacity in a pool is allocated in 4-GiB or 8-GiB increments, depending on your drive type. Any capacity that is not a multiple of 4- or 8-GiB is allocated but not usable. To make sure that the entire capacity is usable, specify the capacity in 4-GiB or 8-GiB increments. If unusable capacity exists, the only way to regain it is to increase the capacity of the volume.</p>
Volume Block Size (EF300 and EF600 only)	<p>Shows the block sizes that can be created for the volume:</p> <ul style="list-style-type: none">• 512 — 512 bytes• 4K — 4,096 bytes

Field	Description
Segment Size	<p>Shows the setting for segment sizing, which only appears for volumes in a volume group. You can change the segment size to optimize performance.</p> <p>Allowed segment size transitions — System Manager determines the segment size transitions that are allowed. Segment sizes that are inappropriate transitions from the current segment size are unavailable on the drop-down list. Allowed transitions usually are double or half of the current segment size. For example, if the current volume segment size is 32 KiB, a new volume segment size of either 16 KiB or 64 KiB is allowed.</p> <p>SSD Cache-enabled volumes — You can specify a 4-KiB segment size for SSD Cache-enabled volumes. Make sure you select the 4-KiB segment size only for SSD Cache-enabled volumes that handle small-block I/O operations (for example, 16 KiB I/O block sizes or smaller). Performance might be impacted if you select 4 KiB as the segment size for SSD Cache-enabled volumes that handle large block sequential operations.</p> <p>Amount of time to change segment size — The amount of time to change a volume's segment size depends on these variables:</p> <ul style="list-style-type: none"> • The I/O load from the host • The modification priority of the volume • The number of drives in the volume group • The number of drive channels • The processing power of the storage array controllers <p>When you change the segment size for a volume, I/O performance is affected, but your data remains available.</p>
Secure-capable	<p>Yes appears next to "Secure-capable" only if the drives in the pool or volume group are secure-capable.</p> <p>Drive Security prevents unauthorized access to the data on a drive that is physically removed from the storage array. This option is available only when the Drive Security feature has been enabled, and a security key is set up for the storage array.</p> <p>A pool or volume group can contain both secure-capable and non-secure-capable drives, but all drives must be secure-capable to use their encryption capabilities.</p>

Field	Description
DA	<p>Yes appears next to "DA" only if the drives in the pool or volume group support Data Assurance (DA).</p> <p>DA increases data integrity across the entire storage system. DA enables the storage array to check for errors that might occur as data is transferred through the controllers down to the drives. Using DA for the new volume ensures that any errors are detected.</p>
Resource provisioned (EF300 and EF600 only)	<p>Yes appears next to "Resource provisioned" only if the drives support this option. 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>

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

Field details

Field	Description
Volume Name	<p>A volume is assigned a default name by System Manager during the volume creation sequence. You can either accept the default name or provide a more descriptive one indicating the type of data stored in the volume.</p>
Reported Capacity	<p>Define the capacity of the new volume and the capacity units to use (MiB, GiB, or TiB). For Thick volumes, the minimum capacity is 1 MiB, and the maximum capacity is determined by the number and capacity of the drives in the pool or volume group.</p> <p>Keep in mind that storage capacity is also required for copy services (snapshot images, snapshot volumes, volume copies, and remote mirrors); therefore, do not allocate all of the capacity to standard volumes.</p> <p>Capacity in a pool is allocated in 4-GiB or 8-GiB increments, depending on your drive type. Any capacity that is not a multiple of 4- or 8-GiB is allocated but not usable. To make sure that the entire capacity is usable, specify the capacity in 4-GiB or 8-GiB increments. If unusable capacity exists, the only way to regain it is to increase the capacity of the volume.</p>
Volume Type	<p>Volume type indicates the type of volume that was created for an application-specific workload.</p>
Volume Block Size (EF300 and EF600 only)	<p>Shows the block sizes that can be created for the volume:</p> <ul style="list-style-type: none">• 512 — 512 bytes• 4K — 4,096 bytes

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

Field	Description
Secure-capable	<p>Yes appears next to "Secure-capable" only if the drives in the pool or volume group are secure-capable.</p> <p>Drive security prevents unauthorized access to the data on a drive that is physically removed from the storage array. This option is available only when the drive security feature has been enabled, and a security key is set up for the storage array.</p> <p>A pool or volume group can contain both secure-capable and non-secure-capable drives, but all drives must be secure-capable to use their encryption capabilities.</p>
DA	<p>Yes appears next to "DA" only if the drives in the pool or volume group support Data Assurance (DA).</p> <p>DA increases data integrity across the entire storage system. DA enables the storage array to check for errors that might occur as data is transferred through the controllers down to the drives. Using DA for the new volume ensures that any errors are detected.</p>
Resource provisioned (EF300 and EF600 only)	<p>Yes appears next to "Resource Provisioned" only if the drives support this option. 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>

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

Step 4: Review volume configuration

Review a summary of the volumes you intend to create and make any necessary changes.

