



Create and manage volumes

ONTAP 9

NetApp
September 18, 2024

Table of Contents

- Create and manage volumes 1
 - Create a volume 1
 - Enable large volume and large file support 2
 - SAN volumes 3
 - Determine file and inode usage for a volume 16
 - Control and monitor FlexVol volume I/O performance with Storage QoS 16
 - Delete a FlexVol volume 17
 - Protection against accidental volume deletion 18
 - Commands for managing FlexVol volumes 18
 - Commands for displaying space usage information 19

Create and manage volumes

Create a volume

You can create a volume and specify its junction point and other properties by using the `volume create` command.

About this task

A volume must include a *junction path* for its data to be made available to clients. You can specify the junction path when you create a new volume. If you create a volume without specifying a junction path, you must *mount* the volume in the SVM namespace using the `volume mount` command.

Before you begin

- The SVM for the new volume and the aggregate that will supply the storage to the volume must already exist.
- If the SVM has a list of associated aggregates, the aggregate must be included in the list.
- Beginning in ONTAP 9.13.1, you can create volumes with capacity analytics and Activity Tracking enabled. To enable capacity or Activity Tracking, issue the `volume create` command with `-analytics-state` or `-activity-tracking-state` set to `on`.

To learn more about capacity analytics and Activity Tracking, see [Enable File System Analytics](#).

Steps

1. Create a volume:

```
volume create -vserver svm_name -volume volume_name -aggregate aggregate_name
-size {integer[KB|MB|GB|TB|PB]} -security-style {ntfs|unix|mixed} -user
user_name_or_number -group group_name_or_number -junction-path junction_path
[-policy export_policy_name]
```

The `-security` style, `-user`, `-group`, `-junction-path`, and `-policy` options are for NAS namespaces only.

The choices for `-junction-path` are the following:

- Directly under root, for example, `/new_vol`

You can create a new volume and specify that it be mounted directly to the SVM root volume.

- Under an existing directory, for example, `/existing_dir/new_vol`

You can create a new volume and specify that it be mounted to an existing volume (in an existing hierarchy), expressed as a directory.

If you want to create a volume in a new directory (in a new hierarchy under a new volume), for example, `/new_dir/new_vol`, then you must first create a new parent volume that is junctioned to the SVM root volume. You would then create the new child volume in the junction path of the new parent volume (new directory).

2. Verify that the volume was created with the desired junction point:

```
volume show -vserver svm_name -volume volume_name -junction
```

Examples

The following command creates a new volume named `users1` on the SVM `vs1.example.com` and the aggregate `aggr1`. The new volume is made available at `/users`. The volume is 750 GB in size, and its volume guarantee is of type `volume` (by default).

```
cluster1::> volume create -vserver vs1.example.com -volume users1
-aggregate aggr1 -size 750g -junction-path /users
[Job 1642] Job succeeded: Successful

cluster1::> volume show -vserver vs1.example.com -volume users1 -junction
```

		Junction		Junction
Vserver	Volume	Active	Junction Path	Path Source
vs1.example.com	users1	true	/users	RW_volume

The following command creates a new volume named “home4” on the SVM “vs1.example.com” and the aggregate “aggr1”. The directory `/eng/` already exists in the namespace for the `vs1` SVM, and the new volume is made available at `/eng/home`, which becomes the home directory for the `/eng/` namespace. The volume is 750 GB in size, and its volume guarantee is of type `volume` (by default).

```
cluster1::> volume create -vserver vs1.example.com -volume home4
-aggregate aggr1 -size 750g -junction-path /eng/home
[Job 1642] Job succeeded: Successful

cluster1::> volume show -vserver vs1.example.com -volume home4 -junction
```

		Junction		Junction
Vserver	Volume	Active	Junction Path	Path Source
vs1.example.com	home4	true	/eng/home	RW_volume

Enable large volume and large file support

Beginning with ONTAP 9.12.1 P2, you can create a new volume or modify an existing volume to enable support for a maximum volume size of 300TB and a maximum file (LUN) size of 128TB.

Before you begin

- ONTAP 9.12.1 P2 or later is installed on the cluster.
- If you are enabling large volume support on the source cluster in a SnapMirror relationship, you must have ONTAP 9.12.1 P2 or later installed on the cluster hosting the source volume as well as the cluster hosting the destination volume.
- You are a cluster or SVM administrator.

