



Manage volumes

SANtricity 11.6

NetApp

February 12, 2024

Table of Contents

- Manage volumes 1
 - Increase capacity of a volume 1
 - Change settings for a volume 2
 - Initialize volumes 7
 - Redistribute volumes 8
 - Change controller ownership of a volume 9
 - Change cache settings for a volume 9
 - Change media scan settings for a volume 11
 - Delete volume 13

Manage volumes

Increase capacity of a volume

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

Before you begin

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

About this task

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



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

Steps

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

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

3. Select **Yes** to continue.

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

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

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

Results

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

After you finish

After you expand the volume capacity, you must manually increase the file system size to match. How you do

this depends on the file system you are using. See your host operating system documentation for details.

Change settings for a volume

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

Before you begin

The volume you want to change is in Optimal status.


Steps

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

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

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

Field Details

Setting	Description
Name	Displays the name of the volume. Change the name of a volume when the current name is no longer meaningful or applicable.
Capacities	<p>Displays the reported and allocated capacity for the selected volume.</p> <p>Reported capacity and allocated capacity are the same for thick volumes, but are different for thin volumes. For a thick volume, the physically allocated space is equal to the space that is reported to the host. For a thin volume, reported capacity is the capacity that is reported to the hosts, whereas allocated capacity is the amount of drive space that is currently allocated for writing data.</p>
Pool / Volume group	Displays the name and RAID level of the pool or volume group. Indicates whether the pool or volume group is secure-capable and secure-enabled.
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>  <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	<p>Shows the SSD Cache setting, which you can enable on compatible volumes as a way to improve read-only performance. Volumes are compatible if they share the same Drive Security and Data Assurance capabilities.</p> <p>The SSD Cache feature uses a single or multiple Solid State Disks (SSDs) to implement a read cache. Application performance is improved because of the faster read times for SSDs. Because the read cache is in the storage array, caching is shared across all applications using the storage array. Simply select the volume that you want to cache, and then caching is automatic and dynamic.</p>

5. Click **Save**.

Results

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

After you finish

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

Initialize volumes

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

Before you begin

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



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

About this task

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

The following types of volumes *cannot* be initialized:

- Base volume of a snapshot volume
- Primary volume in a mirror relationship
- Secondary volume in a mirror relationship

- Source volume in a volume copy
- Target volume in a volume copy
- Volume that already has an initialization in progress

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

Steps

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

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

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

Results

System Manager performs the following actions:

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

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

Redistribute volumes

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

Before you begin

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

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

About this task

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

Steps

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

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

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

Results

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

Change controller ownership of a volume

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

Before you begin

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

About this task

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

Steps

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

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

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

Results

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

Change cache settings for a volume

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

About this task

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

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

This action might occur when a new battery is charging, when a controller has been removed, or if a

mismatch in cache sizes has been detected by the controller. After the condition has cleared, the cache properties selected in the dialog box become active. If the selected cache properties do not become active, contact technical support.

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


Steps

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

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


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

Field Details

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

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

Field Details

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

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

Change media scan settings for a volume

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

Before you begin

Understand the following:

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

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

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

About this task

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

You can perform the following actions:

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

Steps

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

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

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

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

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

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

6. Click **Save**.

Results

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

Delete volume

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

Before you begin

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

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

About this task

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

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



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



When a volume exceeds a given size (currently 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.

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.