Steps

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

Results

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

After you finish

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

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

Add volumes to workload

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

About this task

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

Steps

1. Select **Storage > Volumes**.

2. Select the **Applications & Workloads** tab.

The Applications & Workloads view appears.

3. Select **Add to Workload**.

The Select Workload dialog box appears.

4. Do one of the following actions:

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

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

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

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

The Select Volumes dialog box appears.

6. Select the volumes you want to add to the workload.

7. Review the volumes that you want to add to the selected workload.

8. When you are satisfied with your workload configuration, click **Finish**.

Manage volumes

Increase capacity of a volume

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

Before you begin

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

About this task

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



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

Steps

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

The Confirm Increase Capacity dialog box appears.

3. Select **Yes** to continue.

The Increase Reported Capacity dialog box appears.

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

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

Results

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

After you finish

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

Initialize volumes

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

Before you begin

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



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

About this task

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

The following types of volumes *cannot* be initialized:

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

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

Steps

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

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

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

Results

System Manager performs the following actions:

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

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

Redistribute volumes

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

Before you begin

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

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

About this task

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

Steps

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

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

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

Results

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

Change controller ownership of a volume

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

Before you begin

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

About this task

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

Steps

1. Select **Storage > Volumes**.

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

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

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

Results

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

Delete volume

Typically, you delete volumes if the volumes were created with the wrong parameters or capacity, no longer meet storage configuration needs, or are snapshot images that are no longer needed for backup or application testing.

Deleting a volume increases the free capacity in the pool or volume group. You can select one or more volumes to delete.

Before you begin

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

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

About this task

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

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



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



When a volume exceeds a given size (currently 128 TB) the delete is being performed in background and the freed space may not be immediately available.

Steps

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

The Delete Volumes dialog box appears.

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

Results

System Manager performs the following actions:

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

Change allocated capacity limit for a thin volume

For thin volumes capable of allocating space on demand, you can change the limit that restricts the allocated capacity to which a thin volume can automatically expand.

You also can change the percentage point at which an alert (warning threshold exceeded) is sent to the Notifications area on the Home page when a thin volume is near the allocated capacity limit. You can choose to enable or disable this alert notification.



This feature is not available on the EF600 or EF300 storage system.

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

When changing a thin volume's allocated capacity limit and warning threshold, you must take into account the space to be consumed by both the volume's user data and copy services data.

Steps

1. Select **Storage › Volumes**.
2. Select the **Thin Volume Monitoring** tab.

The Thin Volume Monitoring view appears.

3. Select the thin volume that you want to change, and then select **Change Limit**.

The Change Limit dialog box appears. The allocated capacity limit and warning threshold setting for the thin volume you selected appear in this dialog box.

4. Change the allocated capacity limit and warning threshold as needed.

Field details

Setting	Description
Change allocated capacity limit to...	The threshold at which writes fail, preventing the thin volume from consuming additional resources. This threshold is a percentage of the volume's reported capacity size.
Alert me when... (warning threshold)	Select the check box if you want the system to generate an alert when a thin volume is near the allocated capacity limit. The alert is sent to the Notifications area on the Home page. This threshold is a percentage of the volume's reported capacity size. Clear the check box to disable the warning threshold alert notification.

5. Click **Save**.

Manage settings

Change settings for a volume

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

Before you begin

The volume you want to change is in Optimal status.