Create a new volume

Step

1. Create a volume with large volume and file support enabled:

```
volume create -vserver <svm_name> -volume <volume_name> -aggregate  
<aggregate_name> -is-large-size-enabled true
```

Example

The following example creates a new volume with large volume and file size support enabled.

```
volume create -vserver vs1 -volume big_vol1 -aggregate aggr1 -is-large  
-size-enabled true
```

Modify an existing volume

Step

1. Modify a volume to enable large volume and file support:

```
volume modify -vserver <svm_name> -volume <volume_name> -is-large-size  
-enabled true
```

Example

The following example modifies an existing volume to support large volume and file size.

```
volume modify -vserver vs2 -volume data_vol -is-large-size-enabled true
```

Related information

- [Create a volume](#)
- [Command Reference](#)

SAN volumes

Overview of SAN volume provisioning

ONTAP provides several basic options for SAN volume provisioning. Each option uses a different method for managing the volume space and space requirements for the ONTAP block sharing technologies. You should understand how each provisioning option works so you can choose the best option for your environment.



Placing SAN LUNs and NAS shares on the same FlexVol volume is not recommended. You should instead provision separate FlexVol volumes for your SAN LUNs and your NAS shares. This simplifies management and the replication deployments. It also parallels the way FlexVol volumes are supported in Active IQ Unified Manager (formerly OnCommand Unified Manager).

Thin provisioning for volumes

When a thinly provisioned volume is created, ONTAP does not reserve any extra space when the volume is created. As data is written to the volume, the volume requests the storage it needs from the aggregate to accommodate the write operation. Using thin-provisioned volumes enables you to overcommit your aggregate, which introduces the possibility of the volume not being able to secure the space it needs when the aggregate runs out of free space.

You create a thin-provisioned FlexVol volume by setting its `-space-guarantee` option to `none`.

Thick provisioning for volumes

When a thick-provisioned volume is created, ONTAP sets aside enough storage from the aggregate to ensure that any block in the volume can be written to at any time. When you configure a volume to use thick provisioning, you can employ any of the ONTAP storage efficiency capabilities, such as compression and deduplication, to offset the larger upfront storage requirements.

You create a thick-provisioned FlexVol volume by setting its `-space-slo` (service level objective) option to `thick`.

Semi-thick provisioning for volumes

When a volume using semi-thick provisioning is created, ONTAP sets aside storage space from the aggregate to account for the volume size. If the volume is running out of free space because blocks are in use by block-sharing technologies, ONTAP makes an effort to delete protection data objects (Snapshot copies and FlexClone files and LUNs) to free up the space they are holding. As long as ONTAP can delete the protection data objects fast enough to keep pace with the space required for overwrites, the write operations continue to succeed. This is called a “best effort” write guarantee.



You cannot employ storage efficiency technologies such as deduplication, compression, and compaction on a volume that is using semi-thick provisioning.

You create a semi-thick-provisioned FlexVol volume by setting its `-space-slo` (service level objective) option to `semi-thick`.

Use with space-reserved files and LUNs

A space-reserved file or LUN is one for which storage is allocated when it is created. Historically, NetApp has used the term “thin-provisioned LUN” to mean a LUN for which space reservation is disabled (a non-space-reserved LUN).



Non-space-reserved files are not generally referred to as “thin-provisioned files.”

The following table summarizes the major differences in how the three volume provisioning options can be used with space-reserved files and LUNs:

Volume provisioning	LUN/file space reservation	Overwrites	Protection data ²	Storage efficiency ³
Thick	Supported	Guaranteed ¹	Guaranteed	Supported
Thin	No effect	None	Guaranteed	Supported
Semi-thick	Supported	Best effort ¹	Best effort	Not supported

Notes

1. The ability to guarantee overwrites or provide a best-effort overwrite assurance requires that space reservation is enabled on the LUN or file.
2. Protection data includes Snapshot copies, and FlexClone files and LUNs marked for automatic deletion (backup clones).
3. Storage efficiency includes deduplication, compression, any FlexClone files and LUNs not marked for automatic deletion (active clones), and FlexClone subfiles (used for Copy Offload).

Support for SCSI thin-provisioned LUNs

ONTAP supports T10 SCSI thin-provisioned LUNs as well as NetApp thin-provisioned LUNs. T10 SCSI thin provisioning enables host applications to support SCSI features including LUN space reclamation and LUN space monitoring capabilities for blocks environments. T10 SCSI thin provisioning must be supported by your SCSI host software.

