



Volume administration

ONTAP 9

NetApp
February 06, 2026

This PDF was generated from <https://docs.netapp.com/us-en/ontap/volumes/manage-volumes-task.html> on February 06, 2026. Always check docs.netapp.com for the latest.

Table of Contents

Volume administration	1
Volume and LUN management with System Manager	1
Manage volumes	1
Manage LUNs with ONTAP System Manager	7
Expand storage with ONTAP System Manager	10
Save storage space using compression, compaction, and deduplication with ONTAP System Manager	11
Balance loads by moving LUNs with ONTAP System Manager	11
Balance loads by moving volumes to another tier with ONTAP System Manager	12
Use Ansible Playbooks to add or edit volumes or LUNs with ONTAP System Manager	12
Manage storage efficiency policies with ONTAP System Manager	14
Manage resources using quotas with ONTAP System Manager	16
Set quotas to limit resource use with ONTAP System Manager	16
Clone volumes and LUNs for testing with ONTAP System Manager	17
Search, filter, and sort information in ONTAP System Manager	19
Logical storage management with the CLI	21
Logical storage management overview with the CLI	21
Create and manage volumes	22
Move and copy volumes	41
Use FlexClone volumes to create efficient copies of your FlexVol volumes	49
Use FlexClone files and FlexClone LUNs to create efficient copies of files and LUNs	54
Use qtrees to partition your FlexVol volumes	63
Logical space reporting and enforcement for volumes	68
Use quotas to restrict or track resource usage	75
Use deduplication, data compression, and data compaction to increase storage efficiency	117
Rehost a volume from one SVM to another SVM	147
Recommended volume and file or LUN configuration combinations	154
Cautions and considerations for changing file or directory capacity	159
Features supported by FlexClone files and FlexClone LUNs	161
FlexGroup volumes management	164
Learn about ONTAP FlexGroup volumes management with the CLI	164
Learn about ONTAP FlexGroup volumes	164
Supported and unsupported configurations for ONTAP FlexGroup volumes	165
FlexGroup volume setup	170
Manage FlexGroup volumes	178
Data protection for FlexGroup volumes	214
Manage data protection operations for FlexGroup volumes	232
Convert FlexVol volumes to FlexGroup volumes	250
FlexCache volumes management	257
Learn about ONTAP FlexCache volumes	257
Supported and unsupported features for ONTAP FlexCache volumes	259
Guidelines for sizing ONTAP FlexCache volumes	264
Create ONTAP FlexCache volumes	265
FlexCache write-back	270

FlexCache duality	286
Manage FlexCache volumes	293
FlexCache for hotspot remediation	300

Volume administration

Volume and LUN management with System Manager

Manage volumes

Manage ONTAP volumes with System Manager

After you display a list of volumes in System Manager, you can perform various actions to manage the volumes.

Some volume types are not available using System Manager, including the following volumes:

- Vol0
- DEL and TMP type volumes
- FlexGroup constituents
- Replicated volumes in a MetroCluster configuration

Steps

1. In System Manager, click **Storage > Volumes**.

The list of volumes is displayed.


2. You can perform the following:

To perform this task...	Take these actions...
Add a volume	Click  Add . See Add a volume .
Manage multiple volumes	<p>Check the boxes next to the volumes.</p> <ul style="list-style-type: none">• Click  Delete to delete the selected volumes.• Click  Protect to assign a protection policy to the selected volumes.• Click  More to select one of the following actions to perform for all selected volumes:<ul style="list-style-type: none">◦ Enable quota◦ Take offline◦ Move◦ Show Deleted Volumes

Manage a single volume	<p>Next to the volume, click , then select one of the following actions to perform:</p> <ul style="list-style-type: none"> • Edit • Resize (Beginning with ONTAP 9.10.1, and only for online volumes and DP FlexVol volumes) • Delete • Clone • Take Offline (or Bring Online) • Enable Quota (or Disable Quota) • Edit Export Policy • Edit Mount Path • Move • Edit Cloud Tier Settings • Protect
Rename a volume	<p>You can rename a volume from the overview page.</p> <p>Click  next to the volume name, and then modify the name of the volume.</p>

Add a volume

You can create a volume and add it to an existing storage VM that is configured for NFS or SMB service.



Before you begin

- A storage VM that is configured for NFS or SMB service should exist in the cluster.
- Beginning with ONTAP 9.13.1, you can enable capacity analytics and Activity Tracking by default on new volumes. In System Manager, you can manage default settings at the cluster or storage VM level. For more information see [Enable File System Analytics](#).

Steps

1. Go to **Storage > Volumes**.
2. Select **+ Add**.
3. Specify a name and size for the volume.
4. Perform one of the following steps:

Select this button...	To perform this action...
Save	The volume is created and added using the system defaults. No additional steps are required.
More Options	Proceed to Step 5 to define the specifications for the volume.

5. The volume name and size are shown if you previously specified them. Otherwise, enter the name and size.
6. Select a storage VM from the pull-down list.

Only storage VMs configured with the NFS protocol are listed. If only one storage VM configured with the

NFS protocol is available, the **Storage VM** field is not shown.

7. To add a cache for the remote volume, select **Add a cache for remote volume** and specify the following values:
 - Select a cluster.
 - Select a storage VM.
 - Select the volume that you want to be a cache volume.
8. In the **Storage and Optimization** section, specify the following values:
 - a. The capacity of the volume is already shown, but you can modify it.
 - b. In the **Performance Service Level** field, select a service level:

When you select this service level...	This occurs...
An existing service level, such as "Extreme", "Performance", or "Value". Only the service levels that are valid for the system platform (AFF, FAS, or others) are displayed.	A local tier or tiers are automatically chosen. Proceed to Step 9 .
Custom	Proceed to Step 8c to define a new service level.

- c. Beginning with ONTAP 9.9.1, you can use System Manager to manually select the local tier on which you want to place the volume you are creating (if you have selected the "Custom" service level).



This option is not available if you select **Add as a cache for a remote volume** or **Distribute volume data across the cluster** (see below).

When you make this choice...	You perform these steps...
Manual placement	Manual placement is enabled. The Distribute volume data across the cluster selection is disabled (see below). Proceed to Step 8d to complete the process.
No selection	Manual placement is not enabled. The local tier is automatically selected. Proceed to Step 9 .

- d. Select a local tier from the pull-down menu.
- e. Select a QoS policy.

Select "Existing" to choose from a list of existing policies, or select "New" to enter the specifications of a new policy.

9. In the **Optimization options** section, determine if you want to distribute the volume data across the cluster:

When you make this choice...	This occurs...

Distribute volume data across the cluster	The volume you are adding becomes a FlexGroup volume. This option is not available if you previously selected Manual placement .
No selection	The volume you are adding becomes a FlexVol volume by default.

10. In the **Access Permissions** section, specify the access permissions for the protocols for which the volume is configured.

Beginning with ONTAP 9.11.1, the new volume will not be shareable by default. You can specify the default access permissions by ensuring the following check boxes are checked:

- **Export via NGS**: Creates the volume with the default export policy that grants users full access to the data.
- **Share via SMB/CIFS**: Creates a share with an auto-generated name, which you can edit. Access is granted to Everyone. Also, you can specify the permission level.

11. In the **Protection** section, specify the protections for the volume.

- Beginning with ONTAP 9.12.1, you can select **Enable Snapshot (Local)** and choose a Snapshot policy rather than using the default.
- If you select **Enable SnapMirror (Local or Remote)**, then specify the protection policy and settings for the destination cluster from the pull-down lists.

12. Select **Save**.

The volume is created and added to the cluster and storage VM.



You can also save the specifications of this volume to an Ansible Playbook. For more details, go to [Use Ansible Playbooks to add or edit volumes or LUNs](#).

Assign tags to volumes

Beginning with ONTAP 9.14.1, you can use System Manager to assign tags to volumes to identify objects as belonging to a category, such as projects or cost centers.

About this task

You can assign a tag to a volume. First, you need to define and add the tag. Then, you can also edit or delete the tag.

Tags can be added when you create a volume, or they can be added later.

You define a tag by specifying a key and associating a value to it using the format “key:value”. For example: “dept:engineering” or “location:san-jose”.

The following should be considered when you create tags:

- Keys have a minimum length of one character and cannot be null. Values can be null.
- A key can be paired with multiple values by separating the values with a comma, for example, “location:san-jose,toronto”
- Tags can be used for multiple resources.

- Keys must start with a lowercase letter.
- Tags assigned to volumes will be deleted when the volume is deleted.
- Tags are not recovered if a volume is recovered from the recovery queue.
- Tags are retained if the volume is moved or cloned.
- Tags assigned to storage VMs in a disaster recovery relationship are replicated on the volume on the partner site.

Steps

To manage tags, perform the following steps:

1. In System Manager, click **Volumes**, and then select the volume to which you want to add a tag.

The tags are listed in the **Tags** section.

2. Click **Manage Tags** to modify existing tags or add new ones.

You can add, edit, or delete the tags.

To perform this action...	Perform these steps...
Add a tag	<ol style="list-style-type: none"> Click Add Tag. Specify a key and its value or values (separate multiple values with commas). Click Save.
Edit a tag	<ol style="list-style-type: none"> Modify the content in the Key and Values (optional) fields. Click Save.
Delete a tag	<ol style="list-style-type: none"> Click  next to the tag you want to delete.

Recover deleted volumes

If you have accidentally deleted one or more FlexVol volumes, you can use System Manager to recover these volumes. Beginning with ONTAP 9.8, you can also use System Manager to recover FlexGroup volumes. You can also delete the volumes permanently by purging the volumes.

The volume retention time can be set on a storage VM level. By default, the volume retention time is set to 12 hours.

Selecting deleted volumes

Steps

1. Click **Storage > Volumes**.
2. Click **More > Show Deleted Volumes**.
3. Select the volumes and click the desired action to recover or permanently delete the volumes.

Resetting the volume configurations

Deleting a volume deletes the associated configurations of the volume. Recovering a volume does not reset all the configurations. Perform the following tasks manually after recovering a volume to bring the volume back to its original state:


Steps

1. Rename the volume.
2. Set up a junction path (NAS).
3. Create mappings for LUNs in the volume (SAN).
4. Associate a snapshot policy and export policy with the volume.
5. Add new quota policy rules for the volume.
6. Add a QOS policy for the volume.

Manage LUNs with ONTAP System Manager

You can create LUNs and add them to an existing storage VM that is configured with the SAN protocol. You can also group LUNs or rename them.

Add LUNs



Before you Start

A storage VM that is configured for SAN service should exist in the cluster.

Steps

1. Go to **Storage > LUNs**.
2. Click **+ Add**.
3. Specify a prefix that will be used at the start of each LUN name. (If you are creating only one LUN, enter

the LUN name.)

4. Select a storage VM from the pull-down list.

Only storage VMs that are configured for the SAN protocol are listed. If only one storage VM that is configured for the SAN protocol is available, then the **Storage VM** field is not displayed.

5. Indicate how many LUNs you want to create and the size of each LUN.
6. Select the host operating system and LUN format from the pull-down lists.
7. Enter the host initiators, and separate them with commas.
8. Perform one of the following actions:

Click this button...	To perform this action...
Save	The LUNs are created with the specifications you entered. System defaults are used for other specifications. No additional steps are required.
More Options	Proceed to Step 9 to define additional specifications for the LUNs.

9. The LUN prefix is already shown if you previously entered it, but you can modify it. Otherwise, enter the prefix.
10. Select a storage VM from the pull-down list.

Only storage VMs that are configured for the SAN protocol are listed. If only one storage VM that is configured for the SAN protocol is available, then the **Storage VM** field is not displayed.

11. Determine how you want the LUNs to be grouped:

When you make this choice...	This occurs...
Group with related LUNs	The LUNs will be grouped together with related LUNs on an existing volume on the storage VM.
No selection	The LUNs will be grouped together on a volume called "container".

12. In the **Storage and Optimization** section, specify the following values:

- The number and capacity of the LUNs are already shown if you previously entered them, but you can modify them. Otherwise, enter the values.
- In the **Performance Service Level** field, select a service level:

When you select this service level...	This occurs...
An existing service level, such as "Extreme", "Performance", or "Value".	A local tier is automatically chosen. Proceed to Step 13 .
Only the service levels that are valid for the system platform (AFF, FAS, or others) are displayed.	
Custom	Proceed to Step 12c to define a new service level.

c. Beginning with ONTAP 9.9.1, you can use System Manager to manually select the local tier on which you want to place the LUNs you are creating (if you have selected the "Custom" service level).

When you make this choice...	You perform these steps...
Manual placement	Manual placement is enabled. Proceed to Step 12d to complete the process.
No selection	Manual selection is not enabled. The local tier is automatically selected. Proceed to Step 13 .

d. Select a local tier from the pull-down menu.

e. Select a QoS policy.

Select "Existing" to choose from a list of existing policies, or select "New" to enter the specifications of a new policy.

13. In the **Host Information** section, the host operating system and LUN format are already shown, but you can modify them.

14. Under **Host Mapping**, select the type of initiators for the LUNs:

- **Existing initiator group:** Select an initiator group for the list that displays.
- **New initiator group using existing initiator groups:** Specify the name of the new group, and select the group or groups that you want to use to create the new group.
- **Host initiators:** Specify a name from the new initiator group, and click **+Add Initiator** to add initiators to the group.

15. In the **Protection** section, specify the protections for the LUNs.

If you select **Enable SnapMirror (Local or Remote)**, then specify the protection policy and settings for the destination cluster from the pull-down lists.

16. Click **Save**.

The LUNs are created and added to the cluster and storage VM.



You can also save the specifications of these LUNs to an Ansible Playbook. For more details, go to [Use Ansible Playbooks to add or edit volumes or LUNs](#).

Rename a LUN

You can rename a LUN from the overview page.

Steps

1. In System Manager, click **LUNs**.
2. Click next to the name of the LUN you want to rename, then modify the LUN name.
3. Click **Save**.

Expand storage with ONTAP System Manager

Using System Manager, you can increase the size of your volume or LUN so that more space is available to your host. The size of a LUN cannot exceed the size of the containing volume.

Beginning with ONTAP 9.12.1, when you enter the new capacity for a volume, the **Resize Volume** window displays the impact that resizing the volume will have on data space and snapshot reserve.

- [Increase the size of a volume](#)
- [Increase the size of a LUN](#)

Also, you can add a LUN to an existing volume. The processes are different when using System Manager with ONTAP 9.8 and later.

- [Add a LUN to an existing volume \(ONTAP 9.8\)](#)
- [Add a LUN to an existing volume \(ONTAP 9.7\)](#)

Increase the size of a volume

Steps

1. Click **Storage > Volumes**.
2. Hover over the name of the volume you want to increase in size.
3. Click .
4. Select **Edit**.
5. Increase the capacity value.
6. Review the **Existing** and **New** data space and snapshot reserve details.

Increase the size of a LUN

Steps

1. Click **Storage > LUNs**.
2. Hover over the name of the LUN you want to increase in size.
3. Click .
4. Select **Edit**.
5. Increase the capacity value.

Add a LUN to an existing volume (ONTAP 9.8)

Beginning with ONTAP 9.8, you can use System Manager to add a LUN to an existing volume that already has a least one LUN.

Steps

1. Click **Storage > LUNs**.
2. Click **Add+**.
3. Complete the fields in the **Add LUNs** window.

4. Select **More Options**.
5. Select the checkbox labeled **Group with related LUNs**.
6. In the drop-down field, select a LUN that exists on the volume to which you want to add another LUN.
7. Complete the rest of the fields. For **Host Mapping**, click one of the radio buttons:
 - **Existing initiator group** lets you select an existing group from a list.
 - **New initiator group** lets you enter a new group in the field.

Add a LUN to an existing volume (ONTAP 9.7)

To use System Manager with ONTAP 9.7 to add a LUN to an existing volume, you should switch to the Classical View first.

Steps

1. Log in to System Manager in ONTAP 9.7.
2. Click **Classical View**.
3. Select **Storage > LUNs > Create**
4. Specify the details to create the LUN.
5. Specify to which existing volume or qtree the LUN should be added.

Save storage space using compression, compaction, and deduplication with ONTAP System Manager

For volumes on non-AFF clusters, you can run deduplication, data compression, and data compaction together or independently to achieve optimal space savings.

- Deduplication eliminates duplicate data blocks.
- Data compression compresses the data blocks to reduce the amount of physical storage that is required.
- Data compaction stores more data in less space to increase storage efficiency.



These tasks are supported for volumes on non-AFF clusters. All inline storage efficiency features, such as inline deduplication and inline compression, are enabled by default on AFF volumes.

Steps

1. Click **Storage > Volumes**.
2. Next to the name of the volume for which you want to save storage, click
3. Click **Edit** and scroll to **Storage Efficiency**.
4. *Optional:* If you want to enable background deduplication, ensure the checkbox is checked.
5. *Optional:* If you want to enable background compression, specify the storage efficiency policy and ensure the checkbox is checked.
6. *Optional:* If you want to enable inline compression, ensure the checkbox is checked.


Balance loads by moving LUNs with ONTAP System Manager

You can move a LUN to another volume within the storage VM to balance the load, or you

can move it to a volume with a higher performance service level to improve performance.

Move restrictions

- A LUN cannot be moved to a qtree within the same volume.
- A LUN created from a file using the CLI cannot be moved with System Manager.
- LUNs that are online and serving data cannot be moved.
- LUNs cannot be moved if the allocated space in the destination volume cannot contain the LUN (even if autogrow is enabled on the volume).
- LUNs on SnapLock volumes cannot be moved with System Manager.



Steps

1. Click **Storage > LUNs**.
2. Select the LUN that you want to move and click **Move**.
3. Select an existing volume to which you want to move the LUN. If the volume contains qtrees, select the qtree.



While the Move operation is in progress, the LUN is displayed on both the origin and destination volume.

Balance loads by moving volumes to another tier with ONTAP System Manager

Beginning with ONTAP 9.9.1, you can move volumes based on analysis of active and inactive data storage. In ONTAP 9.8, you can also use System Manager to move a volume to another tier to balance the load.

For more information, see [File System Analytics overview](#).

Steps

1. Click **Storage > Volumes**.
2. Select the volume or volumes that you want to move, and then click **Move**.
3. Select an existing tier (aggregate) to which you want to move the volume or volumes.

Use Ansible Playbooks to add or edit volumes or LUNs with ONTAP System Manager

Beginning with ONTAP 9.9.1, you can use Ansible Playbooks with System Manager when you want to add or edit volumes or LUNs.

This feature lets you use the same configuration multiple times or use the same configuration with slight changes when you add or edit volumes or LUNs.

Enable or disable Ansible Playbooks

You can enable or disable the use of Ansible Playbooks with System Manager.

Steps

1. In System Manager, go to the UI settings in the cluster settings page:

Cluster > Settings

2. Under **UI Settings**, change the slider switch to "Enabled" or "Disabled".

Save a volume configuration to an Ansible Playbook

When you create or modify the configuration of a volume, you can save the configuration as Ansible Playbook files.

Steps

1. Add or Edit the volume:

Volume > Add (or Volume > Edit)

2. Specify or edit the configuration values of the volume.
3. Select **Save to Ansible Playbook** to save the configuration to Ansible Playbook files.

A zip file is downloaded that contains the following files:

- **variable.yaml**: The values you entered or modified to add or edit the volume.
- **volumeAdd.yaml** (or **volumeEdit.yaml**): The test cases that are required to create or modify the values when reading the inputs from the **variable.yaml** file.

Save a LUN configuration to an Ansible Playbook

When you create or modify the configuration of a LUN, you can save the configuration as Ansible Playbook files.

Steps

1. Add or Edit the LUN:

LUN > Add (or LUN > Edit)

2. Specify or edit the configuration values of the LUN.
3. Select **Save to Ansible Playbook** to save the configuration to Ansible Playbook files:

A zip file is downloaded that contains the following files:

- **variable.yaml**: The values you entered or modified to add or edit the LUN.
- **lunAdd.yaml** (or **lunEdit.yaml**): The test cases that are required to create or modify the values when reading the inputs from the **variable.yaml** file.

Download Ansible Playbook files from global search results

You can download Ansible Playbook files when you do a global search.

Steps

1. In the search field, enter "volume" or "LUN" or "Playbook".
2. Find the search result, either "Volume Management (Ansible Playbook)" or "LUN Management (Ansible Playbook)".
3. Click on  to download the Ansible Playbook files.

Work with Ansible Playbook files

Ansible Playbook files can be modified and run to specify configurations for volumes and LUNs.

About this task

You use two files to perform an operation (either an "add" or an "edit"):

If you want to...	Use this variable file...	And use this run file...
Add a volume	volumeAdd-variable.yaml	valueAdd.yaml
Edit a volume	volumeEdit-variable.yaml	volumeEdit.yaml
Add a LUN	lunAdd-variable.yaml	lunAdd.yaml
Edit a LUN	lunEdit-variable.yaml	lunEdit.yaml

Steps

1. Modify the variables file.

The file contains the various values that you use to configure the volume or LUN.

- If you do not change the values, leave them commented.
- If you modify the values, remove the commenting.

2. Run the associated run file.

The run file contains the test cases that are required to create or modify the values when reading the inputs from the variable file.

3. Enter your user login credentials.

Manage storage efficiency policies with ONTAP System Manager

Beginning with ONTAP 9.8, you can use System Manager to enable, disable, add, edit, or delete efficiency policies for storage VMs on FAS systems.



This function is not available on AFF systems.

Steps

1. Select **Storage > Storage VMs**
2. Select the storage VM for which you want to manage efficiency policies.
3. On the **Settings** tab, select  in the **Efficiency Policy** section. The efficiency policies for that storage VM are displayed.

You can perform the following tasks:

- **Enable or disable** an efficiency policy by clicking the toggle button in the Status column.
- **Add** an efficiency policy by clicking on **Add+**.
- **Edit** an efficiency policy by clicking on  to the right of the policy name and selecting **Edit**.
- **Delete** an efficiency policy by clicking on  to the right of the policy name and selecting **Delete**.

Efficiency policies list

- **Auto**

Specifies that deduplication is continuously performed in the background. This policy is set for all newly created volumes and for all upgraded volumes that have not been manually configured for background deduplication. If you change the policy to “default” or any other policy, the “auto” policy is disabled.

If a volume moves from a non-AFF system to an AFF system, the “auto” policy is enabled on the destination node by default. If a volume moves from an AFF node to a non-AFF node, the “auto” policy on the destination node is replaced by the “inline-only” policy by default.

- **Policy**

Specifies the name of an efficiency policy.

- **Status**

Specifies the status of an efficiency policy. The status can be one of the following:

- Enabled

Specifies that the efficiency policy can be assigned to a deduplication operation.

- Disabled

Specifies that the efficiency policy is disabled. You can enable the policy by using the status drop-down menu and assign it later to a deduplication operation.

- **Run By**

Specifies whether the storage efficiency policy is run based on a schedule or based on a threshold value (change log threshold).

- **QoS Policy**

Specifies the QoS type for the storage efficiency policy. The QoS type can be one of the following:

- Background

Specifies that the QoS policy is running in the background, which reduces potential performance impact on the client operations.

- Best-effort

Specifies that the QoS policy is running on a best-effort basis, which enables you to maximize the utilization of system resources.

- **Maximum Runtime**

Specifies the maximum run-time duration of an efficiency policy. If this value is not specified, the efficiency policy is run till the operation is complete.

Details area

The area below the efficiency policy list displays additional information about the selected efficiency policy, including the schedule name and the schedule details for a schedule-based policy, and the threshold value for a threshold-based policy.

Manage resources using quotas with ONTAP System Manager

Beginning with ONTAP 9.7, you can configure and manage usage quotas with System Manager.

If you are using the ONTAP CLI to configure and manage usage quotas, refer to [Logical Storage Management](#).

If you are using legacy OnCommand System Manager for ONTAP 9.7 and earlier releases to configure and manage usage quotas, see the following for your release:

- [ONTAP 9.7 and 9.6 Documentation](#)
- [ONTAP 9.5 Documentation](#)
- [ONTAP 9.4 Documentation](#)
- [ONTAP 9.3 Documentation](#)

Quota overview

Quotas provide a way to restrict or track the disk space and number of files used by a user, group, or qtree. Quotas are applied to a specific volume or qtree.

You can use quotas to track and limit resource usage in volumes and provide notification when resource usage reaches specific levels.

Quotas can be soft or hard. Soft quotas cause ONTAP to send a notification when specified limits are exceeded, and hard quotas prevent a write operation from succeeding when specified limits are exceeded.

Set quotas to limit resource use with ONTAP System Manager

Add quotas to limit the amount of disk space the quota target can use.

You can set a hard limit and a soft limit for a quota.

Hard quotas impose a hard limit on system resources; any operation that would result in exceeding the limit fails. Soft quotas send a warning message when resource usage reaches a certain level, but they do not affect data access operations, so you can take appropriate action before the quota is exceeded.

Steps

1. Click **Storage > Quotas**.
2. Click **Add**.

Clone volumes and LUNs for testing with ONTAP System Manager

You can clone volumes and LUNs to create temporary, writable copies for testing. The clones reflect the current, point-in-time state of the data. You can also use clones to give additional users access to data without giving them access to production data.

Before you begin

The FlexClone license should be [installed](#) on the storage system.

Cloning a volume

Create a clone of a volume, as follows:

Steps

1. Click **Storage > Volumes**.
2. Click  next to the name of the volume you want to clone.
3. Select **Clone** from the list.
4. Specify a name for the clone and complete the other selections.
5. Click **Clone** and verify that the volume clone appears in the list of volumes.

Alternatively, you can clone a volume from the **Overview** that displays when you view volume details.

Cloning a LUN

You can create copies of your LUNs by cloning the LUNs in the active volume. These FlexClone LUNs are readable and writeable copies of the original LUNs in the active volume.

A space-reserved FlexClone LUN requires as much space as the space-reserved parent LUN. If the FlexClone LUN is not space-reserved, you must ensure that the volume has enough space to accommodate changes to the FlexClone LUN.



This procedure applies to FAS, AFF, and ASA systems. If you have an ASA r2 system (ASA A1K, ASA A90, ASA A70, ASA A50, ASA A30, ASA A20, or ASA C30), follow [these steps](#) to clone data. ASA r2 systems provide a simplified ONTAP experience specific to SAN-only customers.

Example 1. Steps

System Manager

1. Click **Storage > LUNs**.
2. Click  next to the name of the LUN you want to clone.
3. Select **Clone** from the list.
4. Specify a name for the clone and complete the other selections.
5. Click **Clone** and verify that the LUN clone appears in the list of LUNs.

Alternatively, you can clone a LUN from the **Overview** that displays when you view LUN details.

When you create a LUN clone, System Manager automatically enables the deletion of the clone when space is needed.

CLI

1. Verify that the LUNs are not mapped to an igrup or are written to before making the clone.
2. Use the `lun show` command to verify that the LUN exists.

```
lun show -vserver vs1
```

Vserver	Path	State	Mapped	Type	Size
vs1	/vol/vol1/lun1	online	unmapped	windows	47.07MB

3. Use the `volume file clone create` command to create the FlexClone LUN.

```
volume file clone create -vserver vs1 -volume vol1 -source-path lun1  
-destination-path/lun1_clone
```

If you need the FlexClone LUN to be available for automatic deletion, you include `-autodelete true`. If you are creating this FlexClone LUN in a volume using semi-thick provisioning, you must enable automatic deletion for all FlexClone LUNs.

4. Use the `lun show` command to verify that you created a LUN.

```
lun show -vserver vs1
```

Vserver	Path	State	Mapped	Type	Size
vs1	/vol/volX/lun1	online	unmapped	windows	47.07MB
vs1	/vol/volX/lun1_clone	online	unmapped	windows	47.07MB

Search, filter, and sort information in ONTAP System Manager

You can search for various actions, objects, and information topics in System Manager. You can also search table data for specific entries.

System Manager provides two types of searching:

- [Global searching](#)

When you enter a search argument in the field at the top of each page, System Manager searches throughout the interface to find matches. You can then sort and filter the results.

Beginning with ONTAP 9.12.1, System Manager also provides search results from the NetApp Support Site to provide links to relevant support information.

- [Table-grid searching](#)

Beginning with ONTAP 9.8, when you enter a search argument in the field at the top of a table grid, System Manager searches only the columns and rows of that table to find matches.

Global searching

At the top of each page in System Manager, you can use a global search field to search various objects and actions in the interface. For example, you can search for different objects by name, pages available in the navigator column (on the left side), various action items, like "Add Volume" or "Add License", and links to external help topics. You can also filter and sort the results.



For better results, perform searching, filtering, and sorting one minute after logging in and five minutes after creating, modifying, or deleting an object.