Steps

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

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

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

Field details


Setting	Description
Name	Displays the name of the volume. Change the name of a volume when the current name is no longer meaningful or applicable.
Capacities	<p>Displays the reported and allocated capacity for the selected volume.</p> <p>Reported capacity and allocated capacity are the same for thick volumes, but are different for thin volumes. For a thick volume, the physically allocated space is equal to the space that is reported to the host. For a thin volume, reported capacity is the capacity that is reported to the hosts, whereas allocated capacity is the amount of drive space that is currently allocated for writing data.</p>
Pool / Volume group	Displays the name and RAID level of the pool or volume group. Indicates whether the pool or volume group is secure-capable and secure-enabled.
Host	<p>Displays the volume assignment. You assign a volume to a host or host cluster so it can be accessed for I/O operations. This assignment grants a host or host cluster access to a particular volume or to a number of volumes in a storage array.</p> <ul style="list-style-type: none"> • Assigned to — Identifies the host or host cluster that has access to the selected volume. • LUN — A logical unit number (LUN) is the number assigned to the address space that a host uses to access a volume. The volume is presented to the host as capacity in the form of a LUN. Each host has its own LUN address space. Therefore, the same LUN can be used by different hosts to access different volumes. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p> For NVMe interfaces, this column displays Namespace ID. A namespace is NVM storage that is formatted for block access. It is analogous to a logical unit in SCSI, which relates to a volume in the storage array. The namespace ID is the NVMe controller's unique identifier for the namespace, and can be set to a value between 1 and 255. It is analogous to a logical unit number (LUN) in SCSI.</p> </div>
Identifiers	<p>Displays the identifiers for the selected volume.</p> <ul style="list-style-type: none"> • World-wide identifier (WWID) — A unique hexadecimal identifier for the volume. • Extended unique identifier (EUI) — An EUI-64 identifier for the volume. • Subsystem identifier (SSID) — The storage array subsystem identifier of a volume.

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

Field details

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

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

Setting	Description
SSD Cache	 This feature is not available on the EF600 or EF300 storage system. Shows the SSD Cache setting, which you can enable on compatible volumes as a way to improve read-only performance. Volumes are compatible if they share the same Drive Security and Data Assurance capabilities. The SSD Cache feature uses a single or multiple Solid State Disks (SSDs) to implement a read cache. Application performance is improved because of the faster read times for SSDs. Because the read cache is in the storage array, caching is shared across all applications using the storage array. Simply select the volume that you want to cache, and then caching is automatic and dynamic.

5. Click **Save**.

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

After you finish

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

Change workload settings

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

Steps

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

The Applications & Workloads view appears.
3. Select the workload that you want to change, and then select **View/Edit Settings**.

The Applications & Workloads Settings dialog box appears.
4. **Optional:** Change the user-supplied name of the workload.
5. Click **Save**.

Change cache settings for a volume

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

About this task

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

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

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

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


Steps

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

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


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

Field details

Cache setting	Description
Read Caching	The read cache is a buffer that stores data that has been read from the drives. The data for a read operation might already be in the cache from a previous operation, which eliminates the need to access the drives. The data stays in the read cache until it is flushed.
Write Caching	The write cache is a buffer that stores data from the host that has not yet been written to the drives. The data stays in the write cache until it is written to the drives. Write caching can increase I/O performance.  Cache is automatically flushed after the Write caching is disabled for a volume.

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

Field details

Cache setting	Description
Dynamic Read Cache Prefetch	<p>Dynamic cache read prefetch allows the controller to copy additional sequential data blocks into the cache while it is reading data blocks from a drive to the cache. This caching increases the chance that future requests for data can be filled from the cache. Dynamic cache read prefetch is important for multimedia applications that use sequential I/O. The rate and amount of data that is prefetched into cache is self-adjusting based on the rate and request size of the host reads. Random access does not cause data to be prefetched into cache. This feature does not apply when read caching is disabled.</p> <p>For a thin volume, dynamic cache read prefetch is always disabled and cannot be changed.</p>
Write Caching without Batteries	<p>The write caching without batteries setting lets write caching continue even when the batteries are missing, failed, discharged completely, or not fully charged. Choosing write caching without batteries is not typically recommended, because data might be lost if power is lost. Typically, write caching is turned off temporarily by the controller until the batteries are charged or a failed battery is replaced.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"><p>Possible loss of data — If you select this option and do not have a universal power supply for protection, you could lose data. In addition, you could lose data if you do not have controller batteries and you enable the Write caching without batteries option.</p></div> <p>This setting is available only if you enabled write caching. This setting is not available for thin volumes.</p>
Write Caching with Mirroring	<p>Write caching with mirroring occurs when the data written to the cache memory of one controller is also written to the cache memory of the other controller. Therefore, if one controller fails, the other can complete all outstanding write operations. Write cache mirroring is available only if write caching is enabled and two controllers are present. Write caching with mirroring is the default setting at volume creation.</p> <p>This setting is available only if you enabled write caching. This setting is not available for thin volumes.</p>

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

Change media scan settings for a volume

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

Before you begin

Understand the following:

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

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

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

About this task

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

You can perform the following actions:

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

Steps

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

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

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

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

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

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

6. Click **Save**.

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

Use copy services

Copy Volume overview

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

This function performs a byte-by-byte copy from the source volume to the target volume, making the data on the target volume identical to the data on the source volume.