You use the `ONTAP space-allocation` setting to enable/disable support for the T10 thin provisioning on a LUN. You use the `ONTAP space-allocation enable` setting to enable T10 SCSI thin provisioning on a LUN.

The `[-space-allocation {enabled|disabled}]` command in the ONTAP Command Reference Manual has more information to enable/disable support for the T10 thin provisioning and to enable T10 SCSI thin provisioning on a LUN.

Related information

- [ONTAP command reference](#)

Configure volume provisioning options

You can configure a volume for thin provisioning, thick provisioning, or semi-thick provisioning, depending on your space requirements.

About this task

Setting the `-space-slo` option to `thick` ensures the following:

- The entire volume is preallocated in the aggregate. You cannot use the `volume create` or `volume modify` command to configure the volume's `-space-guarantee` option.
- 100% of the space required for overwrites is reserved. You cannot use the `volume modify` command to configure the volume's `-fractional-reserve` option

Setting the `-space-slo` option to `semi-thick` ensures the following:

- The entire volume is preallocated in the aggregate. You cannot use the `volume create` or `volume modify` command to configure the volume's `-space-guarantee` option.
- No space is reserved for overwrites. You can use the `volume modify` command to configure the volume's `-fractional-reserve` option.
- Automatic deletion of Snapshot copies is enabled.

Step

1. Configure volume provisioning options:

```
volume create -vserver vserver_name -volume volume_name -aggregate
aggregate_name -space-slo none|thick|semi-thick -space-guarantee none|volume
```

The `-space-guarantee` option defaults to `none` for AFF systems and for non-AFF DP volumes. Otherwise, it defaults to `volume`. For existing FlexVol volumes, use the `volume modify` command to configure provisioning options.

The following command configures vol1 on SVM vs1 for thin provisioning:

```
cluster1::> volume create -vserver vs1 -volume vol1 -space-guarantee
none
```

The following command configures vol1 on SVM vs1 for thick provisioning:

```
cluster1::> volume create -vserver vs1 -volume vol1 -space-slo thick
```

The following command configures vol1 on SVM vs1 for semi-thick provisioning:

```
cluster1::> volume create -vserver vs1 -volume vol1 -space-slo semi-
thick
```

Determine space usage in a volume or aggregate

In some cases, enabling a feature in ONTAP might consume more space than you expected. ONTAP helps you determine how space is being consumed by providing three perspectives from which to view space: the volume, a volume's footprint within the aggregate, and the aggregate.

View space allocation

A volume can run out of space due to space consumption or insufficient space within the volume, aggregate, or a combination of both. By seeing a feature-oriented breakdown of space usage from different perspectives, you can assess which features you might want to adjust or turn off, or whether you should take other action (such as increasing the size of the aggregate or volume).

You can view space usage details from any of these perspectives:

- The volume's space usage

This perspective provides details about space usage within the volume, including usage by Snapshot copies.

Use the `volume show-space` command to see a volume's space usage.

Beginning with ONTAP 9.14.1, on volumes with [temperature-sensitive storage efficiency \(TSSE\)](#) enabled, the amount of space used on the volume reported by the `volume show-space -physical used` command includes the space savings realized as a result of TSSE.

- The volume's footprint within the aggregate

This perspective provides details about the amount of space each volume is using in the containing aggregate, including the volume's metadata.

Use the `volume show-footprint` command to see a volume's footprint with the aggregate.

- The aggregate's space usage

This perspective includes totals of the volume footprints of all volumes contained in the aggregate, space reserved for aggregate Snapshot copies, and other aggregate metadata.

WAFL reserves 10% of the total disk space for aggregate level metadata and performance. The space used for maintaining the volumes in the aggregate comes out of the WAFL reserve and cannot be changed.

Beginning with ONTAP 9.12.1, the WAFL reserve for aggregates greater than 30TB is reduced from 10% to 5% for AFF platforms and for the FAS500f platforms. Beginning in ONTAP 9.14.1, this same reduction applies to aggregates on all FAS platforms, resulting in 5% more usable space in the aggregates.

Use the `storage aggregate show-space` command to see the aggregate's space usage.

Certain features, such as tape backup and deduplication, use space for metadata both from the volume and directly from the aggregate. These features show different space usage between the volume and volume footprint perspectives.