Getting search results

The search is not case-sensitive. You can enter a variety of text strings to find the page, actions, or information topics you need. Up to 20 results are listed. If more results are found, you can click **Show more** to view all results. The following examples describe typical searches:

Type of search	Sample search string	Sample search results
By object name	vol_	vol_lun_dest on storage VM: svm0 (Volume) /vol/vol...est1/lun on storage VM: svm0 (LUN) svm0:vol_lun_dest1 role: Destination (Relationship)
By location in interface	volume	Add Volume (Action) Protection – Overview (Page) Recover deleted volume (Help)
By actions	add	Add Volume (Action) Network – Overview (Page) Expand volumes and LUNs (Help)

By help content	san	Storage – Overview (Page) SAN overview (Help) Provision SAN storage for databases (Help)
-----------------	-----	--

Global search results from NetApp Support Site

Beginning with ONTAP 9.12.1, for users who are registered with Active IQ Digital Advisor (also known as Digital Advisor), System Manager displays another column of results that provide links to NetApp Support Site information, including System Manager product information.

Search results contain the following information:

- **Title** of the information which is a link to the document in HTML, PDF, EPUB, or other format.
- **Content type**, which identifies whether it is a product documentation topic, a KnowledgeBase article, or another type of information.
- **Summary description** of the content.
- **Created** date when it was first published.
- **Updated** date when it was last updated.

You can perform the following actions:

Action	Result
Click ONTAP System Manager , then enter text in the search field.	The search results include NetApp Support Site information about System Manager.
Click All products , then enter text in the search field.	The search results include NetApp Support Site information for all NetApp products, not only for System Manager.
Click a search result.	The information from the NetApp Support Site is displayed in a separate browser window or tab.
Click See more results .	If there are more than ten results, you can click See more results after the tenth result to view more results. Each time you click See more results , another ten results are displayed, if available.
Copy the link.	The link is copied to the clipboard. You can paste the link in a file or in a browser window.
Click  .	The panel where the results are displayed is pinned so it remains displayed when you work in another panel.
Click  .	The results panel is no longer pinned and is closed.

Filtering search results

You can narrow the results with filters, as shown in the following examples:

Filter	Syntax	Sample search string
By object type	<type>:<objectName>	volume:vol_2
By object size	<type><size-symbol><number><units>	luns<500mb
By broken disks	"broken disk" or "unhealthy disk"	unhealthy disk
By network interface	<IP address>	172.22.108.21

Sorting search results

When you view all the search results, they are sorted alphabetically. You can sort the results by clicking  Filter and selecting how you want to sort the results.

Table-grid searching

Beginning with ONTAP 9.8, whenever System Manager displays information in a table-grid format, a search button appears at the top of the table.

When you click **Search**, a text field appears in which you can enter a search argument. System Manager searches the entire table and displays only the rows that contain text that matches your search argument.

You can use an asterisk (*) as a "wildcard" character as a substitute for characters. For example, searching for `vol*` might provide rows that contain the following:

- `vol_122_D9`
- `vol_lun_dest1`
- `vol2866`
- `volspec1`
- `volum_dest_765`
- `volume`
- `volume_new4`
- `volume9987`

Logical storage management with the CLI

Logical storage management overview with the CLI

Using the ONTAP CLI, you can create and manage FlexVol volumes, use FlexClone technology to create efficient copies of volumes, files, and LUNs, create qtrees and quotas, and manage efficiency features like deduplication and compression.

You should use these procedures under the following circumstances:

- You want to understand the range of ONTAP FlexVol volume capabilities and storage efficiency features.
- You want to use the command-line interface (CLI), not System Manager or an automated scripting tool.

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 with 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](#). Learn more about `volume create` in the [ONTAP command reference](#).

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
```

Learn more about `volume show` in the [ONTAP command reference](#).

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 in ONTAP

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, maximum [FlexGroup volume](#) size of 60PB, 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.
- Learn more about the commands described in this procedure in the [ONTAP command reference](#).

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_voll -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
```

2. Activate the new configuration settings by remounting the volume:

```
volume unmount -vserver <svm_name> -volume <volume_name>
```

```
volume mount -vserver <svm_name> -volume <volume_name>
```

Related information

- [Create an ONTAP NFS volume](#)
- [ONTAP 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 (snapshots 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 snapshots, 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](#) has more information to enable/disable support for the T10 thin provisioning and to enable T10 SCSI thin provisioning on a LUN.

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 snapshots 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
```

Related information

- [volume create](#)
- [volume modify](#)

Determine space usage in a volume or aggregate in ONTAP

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 snapshots.

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

Learn more about `volume show-space` in the [ONTAP command reference](#).

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.

Learn more about `volume show-footprint` in the [ONTAP command reference](#).

- The aggregate's space usage

This perspective includes totals of the volume footprints of all volumes contained in the aggregate, space reserved for aggregate snapshots, 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 with 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.

Learn more about `storage aggregate show-space` in the [ONTAP command reference](#).

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.

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.

Related Information

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

Enable automatic snapshot and LUN deletion to manage space

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

About this task

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

Step

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

Learn more about `volume snapshot autodelete modify` and defining a policy that meets your needs in the [ONTAP command reference](#).

The following command enables the automatic deletion of snapshots 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 snapshots 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 snapshots work differently than volume-level snapshots and are managed automatically by ONTAP. The option to delete aggregate snapshots is always enabled and helps in managing space utilization.



If the trigger parameter is set to `snap_reserve` for an aggregate, the snapshots 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 in the command will be listed as 0 because these snapshots are automatically deleted. Also, the space used by snapshots 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 creation. If you attempt to create a snapshot and there is insufficient space, the snapshot creation fails, even with autogrow enabled.

- Delete Snapshots, FlexClone files, or FlexClone LUNs.

For example, you can configure ONTAP to automatically delete snapshots that are not linked to snapshots in cloned volumes or LUNs, or you can define which snapshots you want ONTAP to delete first—your oldest or newest snapshots. You can also determine when ONTAP should begin deleting snapshots—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 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 snapshots 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. Learn more about `volume autosize` in the [ONTAP command reference](#).

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 snapshots, 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. Learn more about `volume modify` in the [ONTAP command reference](#).

To specify increasing the size of the volume first (the default), use `volume_grow`. To specify deleting snapshots 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 deletion

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

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

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

Autoshrink functionality and snapshot deletion

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

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

- If both the `grow_shrink` autosize mode and automatic snapshot deletion are enabled, when a volume size shrinks it can trigger an automatic snapshot 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 snapshots to spill out of the reserve and be deleted automatically.

- If the `grow_shrink` autosize mode is enabled and you manually delete a snapshot, 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 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.

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 Used 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 Used 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.

Alert type	EMS level	Configurable?	Definition	Ways to address	Risk if no action taken
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 snapshots • Reducing user data • Disabling space reservation for files or LUNs 	Write operations to unreserved files could fail.
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 snapshots • 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.

Aler t type	EM S Lev el	Con figu rable?	Definition	Ways to address	Risk if no action taken
Nearly full	Deb ug	N	The amount of space allocated for volumes, including their guarantees, has exceeded the threshold set for this alert (95%). The percentage is the Used 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	Deb ug	N	The file system has exceeded the threshold set for this alert (98%). The percentage is the Used 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.
Log ically over all ocat ed	SV C Err or	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 over all ocat ed	Node Err or	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 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 automatic deletion could fall behind, resulting in the volume running out of space.
- If all the above requirements are met and snapshots 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 snapshots 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 snapshots 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-used
```

Example

```
cluster1::> volume show -vserver vs1 -volume vol1 -fields files-used
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 specify an optional minimum throughput limit.

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.

- Beginning with ONTAP 9.18.1, you can assign QoS policies to volumes contained in SVMs that have QoS policies. When you use nested QoS policies, the most restrictive policy is applied.
- Beginning with ONTAP 9.14.0, you can assign policies to qtrees contained in volumes that have QoS policies.

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.

Related information

- [qos policy-group](#)
- [xref:/volumes/ qos policy-group create](#)
- [volume create](#)
- [volume modify](#)
- [qos statistics](#)

Delete a FlexVol volume

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

Before you begin

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



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

Steps

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

```
volume unmount -vserver vserver_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 vserver_name volume_name
```

4. Delete the volume:

```
volume delete -vserver vserver_name volume_name
```

Result

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

Related information

- [snapmirror delete](#)
- [volume unmount](#)
- [volume offline](#)
- [volume delete](#)

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.



Deleting the SVM that contains the deleted volume clears the Volume Recovery Queue (VRQ). Only delete an SVM when you are certain that there is no need to recover volumes owned by the SVM. Volumes in the volume recovery queue cannot exist when the owning SVM is deleted.

Related information

- [How to use the Volume Recovery Queue](#)
- [volume delete](#)
- [volume show](#)

Commands for managing FlexVol volumes in ONTAP

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>

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

Learn more about `volume` in the [ONTAP command reference](#).

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 snapshots.

Beginning with ONTAP 9.18.1, the `storage aggregate show-space` command changes how Logical Referenced Capacity and Logical Unreferenced Capacity is reported. Logical Referenced Capacity reports referenced blocks in all objects and unreferenced blocks in fragmented objects. Logical Unreferenced Capacity reports only unreferenced blocks in objects that have crossed the fullness threshold and are eligible for object deletion and defragmentation.

For example, when you use the default aggregate fullness threshold of 40% for ONTAP S3 and StorageGRID, 60% of the blocks in an object must be unreferenced before the blocks are reported as unreferenced capacity.

In releases earlier than ONTAP 9.18.1, Logical Referenced Capacity reports referenced blocks in all objects (both full and fragmented objects). Logical Unreferenced Capacity reports unreferenced blocks in all objects.

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</code> <code>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	<code>volume snapshot compute-reclaimable (advanced)</code>

To display information about...	Use this command...
The amount of space used by a volume	<pre>volume show -fields size,used,available,percent-used</pre> <pre>volume show-space</pre>
The amount of space used by a volume in the containing aggregate	<pre>volume show-footprint</pre>

Related information

- [storage aggregate show](#)
- [storage aggregate show-space](#)
- [storage aggregate show-status](#)
- [volume snapshot compute-reclaimable](#)
- [volume show](#)

Move and copy volumes

Move a FlexVol volume overview

You can move or copy volumes for capacity utilization, improved performance, and to satisfy service-level agreements. Knowing how moving a FlexVol volume works helps you to determine whether the volume move satisfies service-level agreements and to understand where a volume move is in the volume move process.

FlexVol volumes are moved from one aggregate or node to another within the same storage virtual machine (SVM). A volume move does not disrupt client access during the move.



During the cutover phase of a volume move operation, you cannot create FlexClone files or FlexClone LUNs of a FlexVol volume.

Moving a volume occurs in multiple phases:

- A new volume is made on the destination aggregate.
- The data from the original volume is copied to the new volume.

During this time, the original volume is intact and available for clients to access.

- At the end of the move process, client access is temporarily blocked.

During this time the system performs a final replication from the source volume to the destination volume, swaps the identities of the source and destination volumes, and changes the destination volume to the source volume.

- After completing the move, the system routes client traffic to the new source volume and resumes client access.

The move is not disruptive to client access because the time in which client access is blocked ends before

clients notice a disruption and time out. Client access is blocked for 30 seconds by default. If the volume move operation cannot finish in the time that access is denied, the system aborts this final phase of the volume move operation and allows client access. The system attempts the final phase three times by default. After the third attempt, the system waits an hour before attempting the final phase sequence again. The system runs the final phase of the volume move operation until the volume move is complete.

Considerations and recommendations when moving volumes

There are several considerations and recommendations to be aware of when moving a volume. These are based on the volume you are moving as well as the system configuration such as MetroCluster. You should understand all the relevant issues before moving a volume.

General considerations and recommendations

- If you're upgrading the release family for a cluster, don't move a volume until after you upgrade all of the nodes in the cluster.

This recommendation prevents you from inadvertently attempting to move a volume from a newer release family to an older release family.

- The source volume must be consistent.
- If you have assigned one or more aggregates to the associated storage virtual machine (SVM), the destination aggregate must be one of the assigned aggregates.
- You should only move a volume to a later ONTAP version.
- You cannot move a volume to or from a taken-over CFO aggregate.
- If a volume that contains LUNs isn't NVFAIL enabled before you move it, the volume will be NVFAIL enabled after you move it.
- You can move a volume from a Flash Pool aggregate to another Flash Pool aggregate.
 - The caching policies of that volume are also moved.
 - The move might affect volume performance.
- You can move volumes between a Flash Pool aggregate and a non-Flash Pool aggregate.
 - If you move a volume from a Flash Pool aggregate to a non-Flash Pool aggregate, ONTAP displays a message warning you that the move might affect volume performance and asks whether you want to continue.
 - If you move a volume from a non-Flash Pool aggregate to a Flash Pool aggregate, ONTAP assigns the `auto` caching policy.
- Volumes have the data-at-rest protections of the aggregate they reside on. If you move a volume from an aggregate that consists of NSE drives to one that does not, the volume no longer has NSE data-at-rest protection.
- If you're moving FabricPool optimized volumes from ONTAP 9.13.1 or earlier to ONTAP 9.15.1 or later, see the [NetApp Knowledge Base: CONTAP-307878 - Unexpected reboot during FabricPool optimized volume move if the source ONTAP is less than 9.14.1 and destination is greater than 9.14.1](#).
- Beginning with ONTAP 9.15.1, moving volumes from an A400 system to an A70, A90, or A1K system may cause increased read latency issues. For details and recommended actions, see the [NetApp Knowledge Base: CONTAP-556247 - Slow Compression / Decompression on volumes after being moved from A400 to A70, A90 and A1K](#).

FlexClone volume considerations and recommendations

- FlexClone volumes cannot be offline when they are being moved.
- You can move FlexClone volumes from one aggregate to another aggregate on the same node or another node in the same SVM without initiating the `vol clone split start` command.

By initiating a volume move operation on a FlexClone volume, the clone volume is split during the move process to a different aggregate. After the volume move on the clone volume is complete, the volume that moved no longer appears as a clone, but appears instead as an independent volume without any clone relationship with the previous parent volume.

- FlexClone volume snapshots aren't lost after moving a clone.
- You can move FlexClone parent volumes from one aggregate to another aggregate.

When you move a FlexClone parent volume, a temporary volume is left behind that acts as a parent volume for all FlexClone volumes. No operations are allowed on the temporary volume except to take it offline or to delete it. After all FlexClone volumes are either split or destroyed, the temporary volume is cleaned up automatically.

- After you move a FlexClone child volume, the volume is no longer a FlexClone volume.
- FlexClone move operations are mutually exclusive from FlexClone copy or split operations.
- If a clone-splitting operation is in progress, moving a volume might fail.

You should not move a volume until clone-splitting operations are completed.

MetroCluster considerations and recommendations

- During a volume move in a MetroCluster configuration, when a temporary volume is created on the destination aggregate on the source cluster a record of the temporary volume corresponding to the volume in the mirrored, but unassimilated, aggregate is also created on the surviving cluster.
- If a MetroCluster switchover occurs before the cutover, the destination volume has a record and is a temporary volume (a volume of type TMP).

Move job restarts on the surviving (disaster recovery) cluster, reports a failure, and cleans up all move-related items including the temporary volume. In any event where cleanup cannot be done correctly, an EMS is generated alerting the system administrator to do the necessary cleanup.

- If a MetroCluster switchover occurs after the cutover phase has started but before the move job has completed (that is, the move reached a stage where it can update the cluster to point to the destination aggregate), the move job restarts on the surviving (disaster recovery) cluster and runs to completion.

All move-related items are cleaned up including the temporary volume (original source). In any event where cleanup cannot be done correctly, an EMS is generated alerting the system administrator to do the necessary cleanup.

- Neither forced nor unforced MetroCluster switchbacks are allowed if there are any volume move operations in progress for volumes belonging to the switched over site.

Switchbacks aren't blocked when volume move operations are in progress for volumes local to the surviving site.

- Unforced MetroCluster switchovers are blocked, but forced MetroCluster switchovers aren't blocked if there are any volume move operations in progress.

Requirements for moving volumes in a SAN environment

You need to prepare before moving a volume in a SAN environment.

Before moving a volume containing LUNs or namespaces, you must meet the following requirements:

- For volumes containing one or more LUNs, you should have a minimum of two paths per LUN (LIFs) connecting to each node in the cluster.

This eliminates single points of failure and enables the system to survive component failures.

- For volumes containing namespaces, the cluster must be running ONTAP 9.6 or later.

Volume move is not supported for NVMe configurations running ONTAP 9.5.

Move an ONTAP volume

You can move a FlexVol volume to a different aggregate, node, or both within the same storage virtual machine (SVM) to balance storage capacity after determining that there is a storage capacity imbalance.

About this task

By default, if the cutover operation fails to complete within 30 seconds, it will retry. You can adjust the default behavior by using the `-cutover-window` and `-cutover-action` parameters, both of which require advanced privilege level access.

You must be a cluster administrator to perform this task.

Before you begin

- If you are moving a volume that uses 8K adaptive compression to one of the following platforms, you should [increase the size of the volume's active file system](#) before moving the volume. Data is compressed differently on these platforms so that space is saved at the aggregate level instead of the volume level. Because of this difference, the size of the volume's active file system should be increased by the amount of 8k compression savings to prevent the volume from running out of space during the volume move.
 - AFF and FAS platforms that support dedicated offload processor storage efficiency

Learn more about AFF and FAS platforms that support [dedicated offload processor storage efficiency](#).

- AFF C-Series platforms

See the [Hardware Universe](#) for a full list of C-series platforms.

- If you are moving a data protection mirror and you have not initialized the mirror relationship, use the `snapmirror initialize` command to initialize the mirror relationship. Learn more about `snapmirror initialize` in the [ONTAP command reference](#).

Data protection mirror relationships must be initialized before you can move one of the volumes.

Steps

1. Determine an aggregate to which you can move the volume:

```
volume move target-aggr show
```

The aggregate that you select must have enough space for the volume; that is, the available size is bigger than the volume that you are moving.

The following example shows that the vs2 volume can be moved to any of the listed aggregates:

```
cluster1::> volume move target-aggr show -vserver vs2 -volume user_max
Aggregate Name      Available Size    Storage Type
-----  -----
aggr2              467.9GB        hdd
node12a_aggr3      10.34GB       hdd
node12a_aggr2      10.36GB       hdd
node12a_aggr1      10.36GB       hdd
node12a_aggr4      10.36GB       hdd
5 entries were displayed.
```

Learn more about `volume move target-aggr show` in the [ONTAP command reference](#).

2. Perform a validation check to verify that the volume can be moved to the intended aggregate:

```
volume move start -perform-validation-only
```

Learn more about `volume move start` in the [ONTAP command reference](#).

3. Move the volume:

```
volume move start
```

The following command moves the `user_max` volume on the `vs2` SVM to the `node12a_aggr3` aggregate. The move runs as a background process.

```
cluster1::> volume move start -vserver vs2 -volume user_max
-destination-aggregate node12a_aggr3
```

4. Determine the status of the volume move operation:

```
volume move show
```

The following example shows the state of a volume move that completed the replication phase and is in the cutover phase:

```
cluster1::> volume move show
Vserver    Volume      State      Move Phase  Percent-Complete Time-To-
Complete
-----
-----
vs2        user_max    healthy    cutover      -          -
```

The volume move is complete when it no longer appears in the `volume move show` command output.

Learn more about `volume move show` in the [ONTAP command reference](#).

5. Optionally, view compression savings:

```
volume show-footprint -vserver <SVM> -volume <volume_name>
```



Additional aggregate-level savings might be realized through a post-process conversion scan that runs automatically and shortly after the volume move is completed.

Related information

- [Considerations and recommendations when moving volumes](#)

Increase an ONTAP volume's active file system before migrating from 8k adaptive compression

Platforms that support 8k adaptive compression save space at the volume level. AFF C-Series platforms and platforms that support 32k compression save space at the aggregate level. When migrating a volume from 8k adaptive compression to an AFF C-Series platform or to a platform with 32k compression, you need to increase the size of the volume's active file system by the 8k compression savings. This prevents the volume from running out of free space during the volume move.

The following systems support 32k compression:

Unresolved directive in volumes/increase-volume-active-file-system-size.adoc - include::_include/dedicated-offload-processor-supported-platforms.adoc[]

Learn more about [AFF and FAS platforms that support 32k compression](#).

See the [Hardware Universe](#) for a full list of AFF C-series platforms.

About this task

Perform these steps if you are migrating your data using a volume move operation. If you are migrating your data using a SnapMirror operation, you do not need to manually increase the size of the active file system. SnapMirror destination volumes use volume autosize by default and therefore are not expected to run out of space due to compression savings being realized at the aggregate layer instead of the volume layer.

Before you begin

If logical space reporting and enforcement is not enabled on your volume, you can optionally enable it by

setting the `-is-space-reporting-logical` and `-is-space-enforcement-logical` parameters to **true**. Enabling these settings before the volume move can help you assess if your volume is large enough to accommodate the compression savings loss at the volume layer when you convert from 8k compression. You should enable these settings on the volume. If you enable these settings at the SVM level, they are applied to newly created volumes only.

Steps

1. Verify the volume's current size and snapshot reserve:

```
volume show-space
```

2. Check the volume's compression space savings:

```
volume show -vserver -volume -fields compression-space-saved
```

3. Increase the volume's active filesystem size by the amount shown for `compression-space-saved` plus the snapshot reserve.

```
volume size -vserver <vserver_name> -volume <volume_name> -new-size +<size>
```

Example

IF a volume is 100GB and has a 20% snapshot reserve; then the the active filesystem is 80GB and the snapshot reserve is 20GB. To increase the active filesystem by 20GB, you must add 25GB to the overall volume size; that is, 20GB for the active filesystem and 5GB (20%) for the snapshot reserve.

```
volume size -vserver svm1 -volume volx -size +20GB
```

4. Verify that the size of the volume is increase:

```
volume show -vserver <vserver_name> -volume <volume_name> -fields size
```

Result

Your volume's active file system size is increased and you are ready to move the volume.

What's next?

Perform a [volume move](#) to migrate your data.

Commands for moving volumes in ONTAP

The ONTAP CLI provides specific commands for managing volume movement. Depending on what you need to do, use the following commands to manage quota rules and quota policies:

If you want to...	Use this command...
Abort an active volume move operation.	volume move abort
Show status of a volume moving from one aggregate to another aggregate.	volume move show
Start moving a volume from one aggregate to another aggregate.	volume move start
Manage target aggregates for volume move.	volume move target-aggr
Trigger cutover of a move job.	volume move trigger-cutover
Change the amount of time client access is blocked if the default is not adequate.	volume move start or volume move modify with the -cutover-window parameter. The volume move modify command is an advanced command and the -cutover-window is an advanced parameter.
Determine what the system does if the volume move operation cannot be completed during the time client access is blocked.	volume move start or volume move modify with the -cutover-action parameter. The volume move modify command is an advanced command and the -cutover-action is an advanced parameter.

Related information

- [volume move](#)

Methods for copying a volume

The method you use for copying a volume depends on whether you are copying it to the same aggregate or a different aggregate, and whether you want to retain snapshots from the original volume.

Copying a volume creates a standalone copy of a volume that you can use for testing and other purposes.

The following table lists characteristics of the copy and the methods used to create that copy.

If you want to copy a volume...	Then the method you use is...
Within the same aggregate and you do not want to copy snapshots from the original volume.	Creating a FlexClone volume of the original volume.
To another aggregate and you do not want to copy snapshots from the original volume.	Creating a FlexClone volume of the original volume, and then moving the volume to another aggregate by using the volume move command.

If you want to copy a volume...	Then the method you use is...
To another aggregate and preserve all of the snapshots from the original volume.	Replicating the original volume using SnapMirror, and then breaking the SnapMirror relationship to make a read-write volume copy.

Use FlexClone volumes to create efficient copies of your FlexVol volumes

FlexClone volume use overview

FlexClone volumes are writable, point-in-time copies of a parent FlexVol volume.

FlexClone volumes are space-efficient because they share the same data blocks with their parent FlexVol volumes for common data. The snapshot used to create a FlexClone volume is also shared with the parent volume.

You can clone an existing FlexClone volume to create another FlexClone volume. You can also create a clone of a FlexVol volume containing LUNs and LUN clones.

You can also split a FlexClone volume from its parent volume. Beginning with ONTAP 9.4, for non-guaranteed volumes on AFF systems, the split operation for FlexClone volumes shares the physical blocks and does not copy the data. Therefore, splitting of FlexClone volumes on AFF systems is faster than the FlexClone splitting operation in other FAS systems in ONTAP 9.4 and later releases.

You can create two types of FlexClone volumes: read-write FlexClone volumes and data protection FlexClone volumes. While you can create a read-write FlexClone volume of a regular FlexVol volume, you must use only a SnapVault secondary volume to create a data protection FlexClone volume.

Create a FlexClone volume

You can create a data protection FlexClone volume from a SnapMirror destination volume or from a parent FlexVol volume that is a SnapVault secondary volume. Beginning with ONTAP 9.7, you can create a FlexClone volume from a FlexGroup volume. After you create a FlexClone volume, you cannot delete the parent volume while the FlexClone volume exists.

Before you begin

- The FlexClone license must be installed on the cluster. This license is included with [ONTAP One](#).
- The volume that you want to clone must be online.



Cloning a volume as a FlexClone volume on a different SVM is not supported on MetroCluster configurations.

Create a FlexClone volume of a FlexVol or FlexGroup

Step

- Create a FlexClone volume:

```
volume clone create
```



While creating a read-write FlexClone volume from the read-write parent volume, you do not need to specify the base snapshot. ONTAP creates a snapshot if you do not name any specific snapshot that is to be used as the base snapshot for the clone. You must specify the base snapshot for creating a FlexClone volume when the parent volume is a data protection volume.

Example

- The following command creates a read-write FlexClone volume `vol1_clone` from the parent volume `vol1`:

```
volume clone create -vserver vs0 -flexclone vol1_clone -type RW -parent-volume vol1
```

- The following command creates a data protection FlexClone volume `vol_dp_clone` from the parent volume `dp_vol` by using the base snapshot `snap1`:

```
volume clone create -vserver vs1 -flexclone vol_dp_clone -type DP -parent-volume dp_vol -parent-snapshot snap1
```

Create a FlexClone of any SnapLock type

Beginning with ONTAP 9.13.1, you can specify one of three SnapLock types, compliance, enterprise, non-snaplock, when creating a FlexClone of a RW volume. By default, a FlexClone volume is created with the same SnapLock type as the parent volume. However, you can override the default by using the `snaplock-type` option during FlexClone volume creation.

Using the `non-snaplock` parameter with the `snaplock-type` option, you can create a non-SnapLock type FlexClone volume from a SnapLock parent volume to provide a faster method of bringing data back online when necessary.

Learn more about [SnapLock](#).

Before you begin

You should be aware of the following FlexClone volume limitations when they have a different SnapLock type than the parent volume.

- Only RW-type clones are supported. DP-type clones with a SnapLock type different from the parent volume are not supported.
- Volumes with LUNs cannot be cloned using the `snaplock-type` option set to a value other than 'non-snaplock' because SnapLock volumes do not support LUNs.
- A volume on a MetroCluster mirrored aggregate cannot be cloned with a Compliance SnapLock type because SnapLock Compliance volumes are not supported on MetroCluster mirrored aggregates.
- SnapLock Compliance volumes with Legal-Hold cannot be cloned with a different SnapLock type. Legal-Hold is only supported on SnapLock Compliance volumes.
- SVM DR does not support SnapLock volumes. Attempting to create a SnapLock clone from a volume in an SVM that is part of an SVM DR relationship will fail.
- FabricPool best practices recommend that clones retain the same tiering policy as the parent. However, a SnapLock Compliance clone of a FabricPool-enabled volume cannot have the same tiering policy as the parent. The tiering policy must be set to none. Attempting to create a SnapLock Compliance clone from a parent with a tiering policy other than `none` will fail.

Steps

1. Create a FlexClone volume with a SnapLock type: `volume clone create -vserver svm_name -flexclone flexclone_name -type RW [-snaplock-type {non-snaplock|compliance|enterprise}]`

Example:

```
> volume clone create -vserver vs0 -flexclone vol1_clone -type RW  
-snaplock-type enterprise -parent-volume vol1
```

Split a FlexClone volume from its parent volume

You can split a FlexClone volume from its parent to make the clone a normal FlexVol volume.

The clone splitting operation takes place in the background. Data is accessible on the clone and the parent during the split. Beginning with ONTAP 9.4, space efficiency is preserved. The split process only updates metadata and requires minimal IO. No data blocks are copied.

About this task

- New snapshots of the FlexClone volume cannot be created during the split operation.
- A FlexClone volume cannot be split from the parent volume if it belongs to a data protection relationship or is part of a load-sharing mirror.
- If you take the FlexClone volume offline while splitting is in progress, the split operation is suspended; when you bring the FlexClone volume back online, the splitting operation resumes.
- After the split, both the parent FlexVol volume and the clone require the full space allocation determined by their volume guarantees.
- After a FlexClone volume is split from its parent the two cannot be rejoined.
- Beginning with ONTAP 9.4, for non-guaranteed volumes on AFF systems, the split operation for FlexClone volumes shares the physical blocks and does not copy the data. Therefore, splitting of FlexClone volumes on AFF systems is faster than the FlexClone splitting operation in other FAS systems in ONTAP 9.4 and later. The improved FlexClone splitting operation on AFF systems has the following benefits:
 - Storage efficiency is preserved after splitting the clone from the parent.
 - Existing snapshots are not deleted.
 - The operation is faster.
 - The FlexClone volume can be split from any point in the clone hierarchy.

Before you begin

- You must be a cluster administrator.
- The FlexClone volume must be online when the split operation begins.
- The parent volume must be online for the split to succeed.

Steps

1. Determine the amount of free space required to complete the split operation:

```
volume clone show -estimate -vserver vserver_name -flexclone clone_volume_name  
-parent-volume parent_vol_name
```

The following example provides information about the free space required to split FlexClone volume “clone1” from its parent volume “vol1”:

2. Verify that the aggregate containing the FlexClone volume and its parent has sufficient space:

a. Determine the amount of free space in the aggregate that contains the FlexClone volume and its parent:

storage aggregate show

b. If the containing aggregate does not have enough free space available, add storage to the aggregate;

```
storage aggregate add-disks
```

3. Start the split operation:

```
volume clone split start -vserver vserver name -flexclone clone volume name
```

The following example shows how you can initiate the process to split FlexClone volume "clone1" from its parent volume "vol1":

```
cluster1::> volume clone split start -vserver vs1 -flexclone clone1

Warning: Are you sure you want to split clone volume clone1 in Vserver
vs1 ?
{y|n}: y
[Job 1617] Job is queued: Split clone1.
```

4. Monitor the status of the FlexClone split operation:

```
volume clone split show -vserver vserver name -flexclone clone volume name
```

The following example shows the status of the FlexClone split operation on an AFF system:

```

cluster1::> volume clone split show -vserver vs1 -flexclone clone1
                                         Inodes
Blocks
-----
-----
Vserver      FlexClone      Processed Total      Scanned   Updated      % Inode
% Block
Complete    Complete
vs1         clone1        0          0        411247    153600      0
37

```

5. Verify that the split volume is no longer a FlexClone volume:

```
volume show -volume volume_name -fields clone-volume
```

The value of the `clone-volume` option is “false” for a volume that is not a FlexClone volume.

The following example shows how you can verify whether volume “clone1” that is split from its parent is not a FlexClone volume.