Data copying for greater access

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

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

Change a thin volume to a thick volume

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



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

Backup data

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

Restore snapshot volume data to the base volume

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

Source and target volumes

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

Volume type	Offline volume copy source volume	Online volume copy source volume	Online and offline target volume
Thick volume in a pool	Yes	Yes	Yes
Thick volume in a volume group	Yes	Yes	Yes
Thin volume	Yes	Yes	No
Snapshot volume	Yes ¹	No	No
Snapshot base volume	Yes	No	No
Remote mirror primary volume	Yes ²	No	Yes

¹ You cannot use the snapshot volume copy until after the online copy operation completes.

² If the source volume is a primary volume, the capacity of the target volume must be equal to or greater than the usable capacity of the source volume.

Types of Copy Volume operations

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

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

The offline and online Copy Volume operations are described below.

Offline Copy Volume operation

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

What you need to know about offline copy operations

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

Online Copy Volume operation

The online Copy Volume relationship is between a snapshot volume and a target volume. You can initiate a Copy Volume operation while the source volume is online and available for data writes. This function is achieved by creating a snapshot of the volume and using the snapshot as the actual source volume for the copy.

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

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

What you need to know about online copy operations

What kind of volumes can be used?	<ul style="list-style-type: none">• The volume for which the point-in-time image is created is known as the base volume and must be a standard volume or a thin volume on the storage array.• A target volume can be a standard volume in a volume group or a standard volume in a pool. A target volume cannot be a thin volume or a base volume in a snapshot group.• You can use the online Copy Volume function to copy data from a thin volume to a standard volume in a pool that resides within the same storage array. But you cannot use the Copy Volume function to copy data from a standard volume to a thin volume.
-----------------------------------	--

What you need to know about online copy operations

Base volume performance	<ul style="list-style-type: none">• If the snapshot volume that is used as the copy source is active, the base volume performance is degraded due to copy-on-write operations. When the copy is complete, the snapshot is disabled, and the base volume performance is restored. Although the snapshot is disabled, the reserved capacity volume and copy relationship remain intact.
Types of volumes created	<ul style="list-style-type: none">• A snapshot volume and a reserved capacity volume are created during the online copy operation.• The snapshot volume is not an actual volume containing data; rather, it is a reference to the data that was contained on a volume at a specific time.• For each snapshot that is taken, a reserved capacity volume is created to hold the data for the snapshot. The reserved capacity volume is used only to manage the snapshot image.
Reserved capacity volume	<ul style="list-style-type: none">• Before a data block on the source volume is modified, the contents of the block to be modified are copied to the reserved capacity volume for safekeeping.• Because the reserved capacity volume stores copies of the original data in those data blocks, further changes to those data blocks write only to the source volume.• The online copy operation uses less disk space than a full physical copy because the only data blocks that are stored in the reserved capacity volume are those that have changed since the time of the snapshot.

Copy volume

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

Before you begin

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

About this task

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

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

More about online and offline copies

Online copy

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

Offline copy

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



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

Steps

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

The Copy Volume-Select Target dialog box appears.

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

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

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

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

More about copy priority rates

There are five copy priority rates:

- Lowest
- Low
- Medium
- High
- Highest

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

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

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

Allocate the reserved capacity using the following guidelines:

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

Results

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

After you finish

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

Take action on a Copy Volume operation

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

Steps


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

The Operations in Progress dialog box appears.