Related Information

- [Knowledge Base article: Space Usage](#)
- [Free up 5% of your storage capacity by upgrading to ONTAP 9.12.1](#)

Volume metadata and data metric reporting

Historically, several of the volume space metrics have reported the total data consumed as a combination of two metrics: metadata and user data. Beginning with ONTAP 9.15.1, the metadata and user data metrics are reported separately. Two new metadata counters have been introduced to support this:

- total-metadata

This counter provides the total metadata size inside the volume. It does not include the aggregate resident volume metadata. Reporting it separately helps to determine the logical data allocated by the user.

- total-metadata-footprint

This counter is the sum of volume resident metadata and aggregate resident volume metadata. It provides the total metadata footprint of the volume inside the aggregate. Reporting it separately helps to determine the physical data allocated by the user.

In addition, several existing counters have been updated to remove the metadata component and present only the user data:

- User data
- Volume data footprint

These changes provide a more accurate view of the data consumed by the user. This has several benefits, including the ability to make more precise chargeback decisions.

Delete Snapshot copies automatically

You can define and enable a policy for automatically deleting Snapshot copies and FlexClone LUNs. Automatically deleting Snapshot copies and FlexClone LUNs can help you manage space utilization.

About this task

You can automatically delete Snapshot copies from read-write volumes and FlexClone LUNs from read-write parent volumes. You cannot set up automatic deletion of Snapshot copies from read-only volumes, for example, SnapMirror destination volumes.

Step

1. Define and enable a policy for automatically deleting Snapshot copies by using the `volume snapshot autodelete modify` command.

See the `volume snapshot autodelete modify` man page for information about the parameters that you can use with this command to define a policy that meets your needs.

The following command enables the automatic deletion of Snapshot copies and sets the trigger to `snap_reserve` for the `vol3` volume, which is part of the `vs0.example.com` storage virtual machine (SVM):

```
cluster1::> volume snapshot autodelete modify -vserver vs0.example.com
-volume vol3 -enabled true -trigger snap_reserve
```

The following command enables the automatic deletion of Snapshot copies and of FlexClone LUNs marked for autodeletion for the `vol3` volume, which is part of the `vs0.example.com` storage virtual machine (SVM):

```
cluster1::> volume snapshot autodelete modify -vserver vs0.example.com
-volume vol3 -enabled true -trigger volume -commitment try -delete-order
oldest_first -destroy-list lun_clone,file_clone
```

Aggregate-level Snapshot copies work differently than volume-level Snapshot copies and are managed automatically by ONTAP. The option to delete aggregate Snapshot copies is always enabled and helps in managing space utilization.



If the trigger parameter is set to `snap_reserve` for an aggregate, the Snapshot copies are maintained until the space reserved crosses the threshold capacity. Therefore, even if the trigger parameter is not set to `snap_reserve`, the space used by the Snapshot copy in the command will be listed as 0 because these Snapshot copies are automatically deleted. Also, the space used by Snapshot copies in an aggregate is considered as free and is included in the available space parameter of the command.

Configure volumes to automatically provide more space when they are full

When FlexVol volumes get full, ONTAP can use various methods to attempt to automatically provide more free space for the volume. You choose which methods ONTAP can use, and in which order, depending on the requirements imposed by your application and storage architecture.

About this task

ONTAP can automatically provide more free space for a full volume by using one or both of the following methods:

- Increase the size of the volume (known as *autogrow*).

This method is useful if the volume's containing aggregate has enough space to support a larger volume. You can configure ONTAP to set a maximum size for the volume. The increase is automatically triggered based on the amount of data being written to the volume in relation to the current amount of used space and any thresholds set.

Autogrow is not triggered to support Snapshot copy creation. If you attempt to create a Snapshot copy and there is insufficient space, the Snapshot copy creation fails, even with autogrow enabled.

- Delete Snapshot copies, FlexClone files, or FlexClone LUNs.

For example, you can configure ONTAP to automatically delete Snapshot copies that are not linked to Snapshot copies in cloned volumes or LUNs, or you can define which Snapshot copies you want ONTAP to delete first—your oldest or newest Snapshot copies. You can also determine when ONTAP should begin deleting Snapshot copies—for example, when the volume is nearly full or when the volume's Snapshot reserve is nearly full.