```

cluster1::> volume show -volume clone1 -fields clone-volume
vserver volume **clone-volume**
----- *-----**
vs1      clone1 **false**

```

Related information

- [storage aggregate add-disks](#)

Determine the space used by a FlexClone volume

You can determine the space used by a FlexClone volume based on its nominal size and the amount of space it shares with the parent FlexVol volume. When a FlexClone volume is created, it shares all of its data with its parent volume. Although the nominal size of the FlexVol volume is the same as its parent’s size, it uses very little free space from the aggregate.

About this task

The free space used by a newly-created FlexClone volume is approximately 0.5 percent of its nominal size. This space is used to store the FlexClone volume’s metadata.

New data written to either the parent or the FlexClone volume is not shared between the volumes. The increase in the amount of new data that gets written to the FlexClone volume leads to an increase in the space the FlexClone volume requires from its containing aggregate.

Step

1. Determine the actual physical space used by the FlexClone volume using the `volume show` command.

The following example shows the total physical space used by the FlexClone volume:

```
cluster1::> volume show -vserver vs01 -volume clone_voll -fields
size,used,available,
percent-used,physical-used,physical-used-percent
vserver      volume      size  available  used  percent-used  physical-
used      physical-used-percent
-----  -----  -----  -----  -----  -----  -----
-----  -----
vs01      clone_voll    20MB    18.45MB   564KB      7%      196KB
1%
```

Learn more about `volume show` in the [ONTAP command reference](#).

Considerations for creating a FlexClone volume from a SnapMirror source or destination volume

You can create a FlexClone volume from the source or destination volume in an existing volume SnapMirror relationship. However, doing so could prevent future SnapMirror replication operations from completing successfully.

Replication might not work because when you create the FlexClone volume, you might lock a snapshot that is used by SnapMirror. If this happens, SnapMirror stops replicating to the destination volume until the FlexClone volume is destroyed or is split from its parent. You have two options for addressing this issue:

- If you require the FlexClone volume on a temporary basis and can accommodate a temporary stoppage of the SnapMirror replication, you can create the FlexClone volume and either delete it or split it from its parent when possible.

The SnapMirror replication continues normally when the FlexClone volume is deleted or is split from its parent.

- If a temporary stoppage of the SnapMirror replication is not acceptable, you can create a snapshot in the SnapMirror source volume, and then use that snapshot to create the FlexClone volume. (If you are creating the FlexClone volume from the destination volume, you must wait until that snapshot replicates to the SnapMirror destination volume.)

This method of creating a snapshot in the SnapMirror source volume allows you to create the clone without locking a snapshot that is in use by SnapMirror.

Use FlexClone files and FlexClone LUNs to create efficient copies of files and LUNs

FlexClone file and FlexClone LUN use overview

FlexClone files and FlexClone LUNs are writable, space-efficient clones of parent files and parent LUNs, and help in efficient utilization of the physical aggregate space. FlexClone files and FlexClone LUNs are supported only for FlexVol volumes.

FlexClone files and FlexClone LUNs use 0.4 percent of their size to store the metadata. Clones share the data blocks of their parent files and parent LUNs and occupy negligible storage space until clients write new data either to the parent file or LUN, or to the clone.

Clients can perform all file and LUN operations on both the parent and the clone entities.

You can use multiple methods to delete FlexClone files and FlexClone LUNs.

Create a FlexClone file or FlexClone LUN in ONTAP

You can create space-efficient and time-efficient clones of files and LUNs present in FlexVol volumes or FlexClone volumes by using the `volume file clone create` command.

Before you begin

- The FlexClone license must be installed on the cluster. This license is included with [ONTAP One](#).
- If multiple block ranges are used for sub-LUN cloning or sub-file cloning, the block numbers must not overlap.
- If you are creating a sub-LUN or sub-file on volumes with adaptive compression enabled, the block ranges must not be misaligned.

This means that the source start block number and destination start block number must either be even aligned or odd aligned.

About this task

Depending on the privileges assigned by the cluster administrator, an SVM administrator can create FlexClone files and FlexClone LUNs.

You can specify the autodelete setting for FlexClone files and FlexClone LUNs when you create and modify clones. By default, the autodelete setting is disabled.

You can overwrite an existing FlexClone file or FlexClone LUN when you create a clone by using the `volume file clone create` command with the `-overwrite-destination` parameter.

When the node reaches its maximum split load, the node temporarily stops accepting requests to create FlexClone files and FlexClone LUNs and issues an `EBUSY` error message. When the split load for the node falls below the maximum, the node accepts requests to create FlexClone files and FlexClone LUNs again. You should wait until the node has capacity to create the clones before trying the create request again.

The FlexClone LUN inherits the space reservations attribute of the parent LUN. A space-reserved FlexClone LUN requires as much space as the space-reserved parent LUN. If the FlexClone LUN is not space-reserved, the volume must have enough space to accommodate changes to the clone.

Steps

1. If you are cloning a LUN, verify that the LUN is not mapped or being written to.
2. Create the FlexClone LUN or file:

```
volume file clone create -vserver vserver_name -volume volume_name -source
-path source_path -destination-path destination_path
```

The following example shows how you can create a FlexClone file `file1_clone` of the parent file `file1_source` in the volume `vol1`:

```
cluster1::> volume file clone create -vserver vs0 -volume vol1 -source  
-path /file1_source -destination-path /file1_clone
```

Learn more about `volume file clone create` in the [ONTAP command reference](#).

Create FlexClone LUNs from a snapshot in a volume

You can use a snapshot in your volume to create FlexClone copies of your LUNs. FlexClone copies of LUNs are both readable and writeable.

Before you begin

A FlexClone license must be installed. This license is included with [ONTAP One](#).

About this task

The FlexClone LUN inherits the space reservations attribute of the parent LUN. A space-reserved FlexClone LUN requires as much space as the space-reserved parent LUN. If the FlexClone LUN is not space-reserved, the volume must have enough space to accommodate changes to the clone.

Steps

1. Verify that the LUN is not mapped or being written to.
2. Create a snapshot of the volume that contains the LUNs:

```
volume snapshot create -vserver vserver_name -volume volume_name -snapshot  
snapshot_name
```

You must create a snapshot (the backing snapshot) of the LUN you want to clone.

3. Create the FlexClone LUN from the snapshot:

```
volume file clone create -vserver vserver_name -volume volume_name -source  
-path source_path -snapshot-name snapshot_name -destination-path  
destination_path
```

If you need the FlexClone LUN to be available for automatic deletion, you include `-autodelete true`. If you are creating this FlexClone LUN in a volume using semi-thick provisioning, you must enable automatic deletion for all FlexClone LUNs.

4. Verify that the FlexClone LUN is correct:

```
lun show -vserver vserver_name
```

Vserver	Path	State	Mapped	Type	Size
vs1	/vol/vol1/lun1_clone	online	unmapped	windows	47.07MB
vs1	/vol/vol1/lun1_snap_clone	online	unmapped	windows	47.07MB

View node capacity before creating and deleting FlexClone files and FlexClone LUNs

You should determine whether a node has capacity to receive requests to create and delete FlexClone files and FlexClone LUNs. This can be done by viewing the split load for the node. If the maximum split load is reached, no new requests are accepted until the split load falls below the maximum.

About this task

When the node reaches its maximum split load, an `EBUSY` error message is issued in response to create and delete requests. When the split load for the node falls below the maximum, the node accepts requests to create and delete FlexClone files and FlexClone LUNs again.

A node can accept new requests when the `Allowable Split Load` field displays capacity, and the create request fits in the available capacity.

Steps

1. View how much capacity a node has to create and delete FlexClone files and FlexClone LUNs by using the `volume file clone split load show` command.

In the following example, the split load is displayed for all of the nodes in cluster1. All nodes in the cluster have capacity to create and delete FlexClone files and FlexClone LUNs as indicated by the `Allowable Split Load` field:

```
cluster1::> volume file clone split load show
  Node      Max          Current      Token      Allowable
          Split Load  Split Load Reserved Load Split Load
  -----
  node1      15.97TB      0B          100MB      15.97TB
  node2      15.97TB      0B          100MB      15.97TB
  2 entries were displayed.
```

Related information

- [volume file clone split load show](#)

View space savings with FlexClone files and FlexClone LUNs

You can view the percentage of disk space saved by block sharing on a volume containing FlexClone files and FlexClone LUNs. You might do this as part of capacity planning.

Steps

1. To view the space saving achieved due to FlexClone files and FlexClone LUNs, enter the following command:

```
df -s volname
```

`volname` is the name of the FlexVol volume.



If you run the `df -s` command on a deduplication-enabled FlexVol volume, you can view the space saved by both deduplication and FlexClone files and LUNs.

Example

The following example shows the space saving on a FlexClone volume test1:

```
systemA> df -s test1

Filesystem      used      saved      %saved  Vserver
/vol/test1/      4828      5744      54%    vs1
```

Learn more about the commands described in this procedure in the [ONTAP command reference](#).

Methods to delete FlexClone files and FlexClone LUNs

You can use multiple methods to delete FlexClone files and FlexClone LUNs. Understanding what methods are available helps you plan how to manage clones.

You can use the following methods to delete FlexClone files and FlexClone LUNs:

- You can configure a FlexVol volume to automatically delete clones with autodelete enabled when the free space in a FlexVol volume decreases below a particular threshold.
- You can configure clients to delete clones by using the NetApp Manageability SDK.
- You can use clients to delete clones by using the NAS and SAN protocols.

The slower deletion method is enabled by default because this method does not use the NetApp Manageability SDK. However, you can configure the system to use the faster deletion method when you delete FlexClone files by using the `volume file clone` deletion commands.

How a FlexVol volume can reclaim free space with autodelete setting

FlexVol volumes and reclaiming free space with autodelete overview

You can enable the autodelete setting of a FlexVol volume to automatically delete FlexClone files and FlexClone LUNs. By enabling autodelete, you can reclaim a target amount of free space in the volume when a volume is nearly full.

You can configure a volume to automatically start deleting FlexClone files and FlexClone LUNs when the free space in the volume decreases below a particular threshold value, and automatically stop deleting clones when a target amount of free space in the volume is reclaimed. Although, you cannot specify the threshold value that starts the automatic deletion of clones, you can specify whether a clone is eligible for deletion, and you can specify the target amount of free space for a volume.

A volume automatically deletes FlexClone files and FlexClone LUNs when the free space in the volume decreases below a particular threshold and when *both* of the following requirements are met:

- The autodelete capability is enabled for the volume that contains the FlexClone files and FlexClone LUNs.

You can enable the autodelete capability for a FlexVol volume by using the `volume snapshot autodelete modify` command. You must set the `-trigger` parameter to `volume` or `snap_reserve` for a volume to automatically delete FlexClone files and FlexClone LUNs. Learn more about `volume snapshot autodelete modify` in the [ONTAP command reference](#).

- The autodelete capability is enabled for the FlexClone files and FlexClone LUNs.

You can enable autodelete for a FlexClone file or FlexClone LUN by using the `file clone create` command with the `-autodelete` parameter. As a result, you can preserve certain FlexClone files and FlexClone LUNs by disabling autodelete for the clones and ensuring that other volume settings do not override the clone setting. Learn more about `file clone create` in the [ONTAP command reference](#).

Configure a FlexVol volume to automatically delete FlexClone files and FlexClone LUNs

You can configure a volume to automatically start deleting FlexClone files and FlexClone LUNs when the free space in the volume decreases below a particular threshold value, and automatically stop deleting clones when a target amount of free space in the volume is reclaimed. Although, you cannot specify the threshold value that starts the automatic deletion of clones, you can specify whether a clone is eligible for deletion, and you can specify the target amount of free space for a volume.

A volume automatically deletes FlexClone files and FlexClone LUNs when the free space in the volume decreases below a particular threshold and when *both* of the following requirements are met:

- The autodelete capability is enabled for the volume that contains the FlexClone files and FlexClone LUNs.

You can enable the autodelete capability for a FlexVol volume by using the `volume snapshot autodelete modify` command. You must set the `-trigger` parameter to `volume` or `snap_reserve` for a volume to automatically delete FlexClone files and FlexClone LUNs.

- The autodelete capability is enabled for the FlexClone files and FlexClone LUNs.

You can enable autodelete for a FlexClone file or FlexClone LUN by using the `file clone create` command with the `-autodelete` parameter. As a result, you can preserve certain FlexClone files and FlexClone LUNs by disabling autodelete for the clones and ensuring that other volume settings do not override the clone setting.

Before you begin

- The FlexVol volume must contain FlexClone files and FlexClone LUNs, and be online.
- The FlexVol volume must not be a read-only volume.

Steps

1. Enable automatic deletion of FlexClone files and FlexClone LUNs in the FlexVol volume by using the `volume snapshot autodelete modify` command. Learn more about `volume snapshot autodelete modify` in the [ONTAP command reference](#).
 - For the `-trigger` parameter, you can specify `volume` or `snap_reserve`.
 - For the `-destroy-list` parameter, you must always specify `lun_clone`, `file_clone` regardless of whether you want to delete only one type of clone.The following example shows how you can enable volume `vol1` to trigger the automatic deletion of FlexClone files and FlexClone LUNs for space reclamation until 25% of the volume consists of free

space:

```
cluster1::> volume snapshot autodelete modify -vserver vs1 -volume
vol1 -enabled true -commitment disrupt -trigger volume -target-free
-space 25 -destroy-list lun_clone,file_clone
```

```
Volume modify successful on volume:vol1
```



While enabling FlexVol volumes for automatic deletion, if you set the value of the `-commitment` parameter to `destroy`, all the FlexClone files and FlexClone LUNs with the `-autodelete` parameter set to `true` might be deleted when the free space in the volume decreases below the specified threshold value. However, FlexClone files and FlexClone LUNs with the `-autodelete` parameter set to `false` will not be deleted.

2. Verify that automatic deletion of FlexClone files and FlexClone LUNs is enabled in the FlexVol volume by using the `volume snapshot autodelete show` command. Learn more about `volume snapshot autodelete show` in the [ONTAP command reference](#).

The following example shows that volume `vol1` is enabled for automatic deletion of FlexClone files and FlexClone LUNs:

```
cluster1::> volume snapshot autodelete show -vserver vs1 -volume vol1

Vserver Name: vs1
Volume Name: vol1
Enabled: true
Commitment: disrupt
Defer Delete: user_created
Delete Order: oldest_first
Defer Delete Prefix: (not specified)
Target Free Space: 25%
Trigger: volume
*Destroy List: lun_clone,file_clone*
Is Constituent Volume: false
```

3. Ensure that autodelete is enabled for the FlexClone files and FlexClone LUNs in the volume that you want to delete by performing the following steps:

- a. Enable automatic deletion of a particular FlexClone file or FlexClone LUN by using the `volume file clone autodelete` command. Learn more about `volume file clone autodelete` in the [ONTAP command reference](#).

You can force a specific FlexClone file or FlexClone LUN to be automatically deleted by using the `volume file clone autodelete` command with the `-force` parameter.

The following example shows that automatic deletion of the FlexClone LUN `lun1_clone` contained in volume `vol1` is enabled:

```
cluster1::> volume file clone autodelete -vserver vs1 -clone-path  
/vol/vol1/lun1_clone -enabled true
```

You can enable autodelete when you create FlexClone files and FlexClone LUNs.

b. Verify that the FlexClone file or FlexClone LUN is enabled for automatic deletion by using the `volume file clone show-autodelete` command. Learn more about `volume file clone show-autodelete` in the [ONTAP command reference](#).

The following example shows that the FlexClone LUN `lun1_clone` is enabled for automatic deletion:

```
cluster1::> volume file clone show-autodelete -vserver vs1 -clone  
-path vol/vol1/lun1_clone  
Vserver Name: vs1  
Clone Path: vol/vol1/lun1_clone  
**Autodelete Enabled: true**
```

Learn more about the commands described in this procedure in the [ONTAP command reference](#).

Prevent automatic deletion of a FlexClone file or FlexClone LUN

If you configure a FlexVol volume to automatically delete FlexClone files and FlexClone LUNs, any clone that fits the criteria you specify might be deleted. If you have specific FlexClone files or FlexClone LUNs that you want to preserve, you can exclude them from the automatic FlexClone deletion process.

Before you begin

A FlexClone license must be installed. This license is included with [ONTAP One](#).

About this task

When you create a FlexClone file or FlexClone LUN, by default the autodelete setting for the clone is disabled. FlexClone files and FlexClone LUNs with autodelete disabled are preserved when you configure a FlexVol volume to automatically delete clones to reclaim space on the volume.

If you set the commitment level on the volume to `try` or `disrupt`, you can individually preserve specific FlexClone files or FlexClone LUNs by disabling autodelete for those clones.

However, if you set the commitment level on the volume to `destroy` and the destroy lists include `lun_clone`, `file_clone`, the volume setting overrides the clone setting, and all FlexClone files and FlexClone LUNs can be deleted regardless of the autodelete setting for the clones.

Steps

1. Prevent a specific FlexClone file or FlexClone LUN from being automatically deleted by using the `volume file clone autodelete` command.

The following example shows how you can disable autodelete for FlexClone LUN `lun1_clone` contained in `vol1`:

```
cluster1::> volume file clone autodelete -vserver vs1 -volume vol1  
-clone-path lun1_clone -enable false
```

A FlexClone file or FlexClone LUN with autodelete disabled cannot be deleted automatically to reclaim space on the volume.

2. Verify that autodelete is disabled for the FlexClone file or FlexClone LUN by using the `volume file clone show-autodelete` command.

The following example shows that autodelete is false for the FlexClone LUN `lun1_clone`:

```
cluster1::> volume file clone show-autodelete -vserver vs1 -clone-path  
vol/vol1/lun1_clone  
Name: vs1  
Vserver  
vol/vol1/lun1_clone  
Clone Path:  
Enabled: false  
Autodelete
```

Commands for configuring deletion of FlexClone files

When clients delete FlexClone files without using the NetApp Manageability SDK, you can use the `volume file clone deletion` commands to enable faster deletion of FlexClone files from a FlexVol volume. Extensions for and minimum size of FlexClone files are used to enable faster deletion.

You can use the `volume file clone deletion` commands to specify a list of supported extensions and a minimum size requirement for FlexClone files in a volume. The faster deletion method is used only for FlexClone files that meet the requirements. For FlexClone files that do not meet the requirements, the slower deletion method is used.

When clients delete FlexClone files and FlexClone LUNs from a volume by using the NetApp Manageability SDK, the extension and size requirements do not apply because the faster deletion method is always used.

To...	Use this command...
Add an extension to the supported list of extensions for the volume	<code>volume file clone deletion add-extension</code>
Change the minimum size of FlexClone files that can be deleted from the volume by using the faster deletion method	<code>volume file clone deletion modify</code>
Remove an extension from the supported list of extensions for the volume	<code>volume file clone deletion remove-extension</code>

To...	Use this command...
View the supported list of extensions and the minimum size of FlexClone files that clients can delete from the volume by using the faster deletion method	volume file clone deletion show

Learn more about the commands described in this procedure in the [ONTAP command reference](#).

Related information

- [volume file clone deletion](#)

Use qtrees to partition your FlexVol volumes

Qtrees and ONTAP FlexVol volume partitioning

Qtrees enable you to partition FlexVol volumes into smaller segments that can be managed individually. The volume partitioning enabled by qtrees provides a finer level of control when administering storage by project, user, or group. You can use qtrees to better manage quotas, security style, and CIFS oplocks.



ONTAP creates a default qtree for each volume named **qtree0**. If you do not put data in a specific qtree, it's placed in qtree0.

General limitations

You should be aware of the limitations of qtrees before using them in a production environment. Also review the [Operation and limitations](#) when using the extended qtree performance monitoring feature.

- Qtree names can be no more than 64 characters.
- Certain special characters used in the qtree names, such as commas and spaces, can cause problems with other ONTAP capabilities and should be avoided.
- You cannot move directories between different qtrees. Only files can be moved between qtrees.
- If you create qtree-level and volume-level shares on the same FlexVol or SCVMM pool, the qtrees appear as directories on the FlexVol share. You should be careful not to accidentally delete them.

Commands for managing and configuring qtrees

You can manage and configure qtrees using the ONTAP CLI. Depending on what you want to do, you should use the following commands to administer qtrees.



The command `volume rehost` can cause other concurrent administrative operations targeted at the same volume to fail.

If you want to...	Use this command...
Create a qtree	<code>volume qtree create</code>

Display a filtered list of qtrees	volume qtree show
Delete a qtree	volume qtree delete <div style="display: flex; align-items: center; justify-content: space-between;"> This command will fail unless the qtree is empty or the <code>-force true</code> flag is used. </div>
Modify a qtree's UNIX permissions	volume qtree modify -unix-permissions
Modify a qtree's CIFS oplocks setting	volume qtree oplocks
Modify a qtree's security setting	volume qtree security
Rename a qtree	volume qtree rename
Display a qtree's statistics	volume qtree statistics
Reset a qtree's statistics	volume qtree statistics -reset

Extended qtree performance monitoring

Beginning with ONTAP 9.16.1, you can use the ONTAP REST API to access the extended qtree monitoring capabilities which includes latency metrics and historical statistics.

The ONTAP REST API includes several endpoints related to qtrees. Prior to ONTAP 9.16.1, customers could access real-time statistics for qtrees, including IO operations per second (IOPs) as well as throughput for read, write, and other operations.

The extended qtree performance monitoring available beginning with ONTAP 9.16.1 gives you the ability to monitor real-time latency statistics as well as IOPs and throughput for NFSv3, NFSv4.0, NFSv4.1, NFSv4.2, pNFS (technically a part of NFSv4.1 and NFSv4.2), and CIFS. It also collects and archives statistics to allow viewing of historical performance data.

This extended monitoring provides storage administrators with greater insight into system performance. You can use this data to identify high-use qtrees, potential bottlenecks, and other areas when working to improve quality of service. Being able to analyze these metrics, including trends over a longer period of time, enables you to make more informed data-driven decisions.

Operation and limitations

There are several operational characteristics, including limitations, you should consider before using the extended qtree performance monitoring feature in a production environment.

Remount required

After enabling qtree extended monitoring, you need to remount the affected volume to activate the feature.

Availability of statistics

After enabling extended performance monitoring, the statistical data is not immediately available. This includes IOPS, throughput, and latency statistics. It can take up to five minutes before this data is displayed for a qtree.

Qtrees per cluster

You can enable extended performance monitoring for a maximum of 50,000 qtrees in an ONTAP cluster.

Access extended metrics using the ONTAP REST API

Beginning with ONTAP 9.16.1, you can access the extended qtree performance monitoring feature through the ONTAP REST API. The basic capabilities fall into several categories as described below.

Enable and disable extended performance monitoring

You can access the property `ext_performance_monitoring.enabled` at the endpoint `/api/storage/qtrees` to enable or disable the extended monitoring feature. The POST and PATCH methods are available depending on whether you are creating a new qtree or configuring an existing qtree.

Retrieve global monitoring metrics and settings

Several new global properties have been added to the `/api/storage/qtrees` endpoint. You can retrieve these fields using the GET method.

Retrieve metrics for a specific qtree

You can use the GET method at the endpoint `/api/storage/qtrees/{volume.uuid}/{id}/metrics` to retrieve the new statistics and metrics properties for a specific qtree as defined at a specific volume.

Upgrading and reverting

If you enable the feature in ONTAP 9.16.1, you can upgrade to a subsequent ONTAP release without restrictions. However, there are two scenarios to consider.

Upgrade to 9.16.1 and handling mixed version clusters

The extended performance monitoring feature cannot be used (that is, `ext_performance_monitoring.enabled` cannot be set to `true`) until the effective cluster version (ECV) of the cluster is at 9.16.1.

Revert from 9.16.1

If any qtrees have the property `ext_performance_monitoring.enabled` set to `true`, reverting to 9.15.1 from 9.16.1 is not allowed. The revert operation is blocked. The best practice is to set `ext_performance_monitoring.enabled` to `false` for all qtrees prior to reverting to an earlier ONTAP release.

Learn more

Learn more about the ONTAP REST API, including [what's new with the ONTAP REST API](#), from the ONTAP automation documentation. You should also review the ONTAP automation documentation for details about the ONTAP REST API [qtree endpoints](#).

Obtain a qtree junction path

You can mount an individual qtree by obtaining the junction path or namespace path of the qtree. The qtree path displayed by the CLI command `qtree show -instance` is of the format `/vol/<volume_name>/<qtree_name>`. However, this path does not refer to the junction path or namespace path of the qtree.

Learn more about `qtree show` in the [ONTAP command reference](#).

About this task

You need to know the junction path of the volume to obtain the junction path or namespace path of the qtree.

Steps

1. Use the `vserver volume junction-path` command to obtain the junction path of a volume.

The following example displays the junction path of the volume named `vol1` located on the storage virtual machine (SVM) named `vs0`:

```
cluster1::> volume show -volume vol1 -vserver vs0 -fields junction-path
-----
vs0 vol1 /vol1
```

From the above output, the volume's junction path is `/vol1`. Since qtrees are always rooted at the volume, the junction path or namespace path of the qtree will be `/vol1/qtree1`.

Learn more about `vserver volume junction-path` in the [ONTAP command reference](#).

Directory to qtree conversions

Convert a directory to a qtree

If you have a directory at the root of a FlexVol volume that you want to convert to a qtree, you need to migrate the data contained in the directory to a new qtree with the same name, using your client application.

About this task

The steps you take to convert a directory to a qtree depend on what client you use. The following process outlines the general tasks you need to complete.

Before you begin

You cannot delete a directory if it is associated with an existing CIFS share.

Steps

1. Rename the directory to be made into a qtree.
2. Create a new qtree with the original directory name.
3. Use the client application to move the contents of the directory into the new qtree.
4. Delete the now-empty directory.

Convert a directory to a qtree using a Windows client

To convert a directory to a qtree using a Windows client, you rename the directory, create a qtree on the storage system, and move the contents of the directory to the qtree.

About this task

You must use Windows Explorer for this procedure. You cannot use the Windows command-line interface or the DOS prompt environment.

Steps

1. Open Windows Explorer.
2. Click the folder representation of the directory you want to change.



The directory must reside at the root of its containing volume.

3. From the **File** menu, select **Rename** to give this directory a different name.
4. On the storage system, use the `volume qtree create` command to create a new qtree with the original name of the directory. Learn more about `volume qtree create` in the [ONTAP command reference](#).
5. In Windows Explorer, open the renamed directory folder and select the files inside it.
6. Drag these files into the folder representation of the new qtree.



The more subfolders contained in the folder that you are moving, the longer the move operation takes.

7. From the **File** menu, select **Delete** to delete the renamed, now-empty directory folder.

Convert a directory to a qtree using a UNIX client

To convert a directory to a qtree in UNIX, you rename the directory, create a qtree on the storage system, and move the directory's contents to the qtree.

Steps

1. Open a UNIX client window.
2. Use the `mv` command to rename the directory.