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

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

Action	Description
Stop	<p>You can stop a Copy Volume operation while the operation has a status of In Progress, Pending, or Failed.</p> <p>When the Copy Volume is stopped, all of the mapped hosts have write access to the source volume. If data is written to the source volume, the data on the target volume no longer matches the data on the source volume.</p>

Action	Description
Change priority	You can change the priority of a Copy Volume operation while the operation has a status of In Progress to select the rate at which a Copy Volume operation completes.
Re-copy	<p>You can re-copy a volume when you have stopped a Copy Volume operation and want to start it again or when a Copy Volume operation has failed or halted. The Copy Volume operation starts over from the beginning.</p> <p>The re-copy action overwrites existing data on the target volume and fails all snapshot volumes associated with the target volume, if any exist.</p>
Clear	<p>You can remove the Copy Volume operation while the operation has a status of In Progress, Pending, or Failed.</p> <div style="display: flex; align-items: center; margin-top: 10px;">  <p>Be sure that you want to do this operation before selecting Clear. There is no confirmation dialog.</p> </div>

FAQs

What is a volume?

A volume is a container in which applications, databases, and file systems store data. It is the logical component created for the host to access storage on the storage array.

A volume is created from the capacity available in a pool or a volume group. A volume has a defined capacity. Although a volume might consist of more than one drive, a volume appears as one logical component to the host.

Why am I seeing a capacity over-allocation error when I have enough free capacity in a volume group to create volumes?

The selected volume group might have one or more free capacity areas. A free capacity area is the free capacity that can result from deleting a volume or from not using all available free capacity during volume creation.

When you create a volume in a volume group that has one or more free capacity areas, the volume's capacity is limited to the largest free capacity area in that volume group. For example, if a volume group has a total of 15 GiB free capacity, and the largest free capacity area is 10 GiB, the largest volume you can create is 10 GiB.

If a volume group has free capacity areas, the volume group graph contains a link indicating the number of existing free capacity areas. Select the link to display a pop-over that indicates the capacity of each area.

By consolidating free capacity, you can create additional volumes from the maximum amount of free capacity in a volume group. You can consolidate the existing free capacity on a selected volume group using one of the following methods:

- When at least one free capacity area is detected for a volume group, the "Consolidate free capacity" recommendation appears on the Home page in the Notification area. Click the **Consolidate free capacity**

link to launch the dialog box.

- You can also select **Pools & Volume Groups > Uncommon Tasks > Consolidate volume group free capacity** to launch the dialog box.

If you want to use a specific free capacity area rather than the largest free capacity area, use the Command Line Interface (CLI).

How does my selected workload impact volume creation?

During volume creation, you are prompted for information about a workload's use. The system uses this information to create an optimal volume configuration for you, which can be edited as needed. Optionally, you can skip this step in the volume creation sequence.

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

- **Application-specific** — When you are creating volumes using an application-specific workload, the system may recommend an optimized volume configuration to minimize contention between application workload I/O and other traffic from your application instance. Volume characteristics like I/O type, segment size, controller ownership, and read and write cache are automatically recommended and optimized for workloads that are created for the following application types.
 - Microsoft® SQL Server™
 - Microsoft® Exchange Server™
 - Video surveillance applications
 - VMware ESXi™ (for volumes to be used with Virtual Machine File System)

You can review the recommended volume configuration and edit, add, or delete the system-recommended volumes and characteristics using the Add/Edit Volumes dialog box.

- **Other** (or applications without specific volume creation support) — Other workloads use a volume configuration that you must manually specify when you want to create a workload that is not associated with a specific application, or if there is no built-in optimization for the application you intend to use on the storage array. You must manually specify the volume configuration using the Add/Edit Volumes dialog box.

Why aren't these volumes associated with a workload?

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

Why can't I delete the selected workload?

This workload consists of a group of volumes that were created using the command line interface (CLI) or migrated (imported/exported) from a different storage array. As a result, the volumes in this workload are not associated with an application-specific workload, so

the workload cannot be deleted.

How do application-specific workloads help me manage my storage array?

The volume characteristics of your application-specific workload dictate how the workload interacts with the components of your storage array and helps determine the performance of your environment under a given configuration.

An application is software such as SQL Server or Exchange. You define one or more workloads to support each application. For some applications, the system automatically recommends a volume configuration that optimizes storage. Characteristics such as I/O type, segment size, controller ownership, and read and write cache are included in the volume configuration.

How does providing this information help create storage?

The workload information is used to optimize the volume characteristics such as I/O type, segment size, and read/write cache for the workload selected. These optimized characteristics dictate how your workload interacts with the storage array components.

Based on the workload information you provide, System Manager creates the appropriate volumes and places them on the available pools or volume groups that currently exist on the system. The system creates the volumes and optimizes their characteristics based on the current best practices for the workload you selected.