If you enable both of these methods, you can specify which method ONTAP tries first when a volume is nearly full. If the first method does not provide sufficient additional space to the volume, ONTAP tries the other method next.

By default, ONTAP tries to increase the size of the volume first. In most cases, the default configuration is preferable, because when a Snapshot copy is deleted, it cannot be restored. However, if you need to avoid growing the size of a volume whenever possible, you can configure ONTAP to delete Snapshot copies before increasing the size of the volume.

Steps

1. If you want ONTAP to attempt to increase the size of the volume when it gets full, enable the autogrow capability for the volume by using the `volume autosize` command with `grow` mode.

Remember that when the volume grows, it consumes more free space from its associated aggregate. If you are depending on the volume's ability to grow whenever it needs to, you must monitor the free space in the associated aggregate and add more when needed.

2. If you want ONTAP to delete Snapshot copies, FlexClone files, or FlexClone LUNs when the volume gets full, enable autodelete for those object types.
3. If you enabled both the volume autogrow capability and one or more autodelete capabilities, select the first method that ONTAP should use to provide free space to a volume by using the `volume modify` command with the `-space-mgmt-try-first` option.

To specify increasing the size of the volume first (the default), use `volume_grow`. To specify deleting Snapshot copies first, use `snap_delete`.

Configure volumes to automatically grow and shrink their size

You can configure FlexVol volumes to automatically grow and shrink according to how much space they currently require. Automatic growing helps prevent a volume from running out of space, if the aggregate can supply more space. Automatic shrinking prevents a volume from being larger than needed, freeing space in the aggregate for use by other volumes.

About this task

Autoshrink can only be used in combination with autogrow to meet changing space demands and is not available alone. When autoshrink is enabled, ONTAP automatically manages the shrinking behavior of a volume to prevent an endless loop of autogrow and autoshrink actions.

As a volume grows, the maximum number of files it can contain might be automatically increased. When a volume is shrunk, the maximum number of files it can contain is left unchanged, and a volume cannot be automatically shrunk below the size that corresponds to its current maximum number of files. For this reason, it might not be possible to automatically shrink a volume all the way to its original size.

By default, the maximum size a volume can grow to is 120% of the size at which autogrow is enabled. If you need to ensure that the volume can grow to be larger than that, you must set the maximum size for the volume accordingly.

Before you begin

The FlexVol volume must be online.

Step

1. Configure the volume to grow and shrink its size automatically:

```
volume autosize -vserver SVM_name -volume volume_name -mode grow_shrink
```

The following command enables automatic size changes for a volume called `test2`. The volume is configured to begin shrinking when it is 60% full. The default values are used for when it will begin to grow and its maximum size.

```
cluster1::> volume autosize -vserver vs2 test2 -shrink-threshold-percent 60
vol autosize: Flexible volume "vs2:test2" autosize settings UPDATED.

Volume modify successful on volume: test2
```

Requirements for enabling both autoshrink and automatic Snapshot copy deletion

The autoshrink functionality can be used with automatic Snapshot copy deletion as long as certain configuration requirements are met.

If you want to enable both the autoshrink functionality and automatic Snapshot copy deletion, your configuration must meet the following requirements:

- ONTAP must be configured to attempt to increase volume size before trying to delete Snapshot copies(the `-space-mgmt-try-first` option must be set to `volume_grow`).
- The trigger for automatic Snapshot copy deletion must be volume fullness(the `trigger` parameter must be set to `volume`).

Autoshrink functionality and snapshot copy deletion

Because the autoshrink functionality shrinks the size of a FlexVol volume, it can also affect when volume Snapshot copies are automatically deleted.

The autoshrink functionality interacts with automatic volume Snapshot copy deletion in the following ways:

- If both the `grow_shrink` autosize mode and automatic Snapshot copy deletion are enabled, when a volume size shrinks it can trigger an automatic Snapshot copy deletion.

This is because the Snapshot reserve is based on a percentage of the volume size (5 percent by default), and that percentage is now based on a smaller volume size. This can cause Snapshot copies to spill out of the reserve and be deleted automatically.

- If the `grow_shrink` autosize mode is enabled and you manually delete a Snapshot copy, it might trigger an automatic volume shrinkage.

Address FlexVol volume fullness and overallocation alerts

ONTAP issues EMS messages when FlexVol volumes are running out of space so that you can take corrective action by providing more space for the full volume. Knowing the types of alerts and how to address them helps you ensure your data availability.