```
client: mv /n/user1/vol1/dir1 /n/user1/vol1/olddir
```

3. From the storage system, use the `volume qtree create` command to create a qtree with the original name.

```
system1: volume qtree create /n/user1/vol1/dir1
```

Learn more about `volume qtree create` in the [ONTAP command reference](#).

4. From the client, use the `mv` command to move the contents of the old directory into the qtree.



The more subdirectories contained in a directory that you are moving, the longer the move operation will take.

```
client: mv /n/user1/vol1/olddir/* /n/user1/vol1/dir1
```

5. Use the `rmdir` command to delete the old, now-empty directory.

```
client: rmdir /n/user1/vol1/olddir
```

After you finish

Depending on how your UNIX client implements the `mv` command, file ownership and permissions might not be preserved. If this occurs, update file owners and permissions to their previous values.

Learn more about the commands described in this procedure in the [ONTAP command reference](#).

Logical space reporting and enforcement for volumes

Logical space reporting and enforcement for volumes overview

Beginning with ONTAP 9.4, you can allow the logical space used in a volume and the amount of remaining storage space to be displayed to users. Beginning with ONTAP 9.5, you can limit the amount of logical space consumed by users.

Logical space reporting and enforcement are disabled by default.

The following volume types support logical space reporting and enforcement.

Volume type	Is space reporting supported?	Is space enforcement supported?
FlexVol volumes	Yes, beginning with ONTAP 9.4	Yes, beginning with ONTAP 9.5
SnapMirror destination volumes	Yes, beginning with ONTAP 9.8	Yes, beginning with ONTAP 9.13.1
FlexGroup volumes	Yes, beginning with ONTAP 9.9.1	Yes, beginning with ONTAP 9.9.1
FlexCache volumes	Origin setting is used at the cache	Not applicable

Logical space enforcement

Logical space enforcement ensures that users are notified when a volume is full or nearly full. When you enable logical space enforcement in ONTAP 9.5 and later, ONTAP counts the logical-used blocks in a volume to determine the amount of space that is still available in that volume. If there is no space available in a volume, the system returns an ENOSPC (out-of-space) error message.

Logical space enforcement returns three types of alerts to inform you about the available space in a volume:

- `Monitor.vol.full.inc.sav`: This alert is triggered when 98% of the logical space in the volume has

been used.

- `Monitor.vol.nearFull.inc.sav`: This alert is triggered when 95% of the logical space in the volume has been used.
- `Vol.log.overalloc.inc.sav`: This alert is triggered when the logical space used in the volume is greater than the total size of the volume.

This alert tells you that adding to the size of the volume might not create available space since that space will already be consumed by overallocated logical blocks.



Total (logical space) should be equal to provisioned space excluding snapshot reserve of the volume with logical space enforcement.

For more information, see [Configuring volumes to automatically provide more space when they are full](#).

Logical space reporting

When you enable logical space reporting on a volume, your system can display the amount of logical used and available space in addition to the total space in a volume. In addition, users on Linux and Windows client systems can see logical used and available space instead of physical used and physical available space.

Definitions:

- Physical space refers to the physical blocks of storage available or used in the volume.
- Logical space refers to the usable space in a volume.
- Logical space used is physical space used plus savings from storage efficiency features (such as deduplication and compression) that have been configured.

Beginning with ONTAP 9.5, you can enable logical space enforcement together with space reporting.

When enabled, logical space reporting displays the following parameters with the `volume show` command:

Parameter	Meaning
<code>-logical-used</code>	Displays information only about the volume or volumes that have the specified logical used size. This value includes all the space saved by the storage efficiency features along with the physically used space. This does not include snapshot reserve but does consider snapshot spill.
<code>-logical-used-by-afs</code>	Displays information only about the volume or volumes that have the specified logical size used by the active file system. This value differs from the <code>-logical-used</code> value by the amount of snapshot spill that exceeds the snapshot reserve.
<code>-logical-available</code>	When only logical space reporting is enabled, only physical-available space is displayed. When both space reporting and enforcement are enabled, it displays the amount of free space currently available considering space saved by the storage efficiency features as being used. This does not include the snapshot reserve.

Parameter	Meaning
-logical-used -percent	<p>Displays the percentage of the current <code>-logical-used</code> value with the provisioned size excluding snapshot reserve of the volume.</p> <p>This value can be greater than 100%, because the <code>-logical-used-by-afs</code> value includes efficiency savings in the volume. The <code>-logical-used-by-afs</code> value of a volume does not include snapshot spill as used space. The <code>-physical-used</code> value of a volume includes Snapshot spill as used space.</p>
-used	<p>Displays the amount of space occupied by user data and file system metadata. It differs from <code>physical-used</code> space by the sum of the space that is reserved for future writes and the space that is saved by aggregate storage efficiency. It includes snapshot spill (the amount of space by which snapshots exceed snapshot reserve). It does not include the snapshot reserve.</p>

Enabling logical space reporting in the CLI also allows the Logical Used Space (%) and Logical Space values to display in System Manager

Client systems see logical space displayed as “used” space on the following system displays:

- `df` output on Linux systems
- Space details under Properties using Windows Explorer on Windows systems.



If logical space reporting is enabled without logical space enforcement, the total displayed on client systems can be higher than the provisioned space.

Enable logical space reporting and enforcement

Beginning with ONTAP 9.4, you can enable logical space reporting. Beginning with 9.5, you can enable logical space enforcement, or both reporting and enforcement together.

About this task

In addition to enabling logical space reporting and enforcement at the individual volume level, you can enable them at the SVM level for every volume that supports the functionality. If you enable logical space features for the entire SVM, you can also disable them for individual volumes.

Beginning with ONTAP 9.8, if you enable logical space reporting on a SnapMirror source volume, it is automatically enabled on the destination volume after the transfer.

Beginning with ONTAP 9.13.1, if the enforcement option is enabled on a SnapMirror source volume, the destination will report logical space consumption and will honor its enforcement, enabling better capacity planning.



If you are running an ONTAP release earlier than ONTAP 9.13.1, you should understand that although the enforcement setting is transferred to the SnapMirror destination volume, the destination volume does not support enforcement. As a result, the destination will report logical space consumption but not honor its enforcement.

Learn more about [ONTAP release support for logical space reporting](#).

Steps

Enable one or more of the following:

- Enable logical space reporting for a volume:

```
volume modify -vserver svm_name -volume volume_name -size volume_size -is-space-reporting-logical true
```

- Enable logical space enforcement for a volume:

```
volume modify -vserver svm_name -volume volume_name -size volume_size -is-space-enforcement-logical true
```

- Enable logical space reporting and enforcement together for a volume:

```
volume modify -vserver svm_name -volume volume_name -size volume_size -is-space-reporting-logical true -is-space-enforcement-logical true
```

- Enable logical space reporting or enforcement for a new SVM:

```
vserver create -vserver _svm_name_ -rootvolume root_volume_name_ -rootvolume-security-style unix -data-services {desired-data-services} [-is-space-reporting-logical true] [-is-space-enforcement-logical true]
```

- Enable logical space reporting or enforcement for an existing SVM:

```
vserver modify -vserver _svm_name_ {desired-data-services} [-is-space-reporting-logical true] [-is-space-enforcement-logical true]
```

Manage SVM capacity limits

Beginning with ONTAP 9.13.1, you can set a maximum capacity for a storage VM (SVM). You can also configure alerts when the SVM approaches a threshold capacity level.

About this task

Capacity on an SVM is calculated as the sum of FlexVols, FlexGroup volumes, FlexClones, FlexCache volumes. Volumes impact capacity calculation even if they are restricted, offline, or in the recovery queue after deletion. If you have volumes configured with auto-grow, the maximum autosize value of the volume will be calculated toward the SVM size; without auto-grow, the actual size of the volume will be calculated.

The following table captures how `autosize-mode` parameters impact the capacity calculation.

<code>autosize-mode off</code>	Size parameter will be used for computation
<code>autosize-mode grow</code>	The <code>max-autosize</code> parameter will be used for computation
<code>autosize-mode grow-shrink</code>	The <code>max-autosize</code> parameter will be used for computation

Before you begin

- You must be a cluster administrator to set an SVM limit.

- Beginning with ONTAP 9.16.1, storage limits can be configured for SVMs that contain data protection volumes, including the following data protection types:
 - FlexVol volumes in asynchronous DR without cascade
 - FlexVol volumes in synchronous DR (both sync and strict-sync policies)
 - [Restore](#)
- Storage limits for SVMs is *not* supported for the following configurations:
 - SnapMirror vault relationships
 - SnapMirror active sync
 - FlexGroup volumes
 - Consistency groups
 - SVM DR
 - Cascades
 - MetroCluster
- Beginning with ONTAP 9.16.1, when you create a load-sharing mirror relationship, the destination SVM cannot have a storage limit enabled.
- When you migrate an SVM, the source SVM cannot have a storage limit enabled. To complete the migrate operation, disable the storage limit on the source then complete the migration.
- SVM capacity is distinct from [quotas](#). Quotas cannot exceed the max size.
- You cannot set a storage limit when other operations are in progress on the SVM. Use the `job show vserver <svm_name>` command to see existing jobs. Try running the command again when any jobs have been completed. Learn more about `job show` in the [ONTAP command reference](#).

Capacity impact

When you reach the capacity limit, the following operations will fail:

- Creating a LUN, namespace, or volume
- Cloning a LUN, namespace, or volume
- Modifying a LUN, namespace, or volume
- Increasing the size of a LUN, namespace, or volume
- Expanding a LUN, namespace, or volume
- Rehosting a LUN, namespace, or volume

Set a capacity limit on a new SVM

System Manager

Steps

1. Select **Storage > Storage VMs**.
2. Select  to create the SVM.
3. Name the SVM and select an **Access protocol**.
4. Under **Storage VM settings**, select **Enable maximum capacity limit**.

Provide a maximum capacity size for the SVM.

5. Select **Save**.

CLI

Steps

1. Create the SVM. To set a storage limit, provide a **storage-limit** value. To set a threshold alert for the storage limit, provide a percentage value for **-storage-limit-threshold-alert**.

```
vserver create -vserver <vserver_name> -aggregate <aggregate_name>
-rootvolume <root_volume_name> -rootvolume-security-style
{unix|ntfs|mixed} -storage-limit <value> [GiB|TiB] -storage-limit
-threshold-alert <percentage> [-ipspace <IPspace_name>] [-language
<language>] [-snapshot-policy <snapshot_policy_name>] [-quota-policy
<quota_policy_name>] [-comment <comment>]
```

If you do not provide threshold value, by default an alert will be triggered when the SVM is at 90% capacity. To disable the threshold alert, provide a value of zero.

2. Confirm the SVM was created successfully:

```
vserver show -vserver <vserver_name>
```

3. If you wish to disable the storage limit, set the **-storage-limit** parameter for the SVM to zero:

```
vserver modify -vserver <vserver_name> -storage-limit 0
```

Set or modify a capacity limit on an existing SVM

You can set a capacity limit and threshold alert on an existing SVM or disable a capacity limit.

Once you set the capacity limit, you cannot modify the limit to a value less than the currently allocated capacity.

System Manager

Steps

1. Select **Storage > Storage VMs**.
2. Select the SVM you want to modify. Next to the name of the SVM, select  then **Edit**.
3. To enable a capacity limit, select the box next to **Enable capacity limit**. Enter a value for the **Maximum capacity** and a percentage value for **Alert threshold**.
If you wish to disable the capacity limit, uncheck the box next **Enable capacity limit**.
4. Select **Save**.

CLI

Steps

1. On the cluster hosting the SVM, issue the `vserver modify` command. Provide a numerical value for `-storage-limit` and a percent value for `-storage-limit-threshold-alert`.

```
vserver modify -vserver <vserver_name> -storage-limit <value>
[GiB|TIB] -storage-limit-threshold-alert <percentage>
```

If you do not provide a threshold value, you will have a default alert at 90% capacity. To disable the threshold alert, provide a value of zero.

2. If you wish to disable the storage limit, set the `-storage-limit` for the SVM to zero:

```
vserver modify -vserver <vserver_name> -storage-limit 0
```

Reaching capacity limits

When you reach the maximum capacity or the alert threshold, you can consult the `vserver.storage.threshold` EMS messages or use the **Insights** page in System Manager to learn about possible actions. Possible resolutions include:

- Editing the SVM maximum capacity limits
- Purging the volumes recovery queue to free up space
- Delete snapshot to provide space for the volume

Related information

- [Capacity measurements in System Manager](#)
- [Monitor cluster, tier, and SVM capacity in System Manager](#)
- [vserver create](#)
- [vserver show](#)
- [vserver modify](#)

Use quotas to restrict or track resource usage

Overview of the quota process

Understand quotas, quota rules, and quota policies

Quotas are defined in quota rules specific to FlexVol volumes. These quota rules are collected together in a quota policy for a storage virtual machine (SVM) and activated on each volume on the SVM.

A quota rule is always specific to a volume. Quota rules have no effect until quotas are activated on the volume defined in the quota rule.

A quota policy is a collection of quota rules for all the volumes of an SVM. Quota policies are not shared among SVMs. An SVM can have up to five quota policies, which enable you to have backup copies of quota policies. One quota policy is assigned to an SVM at any given time. When you initialize or resize quotas on a volume, you are activating the quota rules in the quota policy that is currently assigned to the SVM.

A quota is the actual restriction that ONTAP enforces or the actual tracking that ONTAP performs. A quota rule always results in at least one quota, and might result in many additional derived quotas. The complete list of enforced quotas is visible only in quota reports.

Activation is the process of triggering ONTAP to create enforced quotas from the current set of quota rules in the assigned quota policy. Activation occurs on a volume-by-volume basis. The first activation of quotas on a volume is called initialization. Subsequent activations are called either reinitialization or resizing, depending on the scope of the changes.

Benefits of using quotas

You can use quotas to manage and monitor resource usage with FlexVol volumes.

There are several benefits to defining quotas. You can use the default, explicit, derived, and tracking quotas to manage disk usage in the most efficient manner.

Limit resource consumption

You can limit the amount of disk space or the number of files used by a user or group or contained in a qtree.

Track resource usage

The amount of disk space or number of files used by a user, group, or qtree can be tracked without imposing a limit.

Notify users

Notifications can be generated when resource usage reaches specific levels. This warns users when their disk or file usage is too high.

Quota process

Quotas provide a way to restrict or track the disk space and number of files used by a user, group, or qtree. Quotas are applied to a specific FlexVol volume or qtree.

Quotas can be soft or hard. Soft quotas cause ONTAP to send a notification when specified limits are exceeded, and hard quotas prevent a write operation from succeeding when specified limits are exceeded.

When ONTAP receives a request from a user or user group to write to a FlexVol volume, it checks to see whether quotas are activated on that volume for the user or user group and determines the following:

- Whether the hard limit will be reached

If yes, the write operation fails when the hard limit is reached and the hard quota notification is sent.

- Whether the soft limit will be breached

If yes, the write operation succeeds when the soft limit is breached and the soft quota notification is sent.

- Whether a write operation will not exceed the soft limit

If yes, the write operation succeeds and no notification is sent.

Differences among hard, soft, and threshold quotas

Hard quotas prevent operations while soft quotas trigger notifications.

Hard quotas impose a hard limit on system resources; any operation that would result in exceeding the limit fails. The following settings create hard quotas:

- Disk Limit parameter
- Files Limit parameter

Soft quotas send a warning message when resource usage reaches a certain level, but do not affect data access operations, so you can take appropriate action before the quota is exceeded. The following settings create soft quotas:

- Threshold for Disk Limit parameter
- Soft Disk Limit parameter
- Soft Files Limit parameter

Threshold and Soft Disk quotas enable administrators to receive more than one notification about a quota. Typically, administrators set the Threshold for Disk Limit to a value that is only slightly smaller than the Disk Limit, so that the threshold provides a "final warning" before writes start to fail.

About quota notifications

Quota notifications are messages that are sent to the event management system (EMS) and also configured as SNMP traps.

Notifications are sent in response to the following events:

- A hard quota is reached; in other words, an attempt is made to exceed it
- A soft quota is exceeded
- A soft quota is no longer exceeded

Thresholds are slightly different from other soft quotas. Thresholds trigger notifications only when they are exceeded, not when they are no longer exceeded.

Hard-quota notifications are configurable by using the volume quota modify command. You can turn them off

completely, and you can change their frequency, for example, to prevent sending of redundant messages.

Soft-quota notifications are not configurable because they are unlikely to generate redundant messages and their sole purpose is notification.

The following table lists the events that quotas send to the EMS system:

When this occurs...	This event is sent to the EMS...
A hard limit is reached in a tree quota	wafl.quota.qtree.exceeded
A hard limit is reached in a user quota on the volume	wafl.quota.user.exceeded (for a UNIX user) wafl.quota.user.exceeded.win (for a Windows user)
A hard limit is reached in a user quota on a qtree	wafl.quota.userQtree.exceeded (for a UNIX user) wafl.quota.userQtree.exceeded.win (for a Windows user)
A hard limit is reached in a group quota on the volume	wafl.quota.group.exceeded
A hard limit is reached in a group quota on a qtree	wafl.quota.groupQtree.exceeded
A soft limit, including a threshold, is exceeded	quota.softlimit.exceeded
A soft limit is no longer exceeded	quota.softlimit.normal

The following table lists the SNMP traps that quotas generate:

When this occurs...	This SNMP trap is sent...
A hard limit is reached	quotaExceeded
A soft limit, including a threshold, is exceeded	quotaExceeded and softQuotaExceeded
A soft limit is no longer exceeded	quotaNormal and softQuotaNormal



Notifications contain qtree ID numbers rather than qtree names. You can correlate qtree names to ID numbers by using the `volume qtree show -id` command.

Quota targets and types

Every quota has a specific type. The quota target is derived from the type and specifies the user, group, or qtree to which the quota limits are applied.

The following table lists the quota targets, what types of quotas each quota target is associated with, and how each quota target is represented.

Quota target	Quota type	How target is represented	Notes
user	user quota	UNIX user name UNIX UID A file or directory whose UID matches the user Windows user name in pre-Windows 2000 format Windows SID A file or directory with an ACL owned by the user's SID	User quotas can be applied for a specific volume or qtree.
group	group quota	UNIX group name UNIX GID A file or directory whose GID matches the group	Group quotas can be applied for a specific volume or qtree.  ONTAP does not apply group quotas based on Windows IDs.
qtree	tree quota	qtree name	Tree quotas are applied to a particular volume and do not affect qtrees in other volumes.
""	user quota group quota tree quota	Double quotation marks ("")	A quota target of "" denotes a <i>default quota</i> . For default quotas, the quota type is determined by the value of the type field.

Special kinds of quotas

How default quotas work

You can use default quotas to apply a quota to all instances of a given quota type. For example, a default user quota affects all users on the system for the specified FlexVol volume or qtree. In addition, default quotas enable you to modify your quotas easily.

You can use default quotas to automatically apply a limit to a large set of quota targets without having to create separate quotas for each target. For example, if you want to limit most users to 10 GB of disk space, you can specify a default user quota of 10 GB of disk space instead of creating a quota for each user. If you have specific users for whom you want to apply a different limit, you can create explicit quotas for those users. (Explicit quotas—quotas with a specific target or list of targets—override default quotas.)

In addition, default quotas enable you to use resizing rather than reinitialization when you want quota changes to take effect. For example, if you add an explicit user quota to a volume that already has a default user quota, you can activate the new quota by resizing.

Default quotas can be applied to all three types of quota target (users, groups, and qtrees).

Default quotas do not necessarily have specified limits; a default quota can be a tracking quota.

A quota is indicated by a target that is either an empty string ("") or an asterisk (*), depending on the context:

- When you create a quota using the `volume quota policy rule create` command, setting the `-target` parameter to an empty string ("") creates a default quota.

Learn more about `volume quota policy rule create` in the [ONTAP command reference](#).

- In the `volume quota policy rule create` command, the `-qtree` parameter specifies the name of the qtree to which the quota rule applies. This parameter is not applicable for tree type rules. For user or group type rules at the volume level, this parameter should contain "".

- In the output of the `volume quota policy rule show` command, a default quota appears with an empty string ("") as the target.

Learn more about `volume quota policy rule show` in the [ONTAP command reference](#).

- In the output of the `volume quota report` command, a default quota appears with an asterisk (*) as the ID and Quota Specifier.

Learn more about `volume quota report` in the [ONTAP command reference](#).

Default user quota example

The following quota rule uses a default user quota to apply a 50-MB limit on each user for vol1:

```
cluster1::> volume quota policy rule create -vserver vs0 -volume vol1
-priority 1 -policy-name default -type user -target "" -qtree "" -disk-limit 50m

cluster1::> volume quota policy rule show -vserver vs0 -volume vol1

Vserver: vs0          Policy: default          Volume: vol1
                                         Soft          Soft
                                         User        Disk        Disk    Files    Files
                                         Type       Target     Qtree     Mapping   Limit   Limit
                                         Threshold
-----  -----  -----  -----  -----  -----  -----  -----  -----
-----  -----
user    ""        ""        off        50MB      -        -        -
-
```

If any user on the system enters a command that would cause that user's data to take up more than 50 MB in vol1 (for example, writing to a file from an editor), the command fails.

How you use explicit quotas

You can use explicit quotas to specify a quota for a specific quota target, or to override a default quota for a specific target.

An explicit quota specifies a limit for a particular user, group, or qtree. An explicit quota replaces any default quota that is in place for the same target.

When you add an explicit user quota for a user that has a derived user quota, you must use the same user mapping setting as the default user quota. Otherwise, when you resize quotas, the explicit user quota is rejected because it is considered a new quota.

Explicit quotas only affect default quotas at the same level (volume or qtree). For example, an explicit user quota for a qtree does not affect the default user quota for the volume that contains that qtree. However, the explicit user quota for the qtree overrides (replaces the limits defined by) the default user quota for that qtree.

Examples of explicit quotas

The following quota rules define a default user quota that limits all users in vol1 to 50MB of space. However, one user, jsmith, is allowed 80MB of space, because of the explicit quota (shown in bold):

```
cluster1::> volume quota policy rule create -vserver vs0 -volume vol1
-policy-name default -type user -target "" -qtree "" -disk-limit 50m

cluster1::> volume quota policy rule create -vserver vs0 -volume vol1
-policy-name default -type user -target "jsmith" -qtree "" -disk-limit 80m

cluster1::> volume quota policy rule show -vserver vs0 -volume vol1

Vserver: vs0          Policy: default          Volume: vol1
                           User           Disk           Disk           Soft
                           Mapping        Limit        Limit        Files        Files
Type   Target   Qtree      User           Disk           Disk           Soft
      Threshold      Mapping        Limit        Limit        Files        Files
-----  -----  -----  -----  -----  -----  -----  -----  -----
-----  -----
user   ""      ""      off      50MB      -      -      -      -
-
user   jsmith  ""      off      80MB      -      -      -      -
```

The following quota rule restricts the specified user, represented by four IDs, to 550MB of disk space and 10,000 files in the vol1 volume:

```

cluster1::> volume quota policy rule create -vserver vs0 -volume vol1
-policy-name default -type user -target "
jsmith,corp\jsmith,engineering\john smith,S-1-5-32-544" -qtree "" -disk
-limit 550m -file-limit 10000

cluster1::> volume quota policy rule show -vserver vs0 -volume vol1

Vserver: vs0          Policy: default          Volume: vol1
                                         Soft          Soft
                                         User       Disk       Disk   Files   Files
                                         Mapping   Limit   Limit   Limit   Limit
Type   Target   Qtree   Mapping   Limit   Limit   Limit   Limit
Threshold
-----
-----
user   "jsmith,corp\jsmith,engineering\john smith,S-1-5-32-544"
      ""       off      550MB      -     10000      -
-
```

The following quota rule restricts the eng1 group to 150MB of disk space and an unlimited number of files in the proj1 qtree:

```

cluster1::> volume quota policy rule create -vserver vs0 -volume vol2
-policy-name default -type group -target "eng1" -qtree "proj1" -disk-limit
150m

cluster1::> volume quota policy rule show -vserver vs0 -volume vol2

Vserver: vs0          Policy: default          Volume: vol2
                                         Soft          Soft
                                         User       Disk       Disk   Files   Files
                                         Mapping   Limit   Limit   Limit   Limit
Type   Target   Qtree   Mapping   Limit   Limit   Limit   Limit
Threshold
-----
-----
group  eng1      proj1    off      150MB      -      -      -
-
```

The following quota rule restricts the proj1 qtree in the vol2 volume to 750MB of disk space and 75,000 files:

```
cluster1::> volume quota policy rule create -vserver vs0 -volume vol2
-policy-name default -type tree -target "proj1" -disk-limit 750m -file
-limit 75000
```