Before you finish creating volumes for a given workload, you can review the recommended volume configuration and edit, add, or delete the system-recommended volumes and characteristics using the Add/Edit Volumes dialog box.

Refer to your application-specific documentation for best practice information.

What do I need to do to recognize the expanded capacity?

If you increase the capacity for a volume, the host might not immediately recognize the increase in volume capacity.

Most operating systems recognize the expanded volume capacity and automatically expand after the volume expansion is initiated. However, some might not. If your OS does not automatically recognize the expanded volume capacity, you might need to perform a disk rescan or reboot.

After you have expanded the volume capacity, you must manually increase the file system size to match. How you do this depends on the file system you are using.

Refer to your host operating system documentation for additional details.

Why don't I see all my pools and/or volume groups?

Any pool or volume group to which you cannot move the volume does not display in the list.

Pools or volume groups are not eligible for any of the following reasons:

- The Data Assurance (DA) capabilities of a pool or volume group do not match.

- A pool or volume group is in a non-optimal state.
- The capacity of a pool or volume group is too small.

What is segment size?

A segment is the amount of data in kilobytes (KiB) that is stored on a drive before the storage array moves to the next drive in the stripe (RAID group). Segment size applies only to volume groups, not pools.

Segment size is defined by the number of data blocks it contains. When determining segment size, you must know what type of data you will store in a volume. If an application typically uses small, random reads and writes (IOPS), a smaller segment size typically works better. Alternatively, if the application has large, sequential reads and writes (throughput), a large segment size is generally better.

Whether an application uses small random reads and writes, or large sequential reads and writes, the storage array performs better if the segment size is larger than the typical data block chunk size. This normally makes it easier and faster for the drives to access the data, which is important for better storage array performance.

Environments where IOPS performance is important

In an I/O operations per second (IOPS) environment, the storage array performs better if you use a segment size that is larger than the typical data block size (“chunk”) that is read/written to a drive. This ensures that each chunk is written to a single drive.

Environments where throughput is important

In a throughput environment, the segment size should be an even fraction of the total drives for data and the typical data chunk size (I/O size). This spreads the data as a single stripe across the drives in the volume group leading to faster reads and writes.

What is preferred controller ownership?

Preferred controller ownership defines the controller that is designated to be the owning, or primary, controller of the volume.

Controller ownership is very important and should be planned carefully. Controllers should be balanced as closely as possible for total I/Os.

For example, if one controller reads primarily large, sequential data blocks and the other controller has small data blocks with frequent reads and writes, the loads are very different. Knowing which volumes contain what type of data allows you to balance I/O transfers equally over both controllers.

When would I want to use the assign host later selection?

If you want to speed the process for creating volumes, you can skip the host assignment step so that newly created volumes are initialized offline.

Newly created volumes must be initialized. The system can initialize them using one of two modes — either an Immediate Available Format (IAF) background initialization process or an offline process.

When you map a volume to a host, it forces any initializing volumes in that group to transition to background initialization. This background initialization process allows for concurrent host I/O, which can sometimes be

time-consuming.

When none of the volumes in a volume group are mapped, offline initialization is performed. The offline process is much faster than the background process.

What do I need to know about host block size requirements?

For EF300 and EF600 systems, a volume can be set to support a 512-byte or a 4KiB block size (also called "sector size"). You must set the correct value during volume creation. If possible, the system suggests the appropriate default value.

Before setting the volume block size, read the following limitations and guidelines.

- Some operating systems and virtual machines (notably VMware, at this time) require a 512-byte block size and do not support 4KiB, so make sure you know the host requirements before creating a volume. Typically, you can achieve the best performance by setting a volume to present a 4KiB block size; however, ensure that your host allows for 4KiB (or "4Kn") blocks.
- The type of drives you select for your pool or volume group also determines what volume block sizes are supported, as follows:
 - If you create a volume group using drives that write to 512-byte blocks, then you can only create volumes with 512-byte blocks.
 - If you create a volume group using drives that write to 4KiB blocks, then you can create volumes with either 512-byte or 4KiB blocks.
- If the array has an iSCSI host interface card, all volumes are limited to 512-byte blocks (regardless of volume group block size). This is due to a specific hardware implementation.
- You cannot change a block size once it is set. If you need to change a block size, you must delete the volume and re-create it.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.