When a volume is described as *full*, it means that the percentage of the space in the volume available for use by the active file system (user data) has fallen below a (configurable) threshold. When a volume becomes *overallocated*, the space used by ONTAP for metadata and to support basic data access has been exhausted. Sometimes space normally reserved for other purposes can be used to keep the volume functioning, but space reservation or data availability can be at risk.

Overallocation can be either logical or physical. *Logical overallocation* means that space reserved to honor

future space commitments, such as space reservation, has been used for another purpose. *Physical overallocation* means that the volume is running out of physical blocks to use. Volumes in this state are at risk for refusing writes, going offline, or potentially causing a controller disruption.

A volume can be more than 100% full due to space used or reserved by metadata. However, a volume that is more than 100% full might or might not be overallocated. If qtrees-level and volume-level shares exist on the same FlexVol or SCVMM pool, the qtrees appear as directories on the FlexVol share. Therefore, you need to be careful not to delete them accidentally.

The following table describes the volume fullness and overallocation alerts, the actions you can take to address the issue, and the risks of not taking action:

Alert type	EMS level	Configurable?	Definition	Ways to address	Risk if no action taken
Nearly full	Debug	Y	The file system has exceeded the threshold set for this alert (the default is 95%). The percentage is the <code>Used</code> total minus the size of the Snapshot reserve.	<ul style="list-style-type: none"> Increasing volume size Reducing user data 	No risk to write operations or data availability yet.
Full	Debug	Y	The file system has exceeded the threshold set for this alert (the default is 98%). The percentage is the <code>Used</code> total minus the size of the Snapshot reserve.	<ul style="list-style-type: none"> Increasing volume size Reducing user data 	No risk to write operations or data availability yet, but the volume is approaching the stage where write operations could be at risk.
Logically overallocated	SVC Error	N	In addition to the file system being full, the space in the volume used for metadata has been exhausted.	<ul style="list-style-type: none"> Increasing volume size Deleting Snapshot copies Reducing user data Disabling space reservation for files or LUNs 	Write operations to unreserved files could fail.

Alert type	EMS level	Configurable?	Definition	Ways to address	Risk if no action taken
Physically overallocated	Node Error	N	The volume is running out of physical blocks it can write to.	<ul style="list-style-type: none"> • Increasing volume size • Deleting Snapshot copies • Reducing user data 	Write operations are at risk, as well as data availability; the volume could go offline.

Every time a threshold is crossed for a volume, whether the fullness percentage is rising or falling, an EMS message is generated. When the fullness level of the volume falls below a threshold, a `volume ok` EMS message is generated.

Address aggregate fullness and overallocation alerts

ONTAP issues EMS messages when aggregates are running out of space so that you can take corrective action by providing more space for the full aggregate. Knowing the types of alerts and how you can address them helps you ensure your data availability.

When an aggregate is described as *full*, it means that the percentage of the space in the aggregate available for use by volumes has fallen below a predefined threshold. When an aggregate becomes *overallocated*, the space used by ONTAP for metadata and to support basic data access has been exhausted. Sometimes space normally reserved for other purposes can be used to keep the aggregate functioning, but volume guarantees for volumes associated with the aggregate or data availability can be at risk.

Overallocation can be either logical or physical. *Logical overallocation* means that space reserved to honor future space commitments, such as volume guarantees, has been used for another purpose. *Physical overallocation* means that the aggregate is running out of physical blocks to use. Aggregates in this state are at risk for refusing writes, going offline, or potentially causing a controller disruption.

The following table describes the aggregate fullness and overallocation alerts, the actions you can take to address the issue, and the risks of not taking action.