```
cluster1::> volume quota policy rule show -vserver vs0 -volume vol2
```

Vserver: vs0		Policy: default		Volume: vol2							
Type	Target	Qtree	Mapping	User	Disk	Disk	Files	Files	Soft	Soft	
Threshold				Limit		Limit		Limit		Limit	
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	
tree	proj1	""	-		750MB	-	75000	-			
-											

How derived quotas work

A quota enforced as a result of a default quota, rather than an explicit quota (a quota with a specific target), is referred to as a *derived quota*.

The number and location of the derived quotas depends on the quota type:

- A default tree quota on a volume creates derived default tree quotas for every qtree on the volume.
- A default user or group quota creates a derived user or group quota for every user or group that owns a file at the same level (volume or qtree).
- A default user or group quota on a volume creates a derived default user or group quota on every qtree that also has a tree quota.

The settings—including limits and user mapping—of derived quotas are the same as the settings of the corresponding default quotas. For example, a default tree quota with a 20-GB disk limit on a volume creates derived tree quotas with 20-GB disk limits on the qtrees in the volume. If a default quota is a tracking quota (with no limits), the derived quotas are also tracking quotas.

To see derived quotas, you can generate a quota report. In the report, a derived user or group quota is indicated by a Quota Specifier that is either blank or an asterisk (*). A derived tree quota, however, has a Quota Specifier. To identify a derived tree quota, you must look for a default tree quota on the volume with the same limits.

Explicit quotas interact with derived quotas in the following ways:

- Derived quotas are not created if an explicit quota already exists for the same target.
- If a derived quota exists when you create an explicit quota for a target, you can activate the explicit quota by resizing rather than having to perform a full quota initialization.

Use tracking quotas

A tracking quota generates a report of disk and file usage and does not limit resource

usage. When tracking quotas are used, modifying the quota values is less disruptive because you can resize quotas rather than turning them off and back on.

To create a tracking quota, you omit the Disk Limit and Files Limit parameters. This tells ONTAP to monitor disk and files usage for that target at that level (volume or qtree), without imposing any limits. Tracking quotas are indicated in the output of `show` commands and the quota report with a dash ("") for all limits. ONTAP automatically creates tracking quotas when you use the System Manager UI to create explicit quotas (quotas with specific targets). When using the CLI, the storage administrator creates tracking quotas on top of explicit quotas.

You can also specify a *default tracking quota*, which applies to all instances of the target. Default tracking quotas enable you to track usage for all instances of a quota type (for example, all qtrees or all users). In addition, they enable you to use resizing rather than reinitialization when you want quota changes to take effect.

Examples

The output for a tracking rule shows tracking quotas in place for a qtree, user, and group, as shown in the following example for a volume-level tracking rule:

Vserver: vs0			Policy: default				Volume: fv1		
Type	Target	Qtree	User		Disk		Files		Threshold
			Mapping	Limit	Disk	Limit	Files	Limit	
tree	""	""	-	-	-	-	-	-	-
user	""	""	off	-	-	-	-	-	-
group	""	""	-	-	-	-	-	-	-

How quotas are applied

Understanding how quotas are applied enables you to configure quotas properly and set the expected limits.

Whenever an attempt is made to create a file or write data to a file in a FlexVol volume that has quotas enabled, the quota limits are checked before the operation proceeds. If the operation exceeds either the disk limit or the files limit, the operation is prevented.

Quota limits are checked in the following order:

1. The tree quota for that qtree (This check is not relevant if the file is being created or written to qtree0.)
2. The user quota for the user that owns the file on the volume
3. The group quota for the group that owns the file on the volume
4. The user quota for the user that owns the file on the qtree (This check is not relevant if the file is being created or written to qtree0.)
5. The group quota for the group that owns the file on the qtree (This check is not relevant if the file is being created or written to qtree0.)

The quota with the smallest limit might not be the one that is exceeded first. For example, if a user quota for volume vol1 is 100 GB, and the user quota for qtree q2 contained in volume vol1 is 20 GB, the volume limit

could be reached first if that user has already written more than 80 GB of data in volume vol1 (but outside of qtree q2).

Related information

- [How quotas are applied to the root user](#)
- [How quotas are applied to users with multiple IDs](#)

Considerations for assigning quota policies

A quota policy is a grouping of the quota rules for all the FlexVol volumes of an SVM. You must be aware of certain considerations when assigning the quota policies.

- An SVM has one assigned quota policy at any given time. When an SVM is created, a blank quota policy is created and assigned to the SVM. This default quota policy has the name "default" unless a different name is specified when the SVM is created.
- An SVM can have up to five quota policies. If an SVM has five quota policies, you cannot create a new quota policy for the SVM until you delete an existing quota policy.
- When you need to create a quota rule or change quota rules for a quota policy, you can choose either of the following approaches:
 - If you are working in a quota policy that is assigned to an SVM, you need not assign the quota policy to the SVM.
 - If you are working in an unassigned quota policy and then assigning the quota policy to the SVM, you must have a backup of the quota policy that you can revert to if required.

For example, you can make a copy of the assigned quota policy, change the copy, assign the copy to the SVM, and rename the original quota policy.

- You can rename a quota policy even when it is assigned to the SVM.

How quotas work with users and groups

Overview of how quotas work with users and groups

You can specify a user or group as the target of a quota. There are several implementation differences to consider when defining a quota.

A few of the differences you need to be aware of include the following:

- User or group
- UNIX or Windows
- Special users and groups
- Are multiple IDs included

There are also different ways to specify IDs for users based on your environment.

Specify UNIX users for quotas

You can specify a UNIX user for a quota in one of several different formats.

The three formats available when specifying a UNIX user for a quota include the following:

- The user name (such as jsmith).



You cannot use a UNIX user name to specify a quota if that name includes a backslash (\) or an @ sign. This is because ONTAP treats names containing these characters as Windows names.

- The user ID or UID (such as 20).

- The path of a file or directory owned by that user, so that the file's UID matches the user.



If you specify a file or directory name, you must select a file or directory that will last as long as the user account remains on the system.

Specifying a file or directory name for the UID does not cause ONTAP to apply a quota to that file or directory.

Specify Windows users for quotas

You can specify a Windows user for a quota in one of several different formats.

The three formats available when specifying a Windows user for a quota include the following:

- The Windows name in pre-Windows 2000 format.
- The security ID (SID) as displayed by Windows in text form, such as S-1-5-32-544.
- The name of a file or directory that has an ACL owned by that user's SID.

If you specify a file or directory name, you must select a file or directory that will last as long as the user account remains on the system.

For ONTAP to obtain the SID from the ACL, the ACL must be valid.



If the file or directory exists in a UNIX-style qtree, or if the storage system uses UNIX mode for user authentication, ONTAP applies the user quota to the user whose **UID**, not SID, matches that of the file or directory.

Specifying a file or directory name to identify a user for a quota does not cause ONTAP to apply a quota to that file or directory.

How default user and group quotas create derived quotas

When you create default user or group quotas, corresponding derived user or group quotas are automatically created for every user or group that owns files at the same level.

Derived user and group quotas are created in the following ways:

- A default user quota on a FlexVol volume creates derived user quotas for every user that owns a file anywhere on the volume.
- A default user quota on a qtree creates derived user quotas for every user that owns a file in the qtree.
- A default group quota on a FlexVol volume creates derived group quotas for every group that owns a file anywhere on the volume.

- A default group quota on a qtree creates derived group quotas for every group that owns a file in the qtree.

If a user or group does not own files at the level of a default user or group quota, derived quotas are not created for the user or group. For example, if a default user quota is created for qtree proj1 and the user jsmith owns files on a different qtree, no derived user quota is created for jsmith.

The derived quotas have the same settings as the default quotas, including limits and user mapping. For example, if a default user quota has a 50-MB disk limit and has user mapping turned on, any resulting derived quotas also have a 50-MB disk limit and user mapping turned on.

However, no limits exist in derived quotas for three special users and groups. If the following users and groups own files at the level of a default user or group quota, a derived quota is created with the same user-mapping setting as the default user or group quota, but it is only a tracking quota (with no limits):

- UNIX root user (UID 0)
- UNIX root group (GID 0)
- Windows BUILTIN\Administrators group

Since quotas for Windows groups are tracked as user quotas, a derived quota for this group is a user quota that is derived from a default user quota, not a default group quota.

Example of derived user quotas

If you have volume where three users—root, jsmith, and bob—own files, and you create a default user quota on the volume, ONTAP automatically creates three derived user quotas. Therefore, after you reinitialize quotas on the volume, four new quotas appear in the quota report:

```
cluster1::> volume quota report
Vserver: vsl
          ----Disk----  ----Files-----  Quota
Volume   Tree     Type   ID      Used  Limit    Used  Limit
Specifier
-----  -----  -----  -----  -----  -----  -----  -----
-----  -----
vol1      user    *      0B    50MB      0      -    *
vol1      user    root   5B      -        1      -    -
vol1      user    jsmith 30B   50MB     10      -    *
vol1      user    bob    40B   50MB     15      -    *
4 entries were displayed.
```

The first new line is the default user quota that you created, which is identifiable by the asterisk (*) as the ID. The other new lines are the derived user quotas. The derived quotas for jsmith and bob have the same 50-MB disk limit as the default quota. The derived quota for the root user is a tracking quota without limits.

How quotas are applied to the root user

The root user (UID=0) on UNIX clients is subject to tree quotas, but not to user or group quotas. This allows the root user to take actions on behalf of other users that would otherwise be prevented by a quota.

When the root user carries out a file or directory ownership change or other operation (such as the UNIX chown command) on behalf of a user with less privileges, ONTAP checks the quotas based on the new owner but does not report errors or stop the operation, even if the hard quota restrictions of the new owner are exceeded. This can be useful when an administrative action, such as recovering lost data, results in temporarily exceeding quotas.



After the ownership transfer is carried out, however, a client system will report a disk space error if the user attempts to allocate more disk space while the quota is still exceeded.

Related information

- [How quotas are applied](#)
- [How quotas are applied to users with multiple IDs](#)

How quotas work with special Windows groups

There are several special Windows groups that process quotas differently than other Windows groups. You should understand how quotas are applied for these special groups.



ONTAP does not support group quotas based on Windows group IDs. If you specify a Windows group ID as the quota target, the quota is considered to be a user quota.

Everyone

When the quota target is the Everyone group, a file with an ACL showing the owner is Everyone is counted under the SID for Everyone.

BUILTIN\Administrators

When the quota target is the BUILTIN\Administrators group, the entry is considered to be a user quota and is used for tracking only. You cannot impose restrictions on BUILTIN\Administrators. If a member of BUILTIN\Administrators creates a file, the file is owned by BUILTIN\Administrators and is counted under the SID for BUILTIN\Administrators (not the user's personal SID).

How quotas are applied to users with multiple IDs

A user can be represented by multiple IDs. You can define a single user quota for such a user by specifying a list of IDs as the quota target. A file owned by any of these IDs is subject to the restriction of the user quota.

Suppose a user has the UNIX UID 20 and the Windows IDs corp\john_smith and engineering\jsmith. For this user, you can specify a quota where the quota target is a list of the UID and Windows IDs. When this user writes to the storage system, the specified quota applies, regardless of whether the write originates from UID 20, corp\john_smith, or engineering\jsmith.

Note that separate quota rules are considered separate targets, even if the IDs belong to the same user. For example, for the same user you can specify one quota that limits UID 20 to 1GB of disk space and another quota that limits corp\john_smith to 2GB of disk space, even though both IDs represent the same user. ONTAP applies quotas to UID 20 and corp\john_smith separately. In this case, no limits are applied to engineering\jsmith, even though limits are applied to the other IDs used by the same user.

Related information

- [How quotas are applied](#)
- [How quotas are applied to the root user](#)

How ONTAP determines user IDs in a mixed environment

If you have users accessing your ONTAP storage from both Windows and UNIX clients, both Windows and UNIX security are used to determine file ownership. Several factors determine whether ONTAP uses a UNIX or Windows ID when applying user quotas.

If the security style of the qtree or FlexVol volume that contains the file is only NTFS or only UNIX, then the security style determines the type of ID used when applying user quotas. For qtrees with the mixed security style, the type of ID used is determined by whether the file has an ACL.

The following table summarizes what type of ID is used.

Security Style	ACL	No ACL
UNIX	UNIX ID	UNIX ID
Mixed	Windows ID	UNIX ID
NTFS	Windows ID	Windows ID

How quotas work with multiple users

When you place multiple users in the same quota target, the limits defined by the quota are not applied to each individual user. Rather, the quota limits are shared among all users in the quota target.

Unlike with commands for managing objects, such as volumes and qtrees, you cannot rename a quota target, including a multi-user quota. This means that after a multi-user quota is defined, you cannot modify the users in the quota target, and you cannot add users to a target or remove users from a target. If you want to add or remove a user from a multi-user quota, then the quota containing that user must be deleted and a new quota rule with the set of users in the target defined.

 If you combine separate user quotas into one multi-user quota, you can activate the change by resizing quotas. However, if you want to remove users from a quota target with multiple users, or add users to a target that already has multiple users, you must reinitialize quotas before the change takes effect.

Example of more than one user in a quota rule

In the following example, there are two users listed in the quota entry. The two users can use up to 80MB of space combined. If one uses 75MB, then the other one can use only 5MB.

```
cluster1::> volume quota policy rule create -vserver vs0 -volume vol1  
-policy-name default -type user -target "jsmith,chen" -qtree "" -disk  
-limit 80m
```

```
cluster1::> volume quota policy rule show -vserver vs0 -volume vol1
```

Vserver: vs0		Policy: default		Volume: vol1			
Type	Target	Qtree	User Mapping	Disk Limit	Disk Limit	Files Limit	Files Limit
Threshold							
-----	-----	-----	-----	-----	-----	-----	-----
-----	-----	-----	-----	-----	-----	-----	-----
user	"jsmith,chen"	""	off	80MB	-	-	-
-	-	-	-	-	-	-	-

UNIX and Windows name linking for quotas

In a mixed environment, users can log in as either Windows users or UNIX users. You can configure quotas to recognize that a user's UNIX id and Windows ID represent the same user.

Quotas for Windows user name are mapped to a UNIX user name, or vice versa, when both of the following conditions are met:

- The `user-mapping` parameter is set to "on" in the quota rule for the user.
- The user names have been mapped with the `vserver name-mapping` commands.

When a UNIX and Windows name are mapped together, they are treated as the same person for determining quota usage.

How tree quotas work

Overview of how tree quotas work

You can create a quota with a qtree as its target to limit how large the target qtree can become. These quotas are also called *tree quotas*.



You can also create user and group quotas for a specific qtree. In addition, quotas for a FlexVol volume are sometimes inherited by the qtrees contained by that volume.

When you apply a quota to a qtree, the result is similar to a disk partition, except that you can change the qtree's maximum size at any time by changing the quota. When applying a tree quota, ONTAP limits the disk space and number of files in the qtree, regardless of their owners. No users, including root and members of the BUILTIN\Administrators group, can write to the qtree if the write operation causes the tree quota to be exceeded.

The size of the quota does not guarantee any specific amount of available space. The size of the quota can be

larger than the amount of free space available to the qtree. You can use the `volume quota report` command to determine the true amount of available space in the qtree.

Learn more about `volume quota report` in the [ONTAP command reference](#).

How user and group quotas work with qtrees

Tree quotas limit the overall size of the qtree. To prevent individual users or groups from consuming the entire qtree, you specify a user or group quota for that qtree.

Example user quota in a qtree

Suppose you have the following quota rules:

```
cluster1::> volume quota policy rule show -vserver vs0 -volume vol1

Vserver: vs0          Policy: default          Volume: vol1
                                         Soft          Soft
                                         User          Disk          Disk          Files          Files
                                         Mapping       Limit        Limit        Limit        Limit
Type   Target     Qtree      User          Disk          Disk          Files          Files
Threshold
-----
-----
user   ""         ""         off          50MB        -           -           -
45MB
user   jsmith     ""         off          80MB        -           -           -
75MB
```

You notice that a certain user, `kjones`, is taking up too much space in a critical qtree, `proj1`, which resides in `vol1`. You can restrict this user's space by adding the following quota rule:

```
cluster1::> volume quota policy rule create -vserver vs0 -volume vol1  
-policy-name default -type user -target "kjones" -qtree "proj1" -disk  
-limit 20m -threshold 15m
```

```
cluster1::> volume quota policy rule show -vserver vs0 -volume vol1
```

Vserver: vs0			Policy: default			Volume: vol1		
Type	Target	Qtree	User Mapping	Disk Limit	Disk Limit	Files Limit	Files Limit	
Threshold								
-----	-----	-----	-----	-----	-----	-----	-----	
user	""	""	off	50MB	-	-	-	
45MB								
user	jsmith	""	off	80MB	-	-	-	
75MB								
user	kjones	proj1	off	20MB	-	-	-	
15MB								

How default tree quotas on a FlexVol volume create derived tree quotas

When you create a default tree quota on a FlexVol volume, corresponding derived tree quotas are automatically created for every qtree in that volume.

These derived tree quotas have the same limits as the default tree quota. If no additional quotas exist, the limits have the following effects:

- Users can use as much space in a qtree as they are allotted for the entire volume (provided they did not exceed the limit for the volume by using space in the root or another qtree).
- Each of the qtrees can grow to consume the entire volume.

The existence of a default tree quota on a volume continues to affect all new qtrees that are added to the volume. Each time a new qtree is created, a derived tree quota is also created.

Like all derived quotas, derived tree quotas display the following behaviors:

- Are created only if the target does not already have an explicit quota.
- Appear in quota reports but do not appear when you show quota rules with the `volume quota policy rule show` command. Learn more about `volume quota policy rule show` in the [ONTAP command reference](#).

Example of derived tree quotas

You have a volume with three qtrees (proj1, proj2, and proj3) and the only tree quota is an explicit quota on the proj1 qtree limiting its disk size to 10 GB. If you create a default tree quota on the volume and reinitialize quotas on the volume, the quota report now contains four tree quotas:

Volume Specifier	Tree Specifier	Type	ID	----Disk----		----Files----		Quota
				Used	Limit	Used	Limit	
vol1	proj1	tree	1	0B	10GB	1	-	proj1
vol1		tree	*	0B	20GB	0	-	*
vol1	proj2	tree	2	0B	20GB	1	-	proj2
vol1	proj3	tree	3	0B	20GB	1	-	proj3
...								

The first line shows the original explicit quota on the proj1 qtree. This quota remains unchanged.

The second line shows the new default tree quota on the volume. The asterisk (*) Quota Specifier indicates it is a default quota. This quota is a result of the quota rule that you created.

The last two lines show new derived tree quotas for the proj2 and proj3 qtrees. ONTAP automatically created these quotas as a result of the default tree quota on the volume. These derived tree quotas have the same 20-GB disk limit as the default tree quota on the volume. ONTAP did not create a derived tree quota for the proj1 qtree because the proj1 qtree already had an explicit quota.

How default user quotas on a FlexVol volume affect quotas for the qtrees in that volume

If a default user quota is defined for a FlexVol volume, a default user quota is automatically created for every qtree contained by that volume for which an explicit or derived tree quota exists.

If a default user quota on the qtree already exists, it remains unaffected when the default user quota on the volume is created.

The automatically created default user quotas on the qtrees have the same limits as the default user quota you create for the volume.

An explicit user quota for a qtree overrides (replaces the limits applied by) the automatically created default user quota, the same way as it overrides a default user quota on that qtree that was created by an administrator.

How qtree changes affect quotas

When you delete, rename, or change the security style of a qtree, the quotas applied by ONTAP might change, depending on the current quotas being applied.

Qtree deletions and tree quotas

When you delete a qtree, all quotas applicable to that qtree, whether they are explicit or derived, are no longer applied by ONTAP.

Whether the quota rules persist depends on where you delete the qtree:

- If you delete a qtree using ONTAP, the quota rules for that qtree are automatically deleted, including tree

quota rules and any user and group quota rules configured for that qtree.

- If you delete a qtree using your CIFS or NFS client, you must delete any quota rules for that qtree to avoid getting errors when you reinitialize quotas. If you create a new qtree with the same name as the one you deleted, the existing quota rules are not applied to the new qtree until you reinitialize quotas.

How renaming a qtree affects quotas

When you rename a qtree using ONTAP, the quota rules for that qtree are automatically updated. If you rename a qtree using your CIFS or NFS client, you must update any quota rules for that qtree.



If you rename a qtree using your CIFS or NFS client and do not update quota rules for that qtree with the new name before you reinitialize quotas, quotas will not be applied to the qtree. Explicit quotas for the qtree, including tree quotas and user or group quotas for the qtree, might be converted into derived quotas.

Qtree security styles and user quotas

You can apply Access Control Lists (ACLs) on qtrees by using NTFS or mixed security styles, but not by using the UNIX security style. Changing the security style of a qtree might affect how quotas are calculated. You should always reinitialize quotas after you change the security style of a qtree.

If you change the security style of a qtree from NTFS or mixed to UNIX, any ACLs on files in that qtree are ignored and the file usage is charged against the UNIX user IDs.

If you change the security style of a qtree from UNIX to either mixed or NTFS, the previously hidden ACLs become visible. In addition, any ACLs that were ignored become effective again, and the NFS user information is ignored. If no ACL existed before, the NFS information continues to be used in the quota calculation.



To make sure that quota usages for both UNIX and Windows users are properly calculated after you change the security style of a qtree, you must reinitialize quotas for the volume containing that qtree.

Example

The following example shows how a change in the security style of a qtree results in a different user being charged for the usage of a file in the particular qtree.

Suppose NTFS security is in effect on qtree A, and an ACL gives Windows user `corp\joe` ownership of a 5MB file. User `corp\joe` is charged with 5MB of disk space usage for qtree A.

Now you change the security style of qtree A from NTFS to UNIX. After quotas are reinitialized, Windows user `corp\joe` is no longer charged for this file; instead, the UNIX user corresponding to the UID of the file is charged for the file. The UID could be a UNIX user mapped to `corp\joe` or the root user.

How quotas are activated

Overview of how quotas are activated

New quotas and changes to existing quotas must be activated to take effect. The activation is performed at the volume level. Knowing how quota activation works can help you manage your quotas with less disruption.

Quotas are activated either by *initializing* (turning them on) or by *resizing*. Turning off quotas and turning them

on again is called reinitializing.

The length of the activation process and its impact on quota enforcement depends on the type of activation:

- The initialization process involves two parts: a quota on job and a quota scan of the volume's entire file system. The scan begins after the quota on job completes successfully. The quota scan can take some time; the more files that the volume has, the longer it takes. Until the scan is finished, quota activation is not complete and quotas are not enforced.
- The resize process involves only a quota resize job. Resizing takes less time than a quota initialization because it does not involve a quota scan. During a resize process, quotas continue to be enforced.

By default, the quota on and quota resize jobs run in the background, which permits you to use other commands at the same time.

Errors and warnings from the activation process are sent to the event management system. If you use the -foreground parameter with the volume quota on or volume quota resize commands, the command does not return until the job is complete; this is useful if you are reinitializing from a script. To display errors and warnings later, you can use the volume quota show command with the -instance parameter.

Quota activation persists across halts and reboots. The process of quota activation does not affect the availability of the storage system data.

Related information

- [volume quota on](#)
- [volume quota resize](#)
- [volume quota show](#)

Understand when to use resizing

Quota resizing is a useful ONTAP feature. And because resizing is faster than quota initialization, you should use resizing whenever possible. However there are a few restrictions you need to be aware of.

Resizing only works for certain types of quota changes. You can resize quotas when making the following types of changes to the quota rules:

- Changing an existing quota.

For example, changing the limits of an existing quota.

- Adding a quota for a quota target for which a default quota or a default tracking quota exists.
- Deleting a quota for which a default quota or default tracking quota entry is specified.
- Combining separate user quotas into one multi-user quota.



After you have made extensive quota changes, you should perform a full reinitialization to ensure that all of the changes take effect.

 If you attempt to resize and not all of your quota changes can be incorporated by using a resize operation, ONTAP issues a warning. You can determine from the quota report whether your storage system is tracking disk usage for a particular user, group, or qtree. If you see a quota in the quota report, it means that the storage system is tracking the disk space and the number of files owned by the quota target.

Example quotas changes that can be made effective by resizing

Some quota rule changes can be made effective by resizing. Consider the following quotas:

#Quota	Target	type	disk	files	thold	sdisk	sfile
*	user@/vol/vol2		50M	15K			
*	group@/vol/vol2		750M	85K			
*	tree@/vol/vol2		-	-			
jdoe	user@/vol/vol2/		100M	75K			
kbuck	user@/vol/vol2/		100M	75K			

Suppose you make the following changes:

- Increase the number of files for the default user target.
- Add a new user quota for a new user, boris, that needs more disk limit than the default user quota.
- Delete the kbuck user's explicit quota entry; the new user now needs only the default quota limits.

These changes result in the following quotas:

#Quota	Target	type	disk	files	thold	sdisk	sfile
*	user@/vol/vol2		50M	25K			
*	group@/vol/vol2		750M	85K			
*	tree@/vol/vol2		-	-			
jdoe	user@/vol/vol2/		100M	75K			
boris	user@/vol/vol2/		100M	75K			

Resizing activates all of these changes; a full quota reinitialization is not necessary.

When a full quota reinitialization is required

Although resizing quotas is faster, you must do a full quota reinitialization if you make certain small or extensive changes to your quotas.

A full quota reinitialization is necessary in the following circumstances:

- You create a quota for a target that has not previously had a quota (neither an explicit quota nor one derived from a default quota).
- You change the security style of a qtree from UNIX to either mixed or NTFS.
- You change the security style of a qtree from mixed or NTFS to UNIX.

- You remove users from a quota target with multiple users, or add users to a target that already has multiple users.
- You make extensive changes to your quotas.

Example of quotas changes that require initialization

Suppose you have a volume that contains three qtrees and the only quotas in the volume are three explicit tree quotas. You decide to make the following changes:

- Add a new qtree and create a new tree quota for it.
- Add a default user quota for the volume.

Both of these changes require a full quota initialization. Resizing does not make the quotas effective.

How you can view quota information

Overview of viewing quota information

You can use quota reports to view details such as the configuration of quota rules and policies, enforced and configured quotas, and errors that have occurred during quota resizing and reinitialization.

Viewing quota information is useful in situations such as the following:

- Configuring quotas, for example to configure quotas and verify the configurations
- Responding to notifications that disk space or file limits will soon be reached or that they have been reached
- Responding to requests for more space

See what quotas are in effect using the quota report

Because of the various ways that quotas interact, more quotas are in effect than just the ones you have explicitly created. To see what quotas are in effect, you can view the quota report.

The following examples show quota reports for different types of quotas applied on a FlexVol volume vol1, and a qtree q1 contained in that volume:

Example with no user quotas specified for the qtree

In this example, there is one qtree, q1, which is contained by the volume vol1. The administrator has created three quotas:

- A default tree quota limit on vol1 of 400MB
- A default user quota limit on vol1 of 100MB
- An explicit user quota limit on vol1 of 200MB for the user jsmith

The quota rules for these quotas look similar to the following example:

```
cluster1::*> volume quota policy rule show -vserver vs1 -volume vol1
```

Vserver: vs1			Policy: default			Volume: vol1		
Type	Target	Qtree	User Mapping	Disk Limit	Disk Limit	Files Limit	Files Limit	
Threshold								
tree	""	""	-	400MB	-	-	-	
-								
user	""	""	off	100MB	-	-	-	
-								
user	jsmith	""	off	200MB	-	-	-	
-								

The quota report for these quotas looks similar to the following example:

```
cluster1::> volume quota report
```

```
Vserver: vs1
```

Volume	Tree Specifier	Type	ID	----Disk----		----Files----		Quota
				Used	Limit	Used	Limit	
vol1	-	tree	*	0B	400MB	0	-	*
vol1	-	user	*	0B	100MB	0	-	*
vol1	-	user	jsmith	150B	200MB	7	-	jsmith
vol1	q1	tree	1	0B	400MB	6	-	q1
vol1	q1	user	*	0B	100MB	0	-	
vol1	q1	user	jsmith	0B	100MB	5	-	
vol1	-	user	root	0B	0MB	1	-	
vol1	q1	user	root	0B	0MB	8	-	

The first three lines of the quota report display the three quotas specified by the administrator. Since two of these quotas are default quotas, ONTAP automatically creates derived quotas.

The fourth line displays the tree quota that is derived from the default tree quota for every qtree in vol1 (in this example, only q1).

The fifth line displays the default user quota that is created for the qtree as a result of the existence of the default user quota on the volume and the qtree quota.

The sixth line displays the derived user quota that is created for jsmith on the qtree because there is a default user quota for the qtree (line 5) and the user jsmith owns files on that qtree. Note that the limit applied to the user jsmith in the qtree q1 is not determined by the explicit user quota limit (200MB). This is because the

explicit user quota limit is on the volume, so it does not affect limits for the qtree. Instead, the derived user quota limit for the qtree is determined by the default user quota for the qtree (100MB).

The last two lines display more user quotas that are derived from the default user quotas on the volume and on the qtree. A derived user quota was created for the root user on both the volume and the qtree because the root user owned files on both the volume and the qtree. Since the root user gets special treatment in terms of quotas, its derived quotas are tracking quotas only.

Example with user quotas specified for the qtree

This example is similar to the previous one, except that the administrator has added two quotas on the qtree.

There is still one volume, vol1, and one qtree, q1. The administrator has created the following quotas:

- A default tree quota limit on vol1 of 400MB
- A default user quota limit on vol1 of 100MB
- An explicit user quota limit on vol1 for the user jsmith of 200MB
- A default user quota limit on qtree q1 of 50MB
- An explicit user quota limit on qtree q1 for the user jsmith of 75MB

The quota rules for these quotas look like this:

Vserver: vs1			Policy: default			Volume: vol1		
Type	Target	Qtree	User Mapping	Disk Limit	Disk Limit	Files Limit	Files Limit	
<i>Threshold</i>								
tree	""	""	-	400MB	-	-	-	
-								
user	""	""	off	100MB	-	-	-	
-								
user	""	q1	off	50MB	-	-	-	
-								
user	jsmith	""	off	200MB	-	-	-	
-								
user	jsmith	q1	off	75MB	-	-	-	
-								

The quota report for these quotas looks like this:

cluster1::> volume quota report									
Vserver: vs1					----Disk----		----Files----		Quota
Volume Specifier	Tree	Type	ID		Used	Limit	Used	Limit	
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----
vol1	-	tree	*		0B	400MB	0	-	*
vol1	-	user	*		0B	100MB	0	-	*
vol1	-	user	jsmith	2000B	200MB		7	-	jsmith
vol1	q1	user	*		0B	50MB	0	-	*
vol1	q1	user	jsmith		0B	75MB	5	-	jsmith
vol1	q1	tree	1		0B	400MB	6	-	q1
vol1	-	user	root		0B	0MB	2	-	
vol1	q1	user	root		0B	0MB	1	-	

The first five lines of the quota report display the five quotas created by the administrator. Since some of these quotas are default quotas, ONTAP automatically creates derived quotas.

The sixth line displays the tree quota that is derived from the default tree quota for every qtree in vol1 (in this example, only q1).

The last two lines display the user quotas that are derived from the default user quotas on the volume and on the qtree. A derived user quota was created for the root user on both the volume and the qtree because the root user owned files on both the volume and the qtree. Since the root user gets special treatment in terms of quotas, its derived quotas are tracking quotas only.

No other default quotas or derived quotas were created for the following reasons:

- A derived user quota was not created for the jsmith user even though the user owns files on both the volume and the qtree because the user already has explicit quotas at both levels.
- No derived user quotas were created for other users because no other users own files on either the volume or the qtree.
- The default user quota on the volume did not create a default user quota on the qtree because the qtree already had a default user quota.

Why enforced quotas differ from configured quotas

Enforced quotas differ from configured quotas because derived quotas are enforced without being configured but configured quotas are enforced only after they are successfully initialized. Understanding these differences can help you compare the enforced quotas that are shown in quota reports to the quotas that you configured.

Enforced quotas, which appear in quota reports, might differ from the configured quota rules for the following reasons:

- Derived quotas are enforced without being configured as quota rules. ONTAP creates derived quotas automatically in response to default quotas.

- Quotas might not have been reinitialized on a volume after quota rules were configured.
- Errors might have occurred when quotas were initialized on a volume.

Use the quota report to determine which quotas limit writes to a specific file

You can use the volume quota report command with a specific file path to determine which quota limits affect write operations to a file. This can help you understand which quota is preventing a write operation.

Steps

1. Use the volume quota report command with the -path parameter.

Example of showing quotas affecting a specific file

The following example shows the command and output to determine what quotas are in effect for writes to the file file1, which resides in the qtree q1 in the FlexVol volume vol2:

```
cluster1:> volume quota report -vserver vs0 -volume vol2 -path
/vol/vol2/q1/file1
Virtual Server: vs0
                                         ----Disk----  ----Files-----  Quota
Volume  Tree      Type     ID      Used    Limit    Used    Limit
Specifier
-----  -----  -----  -----  -----  -----  -----  -----
-----  -----
vol2    q1        tree    jsmith   1MB    100MB    2     10000   q1
vol2    q1        group   eng      1MB    700MB    2     70000
vol2          group   eng      1MB    700MB    6     70000   *
vol2          user    corp\jsmith
                           1MB    50MB     1      -      *
vol2    q1        user    corp\jsmith
                           1MB    50MB     1      -      *
5 entries were displayed.
```

Commands for displaying information about quotas in ONTAP

You can use commands to display a quota report containing enforced quotas and resource usage, display information about quota state and errors, or about quota policies and quota rules.



You can run the following commands only on FlexVol volumes.

If you want to...	Use this command...
View information about enforced quotas	volume quota report

If you want to...	Use this command...
View resource usage (disk space and number of files) of quota targets	volume quota report
Determine which quota limits are affected when a write to a file is allowed	volume quota report with the -path parameter
Display the quota state, such as on, off, and initializing	volume quota show
View information about quota message logging	volume quota show with the -logmsg parameter
View errors that occur during quota initialization and resizing	volume quota show with the -instance parameter
View information about quota policies	volume quota policy show
View information about quota rules	volume quota policy rule show
View the name of the quota policy that is assigned to a storage virtual machine (SVM, formerly known as Vserver)	vserver show with the -instance parameter

Learn more about `volume quota` in the [ONTAP command reference](#).

When to use the `volume quota policy rule show` and `volume quota report` commands

Although both commands show information about quotas, the `volume quota policy rule show` quickly displays configured quota rules while the `volume quota report` command, which consumes more time and resources, displays enforced quotas and resource usage.

The `volume quota policy rule show` command is useful for the following purposes:

- Check the configuration of quota rules before activating them

This command displays all configured quota rules regardless of whether the quotas have been initialized or resized.

- Quickly view quota rules without affecting system resources

Because it does not display disk and file usage, this command is not as resource intensive as a quota report.

- Display the quota rules in a quota policy that is not assigned to the SVM.

Learn more about `volume quota policy rule show` in the [ONTAP command reference](#).

The `volume quota report` command is useful for the following purposes:

- View enforced quotas, including derived quotas
- View the disk space and number of files used by every quota in effect, including targets affected by derived quotas

(For default quotas, the usage appears as "0" because the usage is tracked against the resulting derived quota.)

- Determine which quota limits affect when a write to a file will be allowed

Add the `-path` parameter to the `volume quota report` command.

 The quota report is resource-intensive operation. If you run it on many FlexVol volumes in the cluster, it might take a long time to complete. A more efficient way would be to view the quota report for a particular volume in an SVM.

Learn more about `volume quota report` in the [ONTAP command reference](#).

Difference in space usage displayed by a quota report and a UNIX client

Overview of the difference in space usage displayed by a quota report and a UNIX client

The value of used disk space displayed in a quota report for a FlexVol volume or qtree can be different from the value displayed by a UNIX client for the same volume or qtree. The difference in these values is because of the different methods followed by the quota report and the UNIX commands for calculating the data blocks in the volume or qtree.

For example, if a volume contains a file that has empty data blocks (to which data is not written), the quota report for the volume does not count the empty data blocks while reporting the space usage. However, when the volume is mounted on a UNIX client and the file is shown as the output of the `ls` command, the empty data blocks are also included in the space usage. Therefore, the `ls` command displays a higher file size when compared to the space usage displayed by the quota report.

Similarly, the space usage values shown in a quota report can also differ from the values shown as a result of UNIX commands such as `df` and `du`.

How a quota report accounts for disk space and file usage

The number of files used and the amount of disk space specified in a quota report for a FlexVol volume or a qtree depend on the count of the used data blocks corresponding to every inode in the volume or the qtree.

The block count includes both direct and indirect blocks used for regular and stream files. The blocks used for directories, Access Control Lists (ACLs), stream directories, and metafiles do not get accounted for in the quota report. In case of UNIX sparse files, empty data blocks are not included in the quota report.

The quota subsystem is designed to consider and include only user controllable aspects of the filesystem. Directories, ACLs, and snapshot space are all examples of space excluded from quota calculations. Quotas are used to enforce limits, not guarantees, and they only operate on the active filesystem. Quota accounting does not count certain filesystem constructs, nor does it account for storage efficiency (such as compression or deduplication).

Learn more about the commands described in this procedure in the [ONTAP command reference](#).

Disparity between ls command and quota report for space usage

When you use the `ls` command to view the contents of a FlexVol volume mounted on a UNIX client, the file sizes displayed in the output could differ from the space usage displayed in the quota report for the volume depending on the type of data blocks for the file.

The output of the `ls` command displays only the size of a file and does not include indirect blocks used by the file. Any empty blocks of the file also get included in the output of the command.

Therefore, if a file does not have empty blocks, the size displayed by the `ls` command might be less than the disk usage specified by a quota report because of the inclusion of indirect blocks in the quota report.

Conversely, if the file has empty blocks, then the size displayed by the `ls` command might be more than the disk usage specified by the quota report.

The output of the `ls` command displays only the size of a file and does not include indirect blocks used by the file. Any empty blocks of the file also get included in the output of the command.

Example of the difference between space usage accounted by the ls command and a quota report

The following quota report shows a limit of 10 MB for a qtree q1:

Volume Specifier	Tree	Type	ID	----Disk----		----Files----		Quota
				Used	Limit	Used	Limit	
vol1	q1	tree	user1	10MB	10MB	1	-	q1
...								

A file present in the same qtree can have a size exceeding the quota limit when viewed from a UNIX client by using the `ls` command, as shown in the following example:

```
[user1@lin-sys1 q1]$ ls -lh
-rwxr-xr-x 1 user1 nfsuser **27M** Apr 09 2013 file1
```

Learn more about `ls` in the [ONTAP command reference](#).

How the df command accounts for file size

The way in which the `df` command reports the space usage depends on two conditions: whether the quotas are enabled or disabled for the volume that contains the qtree, and if quota usage within the qtree is tracked.

When quotas are enabled for the volume that contains the qtree and quota usage within the qtree is tracked,

the space usage reported by the `df` command equals the value specified by the quota report. In this situation, quota usage excludes blocks used by directories, ACLs, stream directories, and metafiles.

When quotas are not enabled on the volume, or when the qtree does not have a quota rule configured, the reported space usage includes blocks used by directories, ACLs, stream directories, and metafiles for the entire volume, including other qtrees within the volume. In this situation, the space usage reported by the `df` command is greater than the expected value reported when quotas are tracked.

When you run the `df` command from the mount point of a qtree for which quota usage is tracked, the command output shows the same space usage as the value specified by the quota report. In most cases, when the tree quota rule has a hard disk-limit, the total size reported by the `df` command equals the disk limit and the space available equals the difference between the quota disk limit and quota usage.

However, in some cases, the space available reported by the `df` command might equal the space available in the volume as a whole. This can occur when there is no hard disk limit configured for the qtree. Beginning with ONTAP 9.9.1, it can also occur when the space available in the volume as a whole is less than the remaining tree quota space. When either of these conditions occur, the total size reported by the `df` command is a synthesized number equal to the quota used within the qtree plus the space available in the FlexVol volume.



This total size is neither the qtree disk limit nor the volume configured size. It can also vary based on your write activity within other qtrees or on your background storage efficiency activity.

Example of space usage accounted by the `df` command and a quota report

The following quota report shows a disk limit of 1 GB for qtree alice, 2 GB for qtree bob, and no limit for qtree project1:

```
C1_vsim1::> quota report -vserver vs0
Vserver: vs0
          -----Disk-----  -----Files-----  Quota
Volume  Tree   Type   ID   Used   Limit   Used   Limit
Specifier
-----  -----  -----  -----  -----  -----  -----  -----
-----  -----
vol2    alice   tree   1           502.0MB  1GB     2      -    alice
vol2    bob     tree   2           1003MB   2GB     2      -    bob
vol2    project1 tree   3           200.8MB  -        2      -    -
project1
vol2              tree   *           0B      -        0      -    *
4 entries were displayed.
```

In the following example, the output of the `df` command on qtrees alice and bob reports the same used space as the quota report, and the same total size (in terms of 1M blocks) as the disk limit. This is because the quota rules for qtrees alice and bob have a defined disk limit and the volume available space (1211 MB) is greater than the tree quota space remaining for qtree alice (523 MB) and qtree bob (1045 MB).

```
linux-client1 [~]$ df -m /mnt/vol2/alice
Filesystem           1M-blocks  Used  Available  Use%  Mounted on
172.21.76.153:/vol2      1024    502        523  50%  /mnt/vol2
```

```
linux-client1 [~]$ df -m /mnt/vol2/bob
Filesystem           1M-blocks  Used  Available  Use%  Mounted on
172.21.76.153:/vol2      2048   1004        1045  50%  /mnt/vol2
```

In the following example, the output of the `df` command on qtree project1 reports the same used space as the quota report, but the total size is synthesized by adding the available space in the volume as a whole (1211 MB) to the quota usage of qtree project1 (201 MB) to give a total of 1412 MB. This is because the quota rule for qtree project1 has no disk limit.

```
linux-client1 [~]$ df -m /mnt/vol2/project1
Filesystem           1M-blocks  Used  Available  Use%  Mounted on
172.21.76.153:/vol2      1412    201        1211  15%  /mnt/vol2
```

The following example shows how the output of the `df` command on the volume as a whole reports the same available space as project1.

 linux-client1 [~]\$ df -m /mnt/vol2
Filesystem 1M-blocks Used Available Use% Mounted on
172.21.76.153:/vol2 2919 1709 1211 59% /mnt/vol2

Learn more about the commands described in this procedure in the [ONTAP command reference](#).

Disparity between `du` command and quota report for space usage

When you run the `du` command to check the disk space usage for a qtree or FlexVol volume mounted on a UNIX client, the usage value might be higher than the value displayed by a quota report for the qtree or volume.

The output of the `du` command contains the combined space usage of all the files through the directory tree beginning at the level of the directory where the command is issued. Because the usage value displayed by the `du` command also includes the data blocks for directories, it is higher than the value displayed by a quota report.

Example of the difference between space usage accounted by the `du` command and a quota report

The following quota report shows a limit of 10MB for a qtree q1:

Volume Specifier	Tree	Type	ID	----Disk----		----Files----		Quota
				Used	Limit	Used	Limit	
vol1	q1	tree	user1	10MB	10MB	1	-	q1
...								

In the following example, the disk space usage as the output of the `du` command shows a higher value than exceeds the quota limit:

```
[user1@lin-sys1 q1]$ du -sh
**11M**      q1
```

Learn more about the commands described in this procedure in the [ONTAP command reference](#).

Examples of quota configuration

These examples help you understand how to configure quotas and read quota reports.

About these examples

For the following examples, assume that you have a storage system that includes an SVM, `vs1`, with one volume, `vol1`.

1. To start setting up quotas, you create a new quota policy for the SVM:

```
cluster1::>volume quota policy create -vserver vs1 -policy-name
quota_policy_vs1_1
```

2. Because the quota policy is new, you assign it to the SVM:

```
cluster1::>vserver modify -vserver vs1 -quota-policy quota_policy_vs1_1
```

Example 1: Default user quota

1. You decide to impose a hard limit of 50MB for each user in `vol1`:

```
cluster1::>volume quota policy rule create -vserver vs1 -policy-name
quota_policy_vs1_1 -volume vol1 -type user -target "" -disk-limit 50MB
-qtree ""
```

2. To activate the new rule, you initialize quotas on the volume:

```
cluster1::>volume quota on -vserver vs1 -volume vol1 -foreground
```

3. You view the quota report:

```
cluster1::>volume quota report
```

The resulting quota report is similar to the following report:

Vserver: vs1				----Disk----		----Files----		Quota
Volume	Tree Specifier	Type	ID	Used	Limit	Used	Limit	
vol1		user	*	0B	50MB	0	-	*
vol1		user	jsmith	49MB	50MB	37	-	*
vol1		user	root	0B	-	1	-	

The first line shows the default user quota that you created, including the disk limit. Like all default quotas, this default user quota does not display information about disk or file usage. In addition to the quota that was created, two other quotas appear. There is one quota for each user that currently owns files on `vol1`. These additional quotas are user quotas that were derived automatically from the default user quota. The derived user quota for the user `jsmith` has the same 50MB disk limit as the default user quota. The derived user quota for the root user is a tracking quota (without limits).

If any user on the system (other than the root user) tries to perform an action that would use more than 50MB in `vol1` (for example, writing to a file from an editor), the action fails.

Example 2: Explicit user quota overriding a default user quota

1. If you need to provide more space in volume `vol1` to the user `jsmith`, then you enter the following command:

```
cluster1::>volume quota policy rule create -vserver vs1 -policy-name
quota_policy_vs1_1 -volume vol1 -type user -target jsmith -disk-limit
80MB -qtree ""
```

This is an explicit user quota, because the user is explicitly listed as the target of the quota rule.

This is a change to an existing quota limit, because it changes the disk limit of the derived user quota for the user `jsmith` on the volume. Therefore, you do not need to reinitialize quotas on the volume to activate the change.

2. To resize quotas:

```
cluster1::>volume quota resize -vserver vs1 -volume vol1 -foreground
```

Quotas remain in effect while you resize, and the resizing process is short.

The resulting quota report is similar to the following report:

```
cluster1::> volume quota report
Vserver: vs1
                                         ----Disk----  ----Files-----  Quota
Volume   Tree      Type     ID      Used    Limit      Used    Limit
Specifier
-----  -----  -----  -----  -----  -----  -----  -----
-----  -----
vol1      user      *       0B     50MB      0       -      *
vol1      user      jsmith   50MB    80MB     37       -      jsmith
vol1      user      root     0B     -         1       -      -
3 entries were displayed.
```

The second line now shows a disk limit of 80MB and a quota specifier of `jsmith`.

Therefore, `jsmith` can use up to 80MB of space on `vol1` even though all other users are still limited to 50MB.

Example 3: Thresholds

Suppose you want to receive a notification when users reach within 5MB of their disk limits.

1. To create a threshold of 45MB for all users, and a threshold of 75MB for `jsmith`, you change the existing quota rules:

```
cluster1::>volume quota policy rule modify -vserver vs1 -policy
quota_policy_vs1_1 -volume vol1 -type user -target "" -qtree ""
-threshold 45MB
cluster1::>volume quota policy rule modify -vserver vs1 -policy
quota_policy_vs1_1 -volume vol1 -type user -target jsmith -qtree ""
-threshold 75MB
```

Because the sizes of the existing rules are changed, you resize quotas on the volume in order to activate the changes. You wait until the resize process is finished.

2. To see the quota report with thresholds, you add the `-thresholds` parameter to the `volume quota report` command:

```

cluster1::>volume quota report -thresholds
Vserver: vs1
                                         ----Disk----  ----Files-----
Volume   Tree      Type     ID      Used   Limit     Used   Limit   Quota
                                         (Thold)
Specifier
-----
vol1          user     *        0B    50MB     (45MB)   0      -     *
vol1          user     jsmith   59MB   80MB     (75MB)   55     -     jsmith
vol1          user     root     0B     -        (--)     1      -     -
3 entries were displayed.

```

The thresholds appear in parentheses in the Disk Limit column.

Learn more about `volume quota report` in the [ONTAP command reference](#).

Example 4: Quotas on qtrees

Suppose you need to partition some space for two projects. You can create two qtrees, named `proj1` and `proj2`, to accommodate those projects within `vol1`.

Currently, users can use as much space in a qtree as they are allotted for the entire volume (provided they did not exceed the limit for the volume by using space in the root or another qtree). In addition, each of the qtrees can grow to consume the entire volume.

1. If you want to ensure that neither qtree grows beyond 20GB, you can create default tree quota on the volume:

```

cluster1:>>volume quota policy rule create -vserver vs1 -policy-name
quota_policy_vs1_1 -volume vol1 -type tree -target "" -disk-limit 20GB

```



The correct type is `tree`, not `qtree`.

2. Because this is a new quota, you cannot activate it by resizing. You reinitialize quotas on the volume:

```

cluster1:>>volume quota off -vserver vs1 -volume vol1
cluster1:>>volume quota on -vserver vs1 -volume vol1 -foreground

```



You must ensure that you wait for about five minutes before reactivating the quotas on each affected volume, as attempting to activate them almost immediately after running the `volume quota off` command might result in errors. Alternatively, you can run the commands to re-initialize the quotas for a volume from the node that contains the particular volume. Learn more about `volume quota off` in the [ONTAP command reference](#).

Quotas are not enforced during the reinitialization process, which takes longer than the resizing process.

When you display a quota report, it has several new lines. Some lines are for tree quotas and some lines are for derived user quotas.

The following new lines are for the tree quotas:

Volume Specifier	Tree	Type	ID	----Disk----		----Files----		Quota				
				Used	Limit	Used	Limit					
<hr/>												
<hr/>												
<hr/>												
<hr/>												
<hr/>												
vol1		tree	*	0B	20GB	0	-	*				
vol1	proj1	tree	1	0B	20GB	1	-	proj1				
vol1	proj2	tree	2	0B	20GB	1	-	proj2				
<hr/>												

The default tree quota that you created appears in the first new line, which has an asterisk (*) in the ID column. In response to the default tree quota on a volume, ONTAP automatically creates derived tree quotas for each qtree in the volume. These are shown in the lines where `proj1` and `proj2` appear in the Tree column.

The following new lines are for derived user quotas:

Volume Specifier	Tree	Type	ID	----Disk----		----Files----		Quota				
				Used	Limit	Used	Limit					
<hr/>												
<hr/>												
<hr/>												
<hr/>												
<hr/>												
vol1	proj1	user	*	0B	50MB	0	-					
vol1	proj1	user	root	0B	-	1	-					
vol1	proj2	user	*	0B	50MB	0	-					
vol1	proj2	user	root	0B	-	1	-					
<hr/>												

Default user quotas on a volume are automatically inherited for all qtrees contained by that volume, if quotas are enabled for qtrees. When you added the first qtree quota, you enabled quotas on qtrees. Therefore, derived default user quotas were created for each qtree. These are shown in the lines where ID is asterisk (*).

Because the root user is the owner of a file, when default user quotas were created for each of the qtrees,

special tracking quotas were also created for the root user on each of the qtrees. These are shown in the lines where ID is root.

Example 5: User quota on a qtree

1. You decide to limit users to less space in the `proj1` qtree than they get in the volume as a whole. You want to keep them from using any more than 10MB in the `proj1` qtree. Therefore, you create a default user quota for the qtree:

```
cluster1::>volume quota policy rule create -vserver vs1 -policy-name
quota_policy_vs1_1 -volume vol1 -type user -target "" -disk-limit 10MB
-qtree proj1
```

This is a change to an existing quota, because it changes the default user quota for the `proj1` qtree that was derived from the default user quota on the volume. Therefore, you activate the change by resizing quotas. When the resize process is complete, you can view the quota report.

The following new line appears in the quota report showing the new explicit user quota for the qtree:

Volume Specifier	Tree	Type	ID	----Disk----		----Files----		Quota
				Used	Limit	Used	Limit	
vol1	proj1	user	*	0B	10MB	0	-	*

However, the user `jsmith` is being prevented from writing more data to the `proj1` qtree because the quota you created to override the default user quota (to provide more space) was on the volume. As you have added a default user quota on the `proj1` qtree, that quota is being applied and limiting all the users' space in that qtree, including `jsmith`.

2. To provide more space to the user `jsmith`, you add an explicit user quota rule for the qtree with an 80MB disk limit to override the default user quota rule for the qtree:

```
cluster1::>volume quota policy rule create -vserver vs1 -policy-name
quota_policy_vs1_1 -volume vol1 -type user -target jsmith -disk-limit
80MB -qtree proj1
```

Because this is an explicit quota for which a default quota already existed, you activate the change by resizing quotas. When the resize process is complete, you display a quota report.

The following new line appears in the quota report:

Volume Specifier	Tree	Type	ID	----Disk----		----Files----		Quota
				Used	Limit	Used	Limit	
vol1	proj1	user	jsmith	61MB	80MB	57	-	jsmith

The final quota report is similar to the following report:

```
cluster1::>volume quota report
Vserver: vs1
Volume   Tree   Type   ID   Used   Limit   Used   Limit   Quota
Specifier
-----  -----  -----  -----  -----  -----  -----  -----  -----
-----  -----
vol1          tree   *      0B    20GB    0      -      *      *
vol1          user   *      0B    50MB    0      -      *      *
vol1          user   jsmith 70MB   80MB   65      -      jsmith
vol1      proj1  tree   1      0B    20GB    1      -      proj1
vol1      proj1  user   *      0B    10MB    0      -      *
vol1      proj1  user   root   0B    -       1      -      -
vol1      proj2  tree   2      0B    20GB    1      -      proj2
vol1      proj2  user   *      0B    50MB    0      -      -
vol1      proj2  user   root   0B    -       1      -      -
vol1          user   root   0B    -       3      -      -
vol1      proj1  user   jsmith 61MB   80MB   57      -      jsmith
11 entries were displayed.
```

The user `jsmith` is required to meet the following quota limits to write to a file in `proj1`:

1. The tree quota for the `proj1` qtree.
2. The user quota on the `proj1` qtree.
3. The user quota on the volume.

Set up quotas on an SVM

You can set up quotas on a new SVM to management and monitor resource utilization.

About this task

At a high level, there several steps involved when configuring quotas including:

1. Create a quota policy
2. Add the quota rules to the policy

3. Assign the policy to the SVM
4. Initialize the quotas on each FlexVol volume on the SVM

Steps

1. Enter the command `vserver show -instance` to display the name of the default quota policy that was automatically created when the SVM was created.

If a name was not specified when the SVM was created, the name is "default". You can use the `vserver quota policy rename` command to give the default policy a name.



You can also create a new policy by using the `volume quota policy create` command.

2. Use the `volume quota policy rule create` command to create *any* of the following quota rules for each volume on the SVM:
 - Default quota rules for all users
 - Explicit quota rules for specific users
 - Default quota rules for all groups
 - Explicit quota rules for specific groups
 - Default quota rules for all qtrees
 - Explicit quota rules for specific qtrees
3. Use the `volume quota policy rule show` command to check that the quota rules are configured correctly.
4. If you are working on a new policy, use the `vserver modify` command to assign the new policy to the SVM.
5. Use the `volume quota on` command to initialize the quotas on each volume on the SVM.

You can monitor the initialization process in the following ways:

- When you use the `volume quota on` command, you can add the `-foreground` parameter to run the quota on job in the foreground. (By default, the job runs in the background.)

When the job runs in the background, you can monitor its progress by using the `job show` command.

- You can use the `volume quota show` command to monitor the status of the quota initialization.

6. Use the `volume quota show -instance` command to check for initialization errors, such as quota rules that failed to initialize.
7. Use the `volume quota report` command to display a quota report so that you can ensure the enforced quotas match your expectations.

Related information

- [vserver show](#)
- [vserver modify](#)
- [job show](#)
- [volume quota](#)

Modify or resize quota limits

You can change or resize the quotas on all affected volumes, which is faster than reinitializing quotas on those volumes.

About this task

You have a storage virtual machine (SVM, formerly known as Vserver) with enforced quotas and you want either to change the size limits of existing quotas or to add or delete quotas for targets that already have derived quotas.

Steps

1. Use the `vserver show` command with the `-instance` parameter to determine the name of the policy that is currently assigned to the SVM.
2. Modify quota rules by performing any of the following actions:
 - Use the `volume quota policy rule modify` command to modify the disk or file limits of existing quota rules.
 - Use the `volume quota policy rule create` command to create explicit quota rules for targets (users, groups, or qtrees) that currently have derived quotas.
 - Use the `volume quota policy rule delete` command to delete explicit quota rules for targets (users, groups, or qtrees) that also have default quotas.
3. Use the `volume quota policy rule show` command to check that the quota rules are configured correctly.
4. Use the `volume quota resize` command on each volume where you changed quotas, to activate the changes on each volume.

You can monitor the resize process in either of the following ways:

- When you use the `volume quota resize` command, you can add the `-foreground` parameter to run the resize job in the foreground. (By default, the job runs in the background.)

When the job runs in the background, you can monitor its progress by using the `job show` command.

5. Use the `volume quota show -instance` command to check for resize errors such as, quota rules that failed to get resized.

In particular, check for “new definition” errors, which occur when you resize quotas after adding an explicit quota for a target that does not already have a derived quota.

6. Use the `volume quota report` command to display a quota report so that you can ensure the enforced quotas match your requirements.

Related information

- [volume quota policy rule](#)
- [volume quota](#)
- [job show](#)

Reinitialize quotas after making extensive changes

After you make extensive changes to existing quota definitions, you must re-initialize the quotas on all affected volumes. An example of this type of change is adding or deleting quotas for targets that have no enforced quotas.

About this task

You have a storage virtual machine (SVM) with enforced quotas and you want to make changes that require a full reinitialization of quotas.

Steps

1. Use the `vserver show` command with the `-instance` parameter to determine the name of the policy that is currently assigned to the SVM.
2. Modify quota rules by performing any of the following actions:

If you want to...	Then...
Create new quota rules	Use the <code>volume quota policy rule create</code> command
Modify the settings of existing quota rules	Use the <code>volume quota policy rule modify</code> command
Delete existing quota rules	Use the <code>volume quota policy rule delete</code> command

3. Use the `volume quota policy rule show` command to check that the quota rules are configured correctly.
4. Re-initialize quotas on each volume where you changed quotas by turning quotas off and then turning quotas on for those volumes.
 - a. Use the `volume quota off` command on each affected volume to deactivate quotas on that volume.
 - b. Use the `volume quota on` command on each affected volume to activate quotas on that volume.



You must ensure that you wait for about five minutes before reactivating the quotas on each affected volume, as attempting to activate them almost immediately after running the `volume quota off` command might result in errors.

Alternatively, you can run the commands to re-initialize the quotas for a volume from the node that contains the particular volume.

You can monitor the initialization process in either of the following ways:

- When you use the `volume quota on` command, you can add the `-foreground` parameter to run the quota on job in the foreground. (By default, the job runs in the background.)

When the job runs in the background, you can monitor its progress by using the `job show` command.

- You can use the `volume quota show` command to monitor the status of the quota initialization.

5. Use the `volume quota show -instance` command to check for initialization errors, such as quota rules that failed to initialize.
6. Use the `volume quota report` command to display a quota report so that you can ensure the enforced quotas match your expectations.

Related information

- [vserver show](#)
- [volume quota policy rule](#)
- [volume quota](#)
- [job show](#)

Commands to manage quota rules and quota policies

The `volume quota policy rule` commands enable you to configure quota rules, and the `volume quota policy` commands and some `vserver` commands enable you to configure quota policies. Depending on what you need to do, use the following commands to manage quota rules and quota policies:



You can run the following commands only on FlexVol volumes.

Commands for managing quota rules

If you want to...	Use this command...
Create a new quota rule	<code>volume quota policy rule create</code>
Delete an existing quota rule	<code>volume quota policy rule delete</code>
Modify an existing quota rule	<code>volume quota policy rule modify</code>
Display information about configured quota rules	<code>volume quota policy rule show</code>

Commands for managing quota policies

If you want to...	Use this command...
Duplicate a quota policy and the quota rules it contains	<code>volume quota policy copy</code>
Create a new, blank quota policy	<code>volume quota policy create</code>
Delete an existing quota policy that is not currently assigned to a storage virtual machine (SVM)	<code>volume quota policy delete</code>
Rename a quota policy	<code>volume quota policy rename</code>

If you want to...	Use this command...
Display information about quota policies	volume quota policy show
Assign a quota policy to an SVM	vserver modify -quota-policy <i>policy_name</i>
Display the name of the quota policy assigned to an SVM	vserver show

Learn more about the commands described in this procedure in the [ONTAP command reference](#).

Related information

- [volume quota policy](#)
- [vserver modify -quota-policy *policy_name*](#)
- [vserver show](#)

Commands to activate and modify quotas in ONTAP

volume quota commands enable you to change the state of quotas and configure message logging of quotas. Depending on what you need to do, you can use the following commands to activate and modify quotas:

If you want to...	Use this command...
Turn quotas on (also called <i>initializing</i> them)	volume quota on
Resize existing quotas	volume quota resize
Turn quotas off	volume quota off
Change the message logging of quotas, turn quotas on, turn quotas off, or resize existing quotas	volume quota modify

Learn more about the commands described in this procedure in the [ONTAP command reference](#).

Related information

- [volume quota on](#)
- [volume quota resize](#)
- [volume quota off](#)
- [volume quota modify](#)

Use deduplication, data compression, and data compaction to increase storage efficiency

Deduplication, data compression, data compaction, and storage efficiency

You can run deduplication, data compression, and data compaction together or independently to achieve optimal space savings on a FlexVol volume. Deduplication eliminates duplicate data blocks. Data compression compresses the data blocks to reduce the amount of physical storage that is required. Data compaction stores more data in less space to increase storage efficiency.



All inline storage efficiency features, such as inline deduplication and inline compression, are enabled by default on AFF volumes.

Enable deduplication on a volume

You can enable deduplication on a FlexVol volume to achieve storage efficiency. You can enable postprocess deduplication on all volumes and inline deduplication on volumes that reside on AFF or Flash Pool aggregates.

If you want to enable inline deduplication on other types of volumes, see the [NetApp Knowledge Base: How to enable volume inline deduplication on Non-AFF \(All Flash FAS\) aggregates](#).

Before you begin

For a FlexVol volume, you must have verified that enough free space exists for deduplication metadata in volumes and aggregates. The deduplication metadata requires a minimum amount of free space in the aggregate. This amount is equal to 3% of the total amount of physical data for all deduplicated FlexVol volumes or data constituents within the aggregate. Each FlexVol volume or data constituent should have 4% of the total amount of physical data's worth of free space, for a total of 7%.



Inline deduplication is enabled by default on AFF systems.

Choices

- Use the `volume efficiency on` command to enable postprocess deduplication. Learn more about `volume efficiency on` in the [ONTAP command reference](#).

The following command enables postprocess deduplication on volume VolA:

```
volume efficiency on -vserver vs1 -volume VolA
```

- Use the `volume efficiency on` command followed by the `volume efficiency modify` command with the `-inline-deduplication` option set to `true` to enable both postprocess deduplication and inline deduplication. Learn more about `volume efficiency modify` in the [ONTAP command reference](#).

The following commands enable both postprocess deduplication and inline deduplication on volume VolA:

```
volume efficiency on -vserver vs1 -volume VolA
```

```
volume efficiency modify -vserver vs1 -volume VolA -inline-dedupe true
```

- Use the `volume efficiency on` command followed by the `volume efficiency modify` command with the `-inline-deduplication` option set to `true` and the `-policy` option set to `inline-only` to enable only inline deduplication.

The following commands enable only inline deduplication on volume VolA:

```
volume efficiency on -vserver vs1 -volume VolA
```

```
volume efficiency modify -vserver vs1 -volume VolA -policy inline-only -inline
-dedupe true
```

After you finish

Verify that the setting has changed by viewing the volume efficiency settings:

```
volume efficiency show -instance
```

Learn more about `volume efficiency show -instance` in the [ONTAP command reference](#).

Disable deduplication on a volume

You can disable postprocess deduplication and inline deduplication independently on a volume.

Before you begin

Stop any volume efficiency operation that is currently active on the volume: `volume efficiency stop`

Learn more about `volume efficiency stop` in the [ONTAP command reference](#).

About this task

If you have enabled data compression on the volume, running the `volume efficiency off` command disables data compression. Learn more about `volume efficiency off` in the [ONTAP command reference](#).

Choices

- Use the `volume efficiency off` command to disable both postprocess deduplication and inline deduplication.

The following command disable both postprocess deduplication and inline deduplication on volume VolA:

```
volume efficiency off -vserver vs1 -volume VolA
```

- Use the `volume efficiency modify` command with the `-policy` option set to `inline only` to disable postprocess deduplication, but inline deduplication remains enabled.

The following command disables postprocess deduplication, but inline deduplication remains enabled on volume VolA:

```
volume efficiency modify -vserver vs1 -volume VolA -policy inline-only
```

- Use the `volume efficiency modify` command with the `-inline-deduplication` option set to `false` to disable inline deduplication only.

The following command disables only inline deduplication on volume VolA:

```
volume efficiency modify -vserver vs1 -volume VolA -inline-deduplication false
```

Learn more about `volume efficiency modify` in the [ONTAP command reference](#).

Automatic volume-level background deduplication on AFF systems

Beginning with ONTAP 9.3, you can configure volume-level background deduplication to run automatically using a predefined `auto` AFF policy. No manual configuration of the schedules is required. The `auto` policy performs continuous deduplication in the background.

The `auto` policy is set for all newly created volumes and for all upgraded volumes that have not been manually configured for background deduplication. You can [change the policy](#) to default or any other policy to disable the feature.

If a volume moves from a non-AFF system to an AFF system, the `auto` policy is enabled on the destination node by default. If a volume moves from an AFF node to a non-AFF node, the `auto` policy on the destination node is replaced by the `inline-only` policy by default.

On AFF, the system monitors all the volumes having the `auto` policy and deprioritizes the volume that has less savings or has frequent overwrites. The deprioritized volumes no longer participate in automatic background deduplication. Change logging on deprioritized volumes is disabled and metadata on the volume is truncated.

Users can promote the deprioritized volume to re-participate in an automatic background deduplication using the `volume efficiency promote` command available at the advanced privilege level.

Learn more about `volume efficiency promote` in the [ONTAP command reference](#).

Manage aggregate-level inline deduplication on AFF systems

Aggregate-level deduplication eliminates duplicate blocks across volumes belonging to the same aggregate. You can perform aggregate-level deduplication inline on AFF systems. The feature is enabled by default for all newly created volumes and all upgraded volumes with `volume inline deduplication` turned on.

About this task

The deduplication operation eliminates duplicate blocks before data is written to disk. Only volumes with the `space guarantee` set to `none` can participate in aggregate-level inline deduplication. This is the default setting on AFF systems.



Aggregate-level inline deduplication is sometimes referred to as cross-volume inline deduplication.

Step

1. Manage aggregate-level inline deduplication on AFF systems:

If you want to...	Use this command
Enable aggregate-level inline deduplication	<pre>volume efficiency modify -vserver vserver_name -volume vol_name -cross -volume-inline-dedupe true</pre>

If you want to...	Use this command
Disable aggregate-level inline deduplication	volume efficiency modify -vserver vserver_name -volume vol_name -cross-volume-inline-dedupe false
Display aggregate-level inline deduplication status	volume efficiency config -volume vol_name

Examples

The following command displays the aggregate-level inline deduplication status:

```
wfit-8020-03-04::> volume efficiency config -volume choke0_wfit_8020_03_0
Vserver:                                     vs0
Volume:                                       choke0_wfit_8020_03_0
Schedule:                                     -
Policy:                                       choke_VE_policy
Compression:                                  true
Inline Compression:                          true
Inline Dedupe:                               true
Data Compaction:                            true
Cross Volume Inline Deduplication:          false
```

Manage aggregate-level background deduplication on AFF systems

Aggregate-level deduplication eliminates duplicate blocks across volumes belonging to the same aggregate. Beginning with ONTAP 9.3, you can perform aggregate-level deduplication in the background on AFF systems. The feature is enabled by default for all newly created volumes and all upgraded volumes with volume background deduplication turned on.

About this task

The operation is triggered automatically when a large enough percentage of the change log has been populated. No schedule or policy is associated with the operation.

Beginning with ONTAP 9.4, AFF users can also run the aggregate-level deduplication scanner to eliminate duplicates of existing data across volumes in the aggregate. You can use the `storage aggregate efficiency cross-volume-dedupe start` command with the `-scan-old-data=true` option to start the scanner:

```
cluster-1::> storage aggregate efficiency cross-volume-dedupe start
-aggregate aggr1 -scan-old-data true
```

Deduplication scanning can be time consuming. You might want to run the operation in off-peak hours.



Aggregate-level background deduplication is sometimes referred to as cross-volume background deduplication.

Learn more about storage aggregate efficiency cross-volume-dedupe start in the [ONTAP command reference](#).

Steps

1. Manage aggregate-level background deduplication on AFF systems:

If you want to...	Use this command
Enable aggregate-level background deduplication	<code>volume efficiency modify -vserver <vserver_name> -volume <vol_name> -cross-volume-background-dedupe true</code>
Disable aggregate-level background deduplication	<code>volume efficiency modify -vserver <vserver_name> -volume <vol_name> -cross-volume-background-dedupe false</code>
Display aggregate-level background deduplication status	<code>aggregate efficiency cross-volume-dedupe show</code>

Learn more about the commands described in this procedure in the [ONTAP command reference](#).

Related information

- [volume efficiency modify](#)
- [aggregate efficiency cross-volume-dedupe show](#)

Learn about ONTAP temperature-sensitive storage efficiency

ONTAP provides temperature-sensitive storage efficiency (TSSE) benefits by assessing how often your volume's data is accessed and mapping that frequency to the degree of compression applied to that data. For cold data that is accessed infrequently, larger data blocks are compressed, and for hot data, which is accessed frequently and is overwritten more often, smaller data blocks are compressed, making the process more efficient.

TSSE is introduced in ONTAP 9.8 and is enabled automatically on newly created thinly provisioned AFF volumes. You can enable temperature-sensitive storage efficiency on existing thinly provisioned AFF volumes and on thinly provisioned non-AFF DP volumes. TSSE is not supported on thickly provisioned volumes.

Temperature-sensitive storage efficiency is not applied on the following platforms:

Unresolved directive in volumes/enable-temperature-sensitive-efficiency-concept.adoc -
include::_include/dedicated-offload-processor-supported-platforms.adoc[]

These platforms use [CPU or dedicated offload processor storage efficiency](#). Compression is performed using either the main CPU or a dedicated offload processor is not based on hot or cold data.



Over time, the amount of space used in your volume might be more pronounced with TSSE compared to 8K adaptive compression. This behavior is expected due to the architectural differences between TSSE and 8K adaptive compression.

Introduction of "default" and "efficient" modes

Beginning with ONTAP 9.10.1, *default* and *efficient* volume-level storage efficiency modes are introduced for AFF systems only. The two modes provide a choice between file compression (default), which is the default mode when creating new AFF volumes, or temperature-sensitive storage efficiency (efficient), which uses auto adaptive compression to provide increased compression savings on cold, infrequently accessed, data.

When upgrading to ONTAP 9.10.1 and later, existing volumes are assigned a storage efficiency mode based on the type of compression currently enabled on the volumes. During an upgrade, volumes with compression enabled are assigned the default mode, and volumes with temperature-sensitive storage efficiency enabled are assigned the efficient mode. If compression is not enabled, storage efficiency mode remains blank.

With ONTAP 9.10.1, [temperature-sensitive storage efficiency must be explicitly set](#) to enable auto adaptive compression. However, other storage efficiency features like data-compaction, auto dedupe schedule, inline deduplication, cross volume inline deduplication, and cross volume background deduplication are enabled by default on AFF platforms for both default and efficient modes.

Both storage efficiency modes (default and efficient) are supported on FabricPool-enabled aggregates and with all tiering policy types.

Temperature-sensitive storage efficiency enabled on C-Series platforms

Temperature-sensitive storage efficiency is enabled by default on AFF C-Series platforms and when migrating thinly provisioned volumes from a non-TSSE platform to a TSSE-enabled C-Series platform using volume move or SnapMirror with the following releases installed on the destination:

- ONTAP 9.12.1P4 and later
- ONTAP 9.13.1 and later

For more information, see [Storage efficiency behavior with volume move and SnapMirror operations](#).

For existing thinly provisioned volumes, temperature-sensitive storage efficiency is not enabled automatically; however, you can [modify the storage efficiency mode](#) manually to change to efficient mode.



Once you change the storage efficiency mode to efficient you cannot change it back.

Improved storage efficiency with sequential packing of contiguous physical blocks

Beginning with ONTAP 9.13.1, temperature-sensitive storage efficiency adds sequential packing of contiguous physical blocks to further improve storage efficiency. Volumes that have temperature-sensitive storage efficiency enabled automatically have sequential packing enabled when you upgrade systems to ONTAP 9.13.1. After sequential packing is enabled, you must [manually repack existing data](#).

Storage efficiency behavior with volume move and SnapMirror operations

The behavior of storage efficiency can be affected by other storage operations that are active or started at the same time. You should be aware of the impact of these operations on storage efficiency.

There are several situations where storage efficiency on a volume can be affected by other operations including volume moves, SnapMirror relationships, FabricPool volumes, and [temperature-sensitive storage](#)

efficiency (TSSE).

FabricPool

The `all` tiering policy is commonly used on data protection volumes to immediately mark data as cold and tier it as soon as possible. There is no waiting for a minimum number of days to pass before the data is made cold and tiered.

Because the `all` tiering policy tiers data as soon as possible, storage efficiencies that rely on background processes, like 32K efficient adaptive compression (TSSE), do not have enough time to be applied. Inline storage efficiencies like 8K compression are applied as normal.

The following table describes the behavior of a source volume and destination volume when you perform one of these operations.

Source volume efficiency	Destination volume default behavior			Default behavior after manually enabling TSSE (after SnapMirror break)		
	Storage efficiency type	New writes	Cold data compression	Storage efficiency type	New writes	Cold data compression
No storage efficiency (likely FAS)	File compression	File compression is attempted inline on newly written data	No cold data compression, data remains as it is	TSSE with cold data scan algorithm as ZSTD	8k inline compression is attempted in TSSE format	<p>File compressed data: N/A</p> <p>Uncompressed data: 32K compression attempted after threshold days met</p> <p>Newly written data: 32K compression attempted after threshold days met</p>
No storage efficiency (likely FAS)	File compression on C-Series platforms using ONTAP 9.11.1P10 or ONTAP 9.12.1P3	No TSSE-enabled cold data compression	File compressed data: N/A	TSSE with cold data scan algorithm as ZSTD	8K inline compression	<p>File compressed data: N/A</p> <p>Uncompressed data: 32K compression attempted after threshold days met</p> <p>Newly written data: 32K compression attempted after threshold days met</p>

No storage efficiency (likely FAS)	TSSE on C-Series platforms using ONTAP 9.12.1P4 and later or ONTAP 9.13.1 and later	8K inline compression is attempted in TSSE format	File compressed data: N/A Uncompressed data: 32K compression attempted after threshold days met Newly written data: 32K compression attempted after threshold days met	TSSE with cold data scan algorithm as ZSTD	8K inline compression is attempted in TSSE format	File compressed data: N/A Uncompressed data: 32K compression attempted after threshold days met Newly written data: 32K compression attempted after threshold days met
File compression group	Same as source	File compression is attempted inline on newly written data	No cold data compression, data remains as it is	TSSE with cold data scan algorithm as ZSTD	8K inline compression is attempted in TSSE format	File compressed data: Not compressed Uncompressed data: 32K compression is attempted after threshold days met Newly written data: 32K compression is attempted after threshold days met
TSSE cold data scan	TSSE using the same compression algorithm as source volume (LZOPro→LZOPro and ZSTD→ZSTD)	8K inline compression attempted in TSSE format	32K compression attempted with LzoPro after threshold days based coldness is met on both existing data and newly written data.	TSSE is enabled. NOTE: LZOPro cold data scan algorithm can be changed to ZSTD.	8K inline compression is attempted in TSSE format	32K compression is attempted after threshold days coldness is met on both existing data and newly written data.

Set storage efficiency mode during volume creation

Beginning with ONTAP 9.10.1, you can set the storage efficiency mode when creating a new AFF volume.

About this task

You can control the storage efficiency mode on a new AFF volume using the parameter `-storage-efficiency-mode`. You can choose between two options to set the storage efficiency mode: `default` or `efficient`. The storage efficiency mode you choose depends on whether you want greater performance or higher storage efficiency on the volume. The parameter `-storage-efficiency-mode` is not supported on non-AFF volumes or on data protection volumes.

Performance mode is set by default when you create new AFF volumes with storage efficiency.

Learn more about temperature-sensitive storage efficiency and storage efficiency modes.

Steps

1. Create a new volume and set the efficiency mode:

```
volume create -vserver <vserver name> -volume <volume name> -aggregate <aggregate name> -size <volume size> -storage-efficiency-mode <efficient|default>
```

Set `-storage-efficiency-mode` to `efficient` for efficiency mode or to `default` for performance mode.

In the following example, `aff_voll` is created with efficiency mode.

```
volume create -vserver vs1 -volume aff_voll -aggregate aff_aggr1 -storage -efficiency-mode efficient -size 10g
```

Change the volume inactive data compression threshold in ONTAP

You can change how frequently ONTAP performs a cold data scan by modifying the coldness threshold on volumes using temperature-sensitive storage efficiency.

Before you begin

You must be a cluster or SVM administrator and use the ONTAP CLI advanced privilege level.

About this task

The coldness threshold can be from 1 to 60 days. The default threshold is 14 days.

Steps

1. Set the privilege level:

```
set -privilege advanced
```

2. Modify inactive data compression on a volume:

```
volume efficiency inactive-data-compression modify -vserver <vserver_name> -volume <volume_name> -threshold-days <integer>
```

Learn more about `volume efficiency inactive-data-compression modify` in the [ONTAP command reference](#).

Check volume efficiency mode

You can use the `volume-efficiency-show` command on an AFF volume to check whether efficiency is set and to view the current efficiency mode.

Step

1. Check the efficiency mode on a volume:

```
volume efficiency show -vserver <vserver name> -volume <volume name> -fields storage-efficiency-mode
```

Learn more about `volume efficiency show` in the [ONTAP command reference](#).

Change volume efficiency mode

Beginning with ONTAP 9.10.1, the volume-level storage efficiency modes *default* and *efficient* are supported for AFF systems only. These modes provide a choice between file compression (default), which is the default mode when creating new AFF volumes, or temperature-sensitive storage efficiency (efficient), which enables temperature-sensitive storage efficiency (TSSE).



TSSE is supported only on thinly provisioned volumes. [Learn more about TSSE](#).

Steps

You can perform this task using ONTAP System Manager or the ONTAP CLI.

System Manager

Beginning with ONTAP 9.10.1, you can use System Manager to enable higher storage efficiency using the temperature-sensitive storage efficiency feature. Performance-based storage efficiency is enabled by default.

1. Click **Storage > Volumes**.
2. Locate the volume on which you want to enable or disable storage efficiency, and click .
3. Click **Edit > Volumes**, and scroll to **Storage Efficiency**.
4. Select **Enable Higher Storage Efficiency**.

CLI

You can use the `volume efficiency modify` command to change the storage efficiency mode for an AFF volume from `default` to `efficient`, or you can set an efficiency mode when volume efficiency is not already set.

1. Change the volume efficiency mode:

```
volume efficiency modify -vserver <vserver name> -volume <volume name> -storage-efficiency-mode <default|efficient>
```

Learn more about `volume efficiency modify` in the [ONTAP command reference](#).

View volume footprint savings with or without temperature-sensitive storage efficiency

Depending on your ONTAP release, you can view the physical footprint savings on each

volume. You might do this to assess the effectiveness of your administrative processes or as part of capacity planning.

About this task

Beginning with ONTAP 9.11.1, you can use the command `volume show-footprint` to view the physical footprint savings on volumes with temperature-sensitive storage efficiency (TSSE) enabled. Beginning with ONTAP 9.13.1, you can use the same command to view the physical footprint savings on volumes not enabled with TSSE.

Steps

1. View the volume footprint savings:

```
volume show-footprint
```

Example output with TSSE enabled

Feature	Used	Used%
Volume Data Footprint	10.15GB	13%
Volume Guarantee	0B	0%
Flexible Volume Metadata	64.25MB	0%
Delayed Frees	235.0MB	0%
File Operation Metadata	4KB	0%
Total Footprint	10.45GB	13%
Footprint Data Reduction	6.85GB	9%
Auto Adaptive Compression	6.85GB	9%
Effective Total Footprint	3.59GB	5%

Example output without TSSE enabled

Vserver : vs0		
Volume : vol_file_cg_75_per_compress		
Feature	Used	Used%
Volume Data Footprint	5.19GB	7%
Volume Guarantee	0B	0%
Flexible Volume Metadata	32.12MB	0%
Delayed Frees	90.17MB	0%
File Operation Metadata	4KB	0%
Total Footprint	5.31GB	7%
Footprint Data Reduction	1.05GB	1%
Data Compaction	1.05GB	1%
Effective Total Footprint	4.26GB	5%

Related information

- [Set storage efficiency mode during volume creation](#)

Enable data compression on a volume

You can enable data compression on a FlexVol volume to achieve space savings by using the `volume efficiency modify` command. You can also assign a compression type to your volume, if you do not want the default compression type. Learn more about `volume efficiency modify` in the [ONTAP command reference](#).

Before you begin

You must have enabled deduplication on the volume.

 • Deduplication only needs to be enabled and does not need to be running on the volume.

- The compression scanner must be used to compress the existing data on the volumes present in AFF platforms.

Enabling deduplication on a volume

About this task

- In HDD aggregates and Flash Pool aggregates, you can enable both inline and postprocess compression or only postprocess compression on a volume.

If you are enabling both, then you must enable postprocess compression on the volume before enabling inline compression.

- In AFF platforms, only inline compression is supported.

Before enabling inline compression, you must enable postprocess compression on the volume. However,

because postprocess compression is not supported in AFF platforms, no postprocess compression takes place on those volumes and an EMS message is generated informing you that postprocess compression was skipped.

- Temperature sensitive storage efficiency is introduced in ONTAP 9.8. With this feature, storage efficiency is applied according to whether data is hot or cold. For cold data, larger data blocks are compressed, and for hot data, which is overwritten more often, smaller data blocks are compressed, making the process more efficient. Temperature sensitive storage efficiency is enabled automatically on newly created thin-provisioned AFF volumes.
- The compression type is automatically assigned based on the aggregate's platform:

Platform/aggregates	Compression type
AFF	Adaptive compression
Flash Pool aggregates	Adaptive compression
HDD aggregates	Secondary compression

Choices

- Use the `volume efficiency modify` command to enable data compression with the default compression type.

The following command enables postprocess compression on volume VolA of SVM vs1:

```
volume efficiency modify -vserver vs1 -volume VolA -compression true
```

The following command enables both postprocess and inline compression on volume VolA of SVM vs1:

```
volume efficiency modify -vserver vs1 -volume VolA -compression true -inline-compression true
```

- Use the `volume efficiency modify` command at the advanced privilege level to enable data compression with a specific compression type.
 - a. Use the `set -privilege advanced` command to change the privilege level to advanced.
 - b. Use the `volume efficiency modify` command to assign a compression type to a volume.

The following command enables postprocess compression and assigns the adaptive compression type to volume VolA of SVM vs1:

```
volume efficiency modify -vserver vs1 -volume VolA -compression true -compression-type adaptive
```

The following command enables both postprocess and inline compression and assigns the adaptive compression type to volume VolA of SVM vs1:

```
volume efficiency modify -vserver vs1 -volume VolA -compression true -compression-type adaptive -inline-compression true
```

- c. Use the `set -privilege admin` command to change the privilege level to admin.

Move between secondary compression and adaptive compression

You can switch between secondary compression and adaptive compression depending on the amount of data reads. Adaptive compression is preferred when there are a high volume of random reads on the system and higher performance is required. Secondary compression is preferred when data is written sequentially and higher compression savings are required.

About this task

The default compression type is selected based on your aggregates and platform.

Steps

1. Disable efficiency on the volume:

```
volume efficiency off
```

For example, the following command disables efficiency on volume vol1:

```
volume efficiency off -vserver vs1 -volume vol1
```

2. Change to the advanced privilege level:

```
set -privilege advanced
```

3. Decompress the compressed data:

```
volume efficiency undo
```

For example, the following command decompresses the compressed data on volume vol1:

```
volume efficiency undo -vserver vs1 -volume vol1 -compression true
```



You must verify that you have sufficient space in the volume to accommodate the decompressed data.

4. Change to the admin privilege level:

```
set -privilege admin
```

5. Verify that the status of the operation is idle:

```
volume efficiency show
```

For example, the following command displays the status of an efficiency operation on volume vol1:

```
volume efficiency show -vserver vs1 -volume vol1
```

6. Enable efficiency for the volume:

```
volume efficiency on
```

For example, the following command enables efficiency on volume vol1:

```
volume efficiency on -vserver vs1 -volume vol1
```

7. Enable data compression, and then set the type of compression:

```
volume efficiency modify
```

For example, the following command enables data compression and sets the compression type as secondary compression on volume vol1:

```
volume efficiency modify -vserver vs1 -volume vol1 -compression true  
-compression-type secondary
```

This step only enables secondary compression on the volume; the data on the volume is not compressed.



- To compress existing data on AFF systems, you must run the background compression scanner.
- To compress existing data on Flash Pool aggregates or HDD aggregates, you must run the background compression.

8. Optional: Enable inline compression:

```
volume efficiency modify
```

For example, the following command enables inline compression on volume vol1:

```
volume efficiency modify -vserver vs1 -volume vol1 -inline-compression true
```

Disable data compression on a volume

You can disable data compression on a volume by using the `volume efficiency modify` command. Learn more about `volume efficiency modify` in the [ONTAP command reference](#).

About this task

If you want to disable postprocess compression, you must first disable inline compression on the volume.

Steps

1. Stop any volume efficiency operation that is currently active on the volume:

```
volume efficiency stop
```

2. Disable data compression:

```
volume efficiency modify
```

Existing compressed data will remain compressed on the volume. Only new writes coming into the volume are not compressed.

Examples

The following command disables inline compression on volume VolA:

```
volume efficiency modify -vserver vs1 -volume VolA -inline-compression false
```

The following command disables both postprocess compression and inline compression on volume VolA:

```
volume efficiency modify -vserver vs1 -volume VolA -compression false -inline-compression false
```

Learn more about `volume efficiency stop` in the [ONTAP command reference](#).

Manage inline data compaction for AFF systems

You can control inline data compaction on AFF systems at the volume level using the `volume efficiency modify` command. Data compaction is enabled by default for all volumes on AFF systems.

Before you begin

Data compaction requires that the volume space guarantee be set to none. This is the default for AFF systems.



The default space guarantee on non-AFF data protection volumes is set to none.

Steps

1. To verify the space guarantee setting for the volume:

```
volume show -vserver vserver_name -volume volume_name -fields space-guarantee
```

2. To enable data compaction:

```
volume efficiency modify -vserver vserver_name -volume volume_name -data-compaction true
```

3. To disable data compaction:

```
volume efficiency modify -vserver vserver_name -volume volume_name -data-compaction false
```

4. To display data compaction status:

```
volume efficiency show -instance
```

Examples

```
cluster1::> volume efficiency modify -vserver vs1 -volume vol1 -data-compaction true
cluster1::> volume efficiency modify -vserver vs1 -volume vol1 -data-compaction false
```

Enable inline data compaction for FAS systems

You can enable inline data compaction on FAS systems with Flash Pool (hybrid) aggregates or HDD aggregates at the volume level by using the `volume efficiency` cluster shell command. Data compaction is disabled by default for volumes created on

FAS systems. Learn more about volume efficiency in the [ONTAP command reference](#).

About this task

To enable inline data compaction on a volume, its `-space-guarantee` option must be set to `none`. Enabling data compaction on a volume on an HDD aggregate uses additional CPU resources.

Steps

1. Change to the advanced privilege level:

```
set -privilege advanced
```

Learn more about `set` in the [ONTAP command reference](#).

2. Check the data compaction state of the volumes and aggregates for the desired node:

```
volume efficiency show -volume <volume_name>
```

Learn more about `volume efficiency show` in the [ONTAP command reference](#).

3. Enable data compaction on volume:

```
volume efficiency modify -volume <volume_name> -data-compaction true
```

Learn more about `volume efficiency modify` in the [ONTAP command reference](#).



If data compaction is set to `false` for either an aggregate or a volume, then compaction fails. Enabling compaction does not compact existing data; only new writes to the system are compacted. The `volume efficiency start` command contains more information about how to compact existing data. Learn more about `volume efficiency start` in the [ONTAP command reference](#).

4. View the compaction statistics:

```
volume efficiency show -volume <volume_name>
```

Inline storage efficiency enabled by default on AFF systems

Storage efficiency features are enabled by default on all newly created volumes on AFF systems. All inline storage efficiency features are enabled by default on all existing and newly created volumes on all AFF systems.

Storage efficiency features include inline deduplication, inline cross-volume deduplication and inline compression, and are enabled by default on AFF systems as shown in the table.



Data compaction behavior on AFF volumes is enabled by default.

Volume conditions	Storage efficiency features enabled by default		
	Inline deduplication	Inline cross-volume deduplication	Inline compression
Cluster upgrade	Yes	Yes	Yes
ONTAP 7-Mode transition to clustered ONTAP	Yes	Yes	Yes
Volume move	Yes	Yes	Yes
Thick-provisioned volumes	Yes	No	Yes
Encrypted volumes	Yes	No	Yes

The following exceptions apply to one or more inline storage efficiency features:

- Only read-write volumes can support default inline storage efficiency enablement.
- Volumes with compression savings are omitted from enabling inline compression.
- Volumes that have postprocess deduplication turned on are omitted from enabling inline compression.
- On volumes where volume efficiency is turned off, the system overrides the existing volume efficiency policy settings and sets it to enable the inline-only policy.

Storage efficiency visualization

Use the `storage aggregate show-efficiency` command to display information about the storage efficiency of all the aggregates in your system.

The `storage aggregate show-efficiency` command has three different views that can be invoked by passing command options.

Learn more about `storage aggregate show-efficiency` in the [ONTAP command reference](#).

Default view

The default view displays the overall ratio for each of the aggregates.