Alert type	EMS Level	Configurable?	Definition	Ways to address	Risk if no action taken
Nearly full	Debug	N	The amount of space allocated for volumes, including their guarantees, has exceeded the threshold set for this alert (95%). The percentage is the <code>Used</code> total minus the size of the Snapshot reserve.	<ul style="list-style-type: none"> • Adding storage to the aggregate • Shrinking or deleting volumes • Moving volumes to another aggregate with more space • Removing volume guarantees (setting them to <code>none</code>) 	No risk to write operations or data availability yet.
Full	Debug	N	The file system has exceeded the threshold set for this alert (98%). The percentage is the <code>Used</code> total minus the size of the Snapshot reserve.	<ul style="list-style-type: none"> • Adding storage to the aggregate • Shrinking or deleting volumes • Moving volumes to another aggregate with more space • Removing volume guarantees (setting them to <code>none</code>) 	Volume guarantees for volumes in the aggregate might be at risk, as well as write operations to those volumes.
Logically overloaded	SV C Error	N	In addition to the space reserved for volumes being full, the space in the aggregate used for metadata has been exhausted.	<ul style="list-style-type: none"> • Adding storage to the aggregate • Shrinking or deleting volumes • Moving volumes to another aggregate with more space • Removing volume guarantees (setting them to <code>none</code>) 	Volume guarantees for volumes in the aggregate are at risk, as well as write operations to those volumes.
Physically overloaded	Node Error	N	The aggregate is running out of physical blocks it can write to.	<ul style="list-style-type: none"> • Adding storage to the aggregate • Shrinking or deleting volumes • Moving volumes to another aggregate with more space 	Write operations to volumes in the aggregate are at risk, as well as data availability; the aggregate could go offline. In extreme cases, the node could experience a disruption.

Every time a threshold is crossed for an aggregate, whether the fullness percentage is rising or falling, an EMS

message is generated. When the fullness level of the aggregate falls below a threshold, an `aggregate ok` EMS message is generated.

Considerations when setting fractional reserve

Fractional reserve, also called *LUN overwrite reserve*, enables you to turn off overwrite reserve for space-reserved LUNs and files in a FlexVol volume. This can help you maximize your storage utilization.



If your environment is negatively affected by write operations failing due to lack of space, you must understand the requirements that this configuration can impose.

The fractional reserve setting is expressed as a percentage; the only valid values are 0 and 100 percent. The fractional reserve setting is an attribute of the volume. Setting fractional reserve to 0 increases your storage utilization. However, an application accessing data residing in the volume could experience a data outage if the volume is out of free space, even with the volume guarantee set to `volume`. With proper volume configuration and use, however, you can minimize the chance of writes failing. ONTAP provides a “best effort” write guarantee for volumes with fractional reserve set to 0 when *all* of the following requirements are met:

- Deduplication is not in use
- Compression is not in use
- FlexClone sub-files are not in use
- All FlexClone files and FlexClone LUNs are enabled for automatic deletion

This is not the default setting. You must explicitly enable automatic deletion, either at creation time or by modifying the FlexClone file or FlexClone LUN after it is created.

- ODX and FlexClone copy offload are not in use
- Volume guarantee is set to `volume`
- File or LUN space reservation is `enabled`
- Volume Snapshot reserve is set to 0
- Volume Snapshot copy automatic deletion is `enabled` with a commitment level of `destroy`, a `destroy list` of `lun_clone`, `vol_clone`, `cifs_share`, `file_clone`, `sfsr`, and a `trigger` of `volume`

This setting also ensures that FlexClone files and FlexClone LUNs are deleted when necessary.



- If all the above requirements are met but your rate of change is high, in rare cases, the Snapshot copy automatic deletion could fall behind, resulting in the volume running out of space.
- If all the above requirements are met and Snapshot copies are not in use, volume writes are guaranteed to not run out of space.

In addition, you can optionally use the volume autogrow capability to decrease the likelihood of volume Snapshot copies needing to be deleted automatically. If you enable the autogrow capability, you must monitor the free space in the associated aggregate. If the aggregate becomes full enough that the volume is prevented from growing, more Snapshot copies will probably be deleted as the free space in the volume is depleted.

If you cannot meet all of the above configuration requirements and you need to ensure that the volume does

not run out of space, you must set the volume's fractional reserve setting to 100. This requires more free space up front, but guarantees that data modification operations will succeed even when the technologies listed above are in use.

The default value and allowed values for the fractional reserve setting depend on the guarantee of the volume:

Volume guarantee	Default fractional reserve	Allowed values
Volume	100	0, 100
None	0	0, 100

Determine file and inode usage for a volume

FlexVol volumes have a maximum number of files that they can contain. You can use a CLI command to determine whether you need to increase the number of (public) inodes for your FlexVol volumes to prevent them from hitting their file limit.

About this task

Public inodes can be either free (they are not associated with a file) or used (they point to a file). The number of free inodes for a volume is the total number of inodes for the volume minus the number of used inodes (the number of files).

If qtree-level and volume-level shares exist on the same FlexVol or SCVMM pool, the qtrees appear as directories on the FlexVol share. Therefore, you need to be careful not to delete them accidentally.

Steps

1. To display inode usage for a volume, enter the following command:

```
volume show -vserver <SVM_name> -volume <volume_name> -fields files
```

Example

```
cluster1::*> volume show -vserver vs1 -volume vol1 -fields files
Vserver Name: vs1
Files Used (for user-visible data): 98
```

Control and monitor FlexVol volume I/O performance with Storage QoS

You can control input/output (I/O) performance to FlexVol volumes by assigning volumes to Storage QoS policy groups. You might control I/O performance to ensure that workloads achieve specific performance objectives or to throttle a workload that negatively impacts other workloads.

About this task

Policy groups enforce a maximum throughput limit (for example, 100 MB/s). You can create a policy group without specifying a maximum throughput, which enables you to monitor performance before you control the workload.

You can also assign SVMs, LUNs, and files to policy groups.

Note the following requirements about assigning a volume to a policy group:

- The volume must be contained by the SVM to which the policy group belongs.

You specify the SVM when you create the policy group.

- If you assign a volume to a policy group, then you cannot assign the volume's containing SVM or any child LUNs or files to a policy group.

For more information about how to use Storage QoS, see the [System Administration Reference](#).

Steps

1. Use the `qos policy-group create` command to create a policy group.
2. Use the `volume create` command or the `volume modify` command with the `-qos-policy-group` parameter to assign a volume to a policy group.
3. Use the `qos statistics` commands to view performance data.
4. If necessary, use the `qos policy-group modify` command to adjust the policy group's maximum throughput limit.

Delete a FlexVol volume

You can delete a FlexVol volume that is no longer required.

What you'll need

No applications must be accessing the data in the volume you want to delete.



If you accidentally delete a volume, see the Knowledge Base article [How to use the Volume Recovery Queue](#).

Steps

1. If the volume has been mounted, unmount it:

```
volume unmount -vserver vsilver_name -volume volume_name
```

2. If the volume is part of a SnapMirror relationship, delete the relationship by using the `snapmirror delete` command.
3. If the volume is online, take the volume offline:

```
volume offline -vserver vsilver_name volume_name
```

4. Delete the volume:

```
volume delete -vserver vsilver_name volume_name
```

Result

The volume is deleted, along with any associated quota policies and qtrees.

Protection against accidental volume deletion

Default volume delete behavior aids the recovery of accidentally deleted FlexVol volumes.

A `volume delete` request against a volume that has type RW or DP (as seen in `volume show` command output) causes that volume to be moved to a partially deleted state. By default, it is retained in a recovery queue for at least 12 hours before being fully deleted.

For more information, see the KnowledgeBase article [How to use the Volume Recovery Queue](#).

Commands for managing FlexVol volumes

The ONTAP CLI provides specific commands for managing FlexVol volumes. Depending on what you need to do, you can use the following commands to manage FlexVol volumes:

If you want to...	Use this command...
Bring a volume online	<code>volume online</code>
Change the size of a volume	<code>volume size</code>
Determine the associated aggregate of a volume	<code>volume show</code>
Determine the associated aggregate for all volumes on a storage virtual machine (SVM)	<code>volume show -vserver -fields aggregate</code>
Determine the format of a volume	<code>volume show -fields block-type</code>
Mount a volume onto another volume using a junction	<code>volume mount</code>
Put a volume into the restricted state	<code>volume restrict</code>
Rename a volume	<code>volume rename</code>
Take a volume offline	<code>volume offline</code>

See the man page for each command for more information.

Commands for displaying space usage information

You use the `storage aggregate` and `volume` commands to see how space is being used in your aggregates and volumes and their Snapshot copies.

To display information about...	Use this command...
Aggregates, including details about used and available space percentages, Snapshot reserve size, and other space usage information	<code>storage aggregate show storage aggregate show-space -fields snap-size-total,used-including-snapshot-reserve</code>
How disks and RAID groups are used in an aggregate, and RAID status	<code>storage aggregate show-status</code>
The amount of disk space that would be reclaimed if you deleted a specific Snapshot copy	<code>volume snapshot compute-reclaimable (advanced)</code>
The amount of space used by a volume	<code>volume show -fields size,used,available,percent-used volume show-space</code>
The amount of space used by a volume in the containing aggregate	<code>volume show-footprint</code>

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.