```
cluster1::> storage aggregate show-efficiency
```

Detailed view

Invoke the detailed view with the `-details` command option. This view displays the following:

- Overall efficiency ratio for each of the aggregates.

- Overall ratio without snapshots.
- Ratio split for the following efficiency technologies: volume deduplication, volume compression, snapshots, clones, data compaction, and aggregate inline deduplication.

```
cluster1::> storage aggregate show-efficiency -details
```

Advanced view

The advanced view is similar to the detailed view and displays both logical and physical used details.

You must run this command at the advanced privilege level. Switch to advanced privilege by using the `set -privilege advanced` command.

The command prompt changes to `cluster1::*>`.

```
cluster1::> set -privilege advanced
```

Invoke the advanced view with the `-advanced` command option.

```
cluster1::*> storage aggregate show-efficiency -advanced
```

To view ratios for a single aggregate individually invoke the `-aggregate aggregate_name` command. This command can be run at the admin level, as well as the advanced privilege level.

```
cluster1::> storage aggregate show-efficiency -aggregate aggr1
```

Learn more about `set -privilege advanced` in the [ONTAP command reference](#).

Create a volume efficiency policy to run efficiency operations

Create a volume efficiency policy

You can create a volume efficiency policy to run deduplication or data compression followed by deduplication on a volume for a specific duration, and specify the job schedule using the `volume efficiency policy create` command.

Before you begin

You must have created a cron schedule using the `job schedule cron create` command. For more information about managing the cron schedules, see the [System administration reference](#). Learn more about `job schedule cron create` in the [ONTAP command reference](#).

About this task

An SVM administrator with default predefined roles cannot manage the deduplication policies. However, the cluster administrator can modify the privileges assigned to an SVM administrator by using any customized roles. For more information about the SVM administrator capabilities, see [Administrator authentication and RBAC](#).



You can run deduplication or data compression operations at a scheduled time, or by creating a schedule with a specific duration, or by specifying a threshold percentage, which waits for the new data to exceed the threshold and then triggers the deduplication or data compression operation. This threshold value is the percentage of the total number of blocks used in the volume. For example, if you set the threshold value on a volume to 20% when the total number of blocks used on the volume is 50%, data deduplication or data compression triggers automatically when new data written on the volume reaches 10% (20% of 50% blocks used). If required, you can obtain the total number of blocks used from the `df` command output.

Steps

1. Use the `volume efficiency policy create` command to create a volume efficiency policy.

Examples

The following command creates a volume efficiency policy named `pol1` that triggers an efficiency operation daily:

```
volume efficiency policy create -vserver vs1 -policy pol1 -schedule daily
```

The following command creates a volume efficiency policy named `pol2` that triggers an efficiency operation when the threshold percentage reaches 20%:

```
volume efficiency policy create -vserver vs1 -policy pol2 -type threshold -start -threshold-percent 20%
```

Learn more about `volume efficiency policy create` in the [ONTAP command reference](#).

Assign a volume efficiency policy to a volume

You can assign an efficiency policy to a volume to run deduplication or data compression operations by using the `volume efficiency modify` command.

Before you begin

Ensure that you [create the volume efficiency policy](#) before you assign it to a volume.

About this task

If an efficiency policy is assigned to a SnapVault secondary volume, only the volume efficiency priority attribute is considered when running volume efficiency operations. The job schedules are ignored and the deduplication operation is run when incremental updates are made to the SnapVault secondary volume.

Step

1. Use the `volume efficiency modify` command to assign a policy to a volume.

Example

The following command assigns the volume efficiency policy named `new_policy` to volume `VolA`:

```
volume efficiency modify -vserver vs1 -volume VolA -policy new_policy
```

Learn more about `volume efficiency modify` in the [ONTAP command reference](#).

Modify a volume efficiency policy

You can modify a volume efficiency policy to run deduplication and data compression for a different duration or change the job schedule using the `volume efficiency policy modify` command. Learn more about `volume efficiency policy modify` in the [ONTAP command reference](#).

Steps

1. Use the `volume efficiency policy modify` command to modify a volume efficiency policy.

Examples

The following command modifies the volume efficiency policy named `policy1` to run every hour:

```
volume efficiency policy modify -vserver vs1 -policy policy1 -schedule hourly
```

The following command modifies a volume efficiency policy named `pol2` to threshold 30%:

```
volume efficiency policy modify -vserver vs1 -policy pol1 -type threshold -start-threshold-percent 30%
```

View a volume efficiency policy in ONTAP

You can view the volume efficiency policy including the name, schedule, duration, and description.

About this task

The command `volume efficiency policy show` is used to display a volume efficiency policy. When you run the command in cluster scope, the cluster-scoped policies are not displayed. However, you can view the cluster-scoped policies in the SVM context. Learn more about `volume efficiency policy show` in the [ONTAP command reference](#).

Steps

1. Use the `volume efficiency policy show` command to view information about a volume efficiency policy.

The output depends on the parameters you specify.

Learn more about `volume efficiency policy show` in the [ONTAP command reference](#).

Examples

The following command displays information about the policies created for the SVM `vs1`:

```
volume efficiency policy show -vserver vs1
```

The following command displays the policies for which the duration is set as 10 hours:

```
volume efficiency policy show -duration 10
```

Disassociate a volume efficiency policy from a volume

You can disassociate a volume efficiency policy from a volume to stop running any further schedule-based deduplication and data compression operations on the volume. Once you disassociate a volume efficiency policy, you have to trigger it manually.

Step

1. Use the `volume efficiency modify` command to disassociate a volume efficiency policy from a volume.

Example

The following command disassociates the volume efficiency policy from volume VolA: `volume efficiency modify -vserver vs1 -volume VolA -policy -`

Learn more about `volume efficiency modify` in the [ONTAP command reference](#).

Delete a volume efficiency policy

You can delete a volume efficiency policy by using the `volume efficiency policy delete` command.

Before you begin

You must have ensured that the policy you want to delete is not associated with any volume.



You cannot delete the *inline-only* and the *default* predefined efficiency policy.

Step

1. Use the `volume efficiency policy delete` command to delete a volume efficiency policy.

Example

The following command deletes a volume efficiency policy named `policy1`: `volume efficiency policy delete -vserver vs1 -policy policy1`

Learn more about `volume efficiency policy delete` in the [ONTAP command reference](#).

Manage volume efficiency operations manually

Manage volume efficiency operations manually overview

You can manage how the efficiency operations run on a volume by running efficiency operations manually.

You can also control how the efficiency operations run based on the following conditions:

- Use checkpoints or not
- Run efficiency operations on existing data or only new data
- Stop efficiency operations if required

You can use the `volume efficiency show` command with `schedule` as value for the `-fields` option to view the schedule assigned to the volumes.

Learn more about `volume efficiency show` in the [ONTAP command reference](#).

Run an efficiency operation manually

You can run efficiency operations on a volume manually. You might do this when

scheduling efficiency operations is not appropriate.

Before you begin

Depending on the efficiency operation you want to run manually, you must have enabled deduplication or both data compression and deduplication on a volume.

About this task

This operation is performed using the `volume efficiency start` command. When temperature-sensitive storage efficiency is enabled on a volume, deduplication is run initially followed by data compression.

Deduplication is a background process that consumes system resources while it is running. If the data does not change often in a volume, it is best to run deduplication less frequently. Multiple concurrent deduplication operations running on a storage system lead to a higher consumption of system resources.

You can run a maximum of eight concurrent deduplication or data compression operations per node. If any more efficiency operations are scheduled, the operations are queued.

Beginning with ONTAP 9.13.1, if temperature-sensitive storage efficiency is enabled on a volume, you can run volume efficiency on existing data to take advantage of sequential packing to further improve storage efficiency.

Run efficiency manually

Steps

1. Start the efficiency operation on a volume: `volume efficiency start`

Example

+
The following command allows you to manually start only deduplication or deduplication followed by logical compression and container compression on the volume VolA

+

```
volume efficiency start -vserver vs1 -volume VolA
```

Repack existing data

To take advantage of sequential data packing introduced in ONTAP 9.13.1 on volumes with temperature-sensitive storage efficiency enabled, you can repack existing data. You must be in advanced privilege mode to use this command.

Steps

1. Set the privilege level: `set -privilege advanced`
2. Repack existing data: `volume efficiency inactive-data-compression start -vserver vserver_name -volume volume_name -scan-mode extended_recompression`

Example

```
volume efficiency inactive-data-compression start -vserver vs1 -volume
vol1 -scan-mode extended_recompression
```

Related information

- [Run efficiency operations manually on existing data](#)

Checkpoints and efficiency operations

Checkpoints are used internally to log the execution process of an efficiency operation. When an efficiency operation is stopped for any reason (such as system halt, system disruption, reboot, or because the last efficiency operation failed or stopped) and checkpoint data exists, the efficiency operation can resume from the latest checkpoint file.

A checkpoint is created:

- in each stage or substage of the operation
- when you run the `sis stop` command
- when the duration expires

Learn more about the commands described in this procedure in the [ONTAP command reference](#).

Resume a halted efficiency operation

If an efficiency operation is halted due to a system halt, system disruption, or reboot, you can resume the efficiency operation from the same point it was halted. This helps to save time and resources by not needing to restart the operation from the beginning.

About this task

If you enabled only deduplication on the volume, deduplication runs on the data. If you enabled both deduplication and data compression on a volume, then data compression runs first, followed by deduplication.

You can view the details of the checkpoint for a volume by using the `volume efficiency show` command. Learn more about `volume efficiency show` in the [ONTAP command reference](#).

By default, the efficiency operations resume from checkpoints. However, if a checkpoint corresponding to a previous efficiency operation (the phase when the `volume efficiency start -scan-old-data` command is run) is older than 24 hours, then the efficiency operation does not resume from the previous checkpoint automatically. In this case, the efficiency operation starts from the beginning. However, if you know that significant changes have not occurred in the volume since the last scan, you can force continuation from the previous checkpoint by using the `-use-checkpoint` option.

Steps

1. Use the `volume efficiency start` command with the `-use-checkpoint` option to resume an efficiency operation.

The following command enables you to resume an efficiency operation on new data on volume VolA:

```
volume efficiency start -vserver vs1 -volume VolA -use-checkpoint true
```

The following command enables you to resume an efficiency operation on existing data on volume VolA:

```
volume efficiency start -vserver vs1 -volume VolA -scan-old-data true -use  
-checkpoint true
```

Learn more about `volume efficiency start` in the [ONTAP command reference](#).

Run an efficiency operation manually on existing data

You can run the efficiency operations manually on the data that exists in non-temperature sensitive storage efficiency volumes prior to enabling deduplication, data compression, or data compaction. You can run these operations with ONTAP versions earlier than ONTAP 9.8.

About this task

This operation is performed using the `volume efficiency start` command with the `-scan-old-data` parameter. The `-compression` option does not work with `-scan-old-data` on temperature sensitive storage efficiency volumes. Inactive data compression runs automatically on pre-existing data for temperature sensitive storage efficiency volumes in ONTAP 9.8 and later.

If you enable only deduplication on a volume, then deduplication runs on the data. If you enable deduplication, data compression, and data compaction on a volume, then data compression runs first, followed by deduplication and data compaction.

When you run data compression on existing data, by default the data compression operation skips the data blocks that are shared by deduplication and the data blocks that are locked by snapshots. If you choose to run data compression on shared blocks, then optimization is turned off and the fingerprint information is captured and used for sharing again. You can change the default behavior of data compression when compressing existing data.

You can run a maximum of eight deduplication, data compression, or data compaction operations concurrently per node. The remaining operations are queued.



Post process compression does not run on AFF platforms. An EMS message is generated to inform you that this operation was skipped.

Learn more about `volume efficiency start` in the [ONTAP command reference](#).

Steps

1. Use the `volume efficiency start -scan-old-data` command to run deduplication, data compression, or data compaction manually on the existing data.

The following command enables you to run these operations manually on the existing data in volume VolA:

```
volume efficiency start -vserver vs1 -volume VolA -scan-old-data true [-  
compression | -dedupe | -compaction ] true
```

Related information

- [Run efficiency operations manually](#)

Manage volume efficiency operations using schedules

Run an efficiency operation based on the amount of new data written

You can modify the efficiency operation schedule to run deduplication or data compression when the number of new blocks written to the volume after the previous efficiency operation exceeds a specified threshold percentage. This applies whether the previous efficiency operation was performed manually or scheduled.

About this task

If the `schedule` option is set to `auto`, the scheduled efficiency operation runs when the amount of new data exceeds the specified percentage. The default threshold value is 20 percent. This threshold value is the percentage of the total number of blocks already processed by the efficiency operation.

Steps

1. Use the `volume efficiency modify` command with the `auto@num` option to modify the threshold percentage value.

`num` is a two-digit number to specify the percentage.

Example

The following command modifies the threshold percentage value to 30 percent for the volume `VolA`:

```
volume efficiency modify -vserver vs1 -volume -VolA -schedule auto@30
```

Related information

- [Run efficiency operations using scheduling](#)
- [volume efficiency modify](#)

Run an efficiency operation using scheduling

You can modify the scheduling of deduplication or data compression operations on a volume. The configuration options of a schedule and volume efficiency policy are mutually exclusive.

About this task

This operation is performed using the `volume efficiency modify` command. Learn more about `volume efficiency modify` in the [ONTAP command reference](#).

Steps

1. Use the `volume efficiency modify` command to modify the scheduling of deduplication or data compression operations on a volume.

Examples

The following command modifies the scheduling of efficiency operations for `VolA` to run at 11 p.m., Monday through Friday:

```
volume efficiency modify -vserver vs1 -volume VolA -schedule mon-fri@23
```

Related information

- Run efficiency operations depending on the amount of new data written

Monitor volume efficiency operations

View efficiency operations and status

You can view whether deduplication or data compression is enabled on a volume. You can also view the status, state, type of compression, and progress of the efficiency operations on a volume.

There are two tasks available. Both use the command `volume efficiency show`.

View efficiency status

Steps

1. View the status of an efficiency operation on a volume: `volume efficiency show`

The following command displays the status of an efficiency operation on volume VolA that is assigned the adaptive compression type:

```
volume efficiency show -instance -vserver vs1 -volume VolA
```

If the efficiency operation is enabled on volume VolA and the operation is idle, then you can see the following in the system output:

```
cluster1::> volume efficiency show -vserver vs1 -volume VolA

Vserver Name: vs1
Volume Name: VolA
Volume Path: /vol/VolA
State: Enabled
Status: Idle
Progress: Idle for 00:03:20
```

Determine if volumes contain sequentially packed data

You can display a list of volumes that have sequential packing enabled, for instance, when you need to revert to an ONTAP release earlier than 9.13.1. You must be in advanced privilege mode to use this command.

Steps

1. Set the privilege level: `set -privilege advanced`
2. List volumes that have sequential packing enabled:

```
volume efficiency show -extended-auto-adaptive-compression true
```

View efficiency space savings

You can view the amount of space savings achieved through deduplication and data

compression on a volume. You might do this to assess the effectiveness of your administrative processes or as part of capacity planning.

About this task

You need to use the command `volume show` to display the space savings on a volume. Note that the space savings in snapshots is not included when calculating the space savings achieved on a volume. Using deduplication does not affect volume quotas. Quotas are reported at the logical level and remain unchanged.

Steps

1. Use the `volume show` command to view space savings achieved on a volume using deduplication and data compression.

Example

The following command enables you to view the space savings achieved by using deduplication and data compression on volume VolA: `volume show -vserver vs1 -volume VolA`

```
cluster1::> volume show -vserver vs1 -volume VolA

          Vserver Name: vs1
          Volume Name: VolA

...
      Space Saved by Storage Efficiency: 115812B
      Percentage Saved by Storage Efficiency: 97%
      Space Saved by Deduplication: 13728B
      Percentage Saved by Deduplication: 81%
      Space Shared by Deduplication: 1028B
      Space Saved by Compression: 102084B
      Percentage Space Saved by Compression: 97%
...
```

Learn more about `volume show` in the [ONTAP command reference](#).

View efficiency statistics of a FlexVol volume

You can view the details of the efficiency operations run on a FlexVol volume. You might do this to assess the effectiveness of your administrative processes or as part of capacity planning.

Steps

1. Use the `volume efficiency stat` command to view the statistics of efficiency operations on a FlexVol volume.

Example

The following command enables you to view the statistics of the efficiency operations on the volume VolA: `volume efficiency stat -vserver vs1 -volume VolA`

```
cluster1::> volume efficiency stat -vserver vs1 -volume VolA

        Vserver Name: vs1
        Volume Name: VolA
        Volume Path: /vol/VolA
    Inline Compression Attempts: 0
```

Learn more about `volume efficiency stat` in the [ONTAP command reference](#).

Stop volume efficiency operations

You can stop a deduplication or postprocess compression operation.

About this task

This operation uses the command `volume efficiency stop`. This command automatically generates a checkpoint.

Steps

1. Use the `volume efficiency stop` command to stop an active deduplication or postprocess compression operation.

If you specify the `-all` option, active and queued efficiency operations are aborted.

Examples

The following command stops the deduplication or postprocess compression operation that is currently active on volume VolA:

```
volume efficiency stop -vserver vs1 -volume VolA
```

The following command aborts both active and queued deduplication or postprocess compression operations on volume VolA:

```
volume efficiency stop -vserver vs1 -volume VolA -all true
```

Learn more about `volume efficiency stop` in the [ONTAP command reference](#).

Additional information about removing space savings from a volume

You can choose to remove the space savings achieved by running efficiency operations on a volume. However, you must have enough space to accommodate a reversal.

There are several related resources available to help you plan and implement the removal of the space savings.

Related information

- [How to see space savings from deduplication, compression, and compaction in ONTAP 9](#)
- [How to undo the storage efficiency savings in ONTAP](#)

Rehost a volume from one SVM to another SVM

Prepare to rehost a volume from one SVM to another SVM

A volume rehost operation enables you to reassign a NAS or SAN volume from one SVM to another SVM without requiring a SnapMirror copy. The exact rehost procedure depends upon the client access protocol used and the volume type. Volume rehost is a disruptive operation for data access and volume management.

Before you can rehost a volume from one SVM to another, the following conditions must be met:

- The volume must be online
- The volume protocol must be SAN or NAS
 - For the NAS protocol volumes, the volume should not be a part of junction-path and must be unmounted
- If the volume is in a SnapMirror relationship, then the relationship must be deleted, followed by releasing the relationship information only, or broken prior to volume rehost
 - You can resynchronize the SnapMirror relationship after the volume rehost operation
- The vserver subtype should be same for both source and destination SVMs
 - Volumes can only be rehosted between SVMs of the same subtype
- The volume cannot be FlexClone or FlexClone Parent
 - FlexClones must be split before rehosting the parent or clone volume

Rehost an SMB volume

You can rehost a volume that serves data using the SMB protocol. To allow clients to continue accessing the data after the rehosting operation, you must manually configure policies and the associated rules.

About this task

- Rehosting is a disruptive operation.
- If the rehosting operation fails, you might need to reconfigure the volume policies and the associated rules on the source volume.
- If the source SVM and destination SVM Active Directory domains differ, you might lose access to the objects on the volume.
- Beginning with ONTAP 9.8, rehosting a volume with NetApp Volume Encryption (NVE) is supported. If you are using an onboard key manager, the encrypted metadata will be modified during the rehost operation. User data is not changed.

If you are using ONTAP 9.8 or earlier, you must unencrypt the volume before performing the rehost operation.

- When the source SVM has local users and groups, the permissions for the files and directories (ACLs) that are set are no longer effective after volume rehost operation.

The same is true for audit ACLs (SACLs)

- After the rehost operation, the following volume policies, policy rules, and configurations are lost from the source volume, and must be manually reconfigured on the rehosted volume:
 - Volume and qtree export policies
 - Antivirus policies
 - Volume efficiency policy
 - Quality of service (QoS) policies
 - Snapshot policies
 - Quota rules
 - ns-switch and name services configuration export policy and rules
 - User and group IDs

Before you begin

- Volume must be online.
- Volume management operations, such as volume move or LUN move, must not be running.
- Data access to the volume that is being rehosted must be stopped.
- The ns-switch and name services configuration of the target SVM must be configured to support data access of the rehosting volume.
- The source SVM and destination SVM must have the same Active Directory and realmDNS domain.
- The user ID and group ID of the volume must be either available in the target SVM or changed on the hosting volume.



If local users and groups are configured, and if there are files and directories on that volume with permissions set for those users or groups, these permissions are no longer effective.

Steps

1. Record information about the CIFS shares to avoid losing information on CIFS shares in case volume rehost operation fails.
2. Unmount the volume from the parent volume:

```
volume unmount
```

3. Switch to the advanced privilege level:

```
set -privilege advanced
```

4. Rehost the volume on the destination SVM:

```
volume rehost -vserver source_svm -volume vol_name -destination-vserver
destination_svm
```

5. Mount the volume under the appropriate junction path in the destination SVM:

```
volume mount
```

6. Create CIFS shares for the rehosted volume:

```
vserver cifs share create
```

7. If the DNS domains differ between the source SVM and destination SVM, create new users and groups.
8. Update the CIFS client with the new destination SVM LIFs and junction path to the rehosted volume.

After you finish

You must manually reconfigure the policies and the associated rules on the rehosted volume.

SMB configuration

SMB and NFS multiprotocol configuration

Rehost an NFS volume

You can rehost a volume that serves data using the NFS protocol. To allow clients to continue accessing the data after the rehosting operation, you must associate the volume with the export policy of the SVM as well as manually configure the policies and associated rules.

About this task

- Rehosting is a disruptive operation.
- If the rehosting operation fails, you might need to reconfigure the volume policies and the associated rules on the source volume.
- Beginning with ONTAP 9.8, rehosting a volume with NetApp Volume Encryption (NVE) is supported. If you are using an onboard key manager, the encrypted metadata will be modified during the rehost operation. User data is not changed.

If you are using ONTAP 9.8 or earlier, you must unencrypt the volume before performing the rehost operation.

- After the rehost operation, the following volume policies, policy rules, and configurations are lost from the source volume, and must be manually reconfigured on the rehosted volume:
 - Volume and qtree export policies
 - Antivirus policies
 - Volume efficiency policy
 - Quality of service (QoS) policies
 - Snapshot policies
 - Quota rules
 - ns-switch and name services configuration export policy and rules
 - User and group IDs

Before you begin

- The volume must be online.
- Volume management operations, such as volume moves or LUN moves, must not be running.
- Data access to the volume that is being rehosted must be stopped.
- The ns-switch and name services configuration of the target SVM must be configured to support data

access of the rehosting volume.

- The user ID and group ID of the volume must be either available in the target SVM or changed on the hosting volume.

Steps

1. Record information about the NFS export policies to avoid losing information on NFS policies in case volume rehost operation fails.

2. Unmount the volume from the parent volume:

```
volume unmount
```

3. Switch to the advanced privilege level:

```
set -privilege advanced
```

4. Rehost the volume on the destination SVM:

```
volume rehost -vserver source_svm -volume volume_name -destination-vserver  
destination_svm
```

The default export policy of the destination SVM is applied to the rehosted volume.

5. Create the export policy:

```
vserver export-policy create
```

6. Update the export policy of the rehosted volume to a user-defined export policy:

```
volume modify
```

7. Mount the volume under the appropriate junction path in the destination SVM:

```
volume mount
```

8. Verify that the NFS service is running on the destination SVM.

9. Resume NFS access to the rehosted volume.

10. Update the NFS client credentials and LIF configurations to reflect the destination SVM LIFs.

This is because the volume access path (LIFs and junction path) has undergone changes.

After you finish

You must manually reconfigure the policies and the associated rules on the rehosted volume. See [NFS configuration](#) for more information.

Rehost a SAN volume

You can rehost a SAN volume that serves data through mapped LUNs. After re-creating the initiator group (igroup) in the destination SVM, volume rehost operation can automatically remap the volume at the same SVM.

About this task

- Rehosting is a disruptive operation.
- If the rehosting operation fails, you might need to reconfigure the volume policies and the associated rules on the source volume.
- Beginning with ONTAP 9.8, rehosting a volume with NetApp Volume Encryption (NVE) is supported. If you are using an onboard key manager, the encrypted metadata will be modified during the rehost operation. User data is not changed.

If you are using ONTAP 9.8 or early, you must unencrypt the volume before performing the rehost operation.

- After the rehost operation, the following volume policies, policy rules, and configurations are lost from the source volume and must be manually reconfigured on the rehosted volume:
 - Antivirus policies
 - Volume efficiency policy
 - Quality of service (QoS) policies
 - Snapshot policies
 - ns-switch and name services configuration export policy and rules
 - User and group IDs

Before you begin

- The volume must be online.
- Volume management operations, such as volume moves or LUN moves, must not be running.
- There must be no active I/O on the volumes or LUNs.
- You must have verified that the destination SVM does not have igroup of the same name but different initiators.

If the igroup has the same name, then you must have renamed the igroup in either one of the SVMs (source or destination).

- You must have enabled the `force-unmap-luns` option.
 - The default value of the `force-unmap-luns` option is `false`.
 - No warning or confirmation message is displayed when you set the `force-unmap-luns` option to `true`.

Steps

1. Record LUN mapping information on target volume:

```
lun mapping show volume volume vserver source_svm
```

This is a precautionary step to avoid losing information about LUN mapping in case the volume rehost fails.

Learn more about `lun mapping show volume` in the [ONTAP command reference](#).

2. Delete igroups associated with the target volume.

3. Rehost the target volume to the destination SVM:

```
volume rehost -vserver source_svm -volume volume_name -destination-vserver
```

destination_svm

4. Map LUNs on the target volume to appropriate igroups:

- Volume rehost preserves LUNs on the target volume, however the LUNs remain unmapped.
- Use the destination SVM port set while mapping LUNs.
- If the `auto-remap-luns` option is set to `true`, the LUNs are mapped automatically after rehost.

Rehost a volume in a SnapMirror relationship

You can rehost a volume defined as part of a SnapMirror relationship. There are several issues you need to consider before rehosting the relationship.

About this task

- Rehosting is a disruptive operation.
- If the rehosting operation fails, you might need to reconfigure the volume policies and the associated rules on the source volume.
- After the rehost operation, the following volume policies, policy rules, and configurations are lost from the source volume and must be manually reconfigured on the rehosted volume:
 - Volume and qtree export policies
 - Antivirus policies
 - Volume efficiency policy
 - Quality of service (QoS) policies
 - Snapshot policies
 - Quota rules
 - ns-switch and name services configuration export policy and rules
 - User and group IDs

Before you begin

- The volume must be online.
- Volume management operations, such as volume moves or LUN moves, must not be running.
- Data access to the volume that is being rehosted must be stopped.
- The ns-switch and name services configuration of the target SVM must be configured to support data access of the rehosting volume.
- The user ID and group ID of the volume must be either available in the target SVM or changed on the hosting volume.

Steps

1. Record the SnapMirror relationship type:

```
snapmirror show
```

This is a precautionary step to avoid losing information about the SnapMirror relationship type in case the volume rehost fails.

2. From the destination cluster, delete the SnapMirror relationship:

```
snapmirror delete
```

Do not break the SnapMirror relationship; otherwise, the data protection capability of the destination volume is lost and the relationship cannot be reestablished after the rehosting operation.

3. From the source cluster, remove the SnapMirror relationship information:

```
snapmirror release -relationship-info-only true
```

Setting the `-relationship-info-only` parameter to `true` removes the source relationship information without deleting the snapshots.

4. If the volume is mounted, unmount it:

```
volume unmount -vserver <source_svm> -volume <vol_name>
```

5. Switch to the advanced privilege level:

```
set -privilege advanced
```

6. Rehost the volume on the destination SVM:

```
volume rehost -vserver <source_svm> -volume <vol_name> -destination-vserver <destination_svm>
```

7. If the SVM peering relation is not present, create the SVM peer relationship between the source SVM and destination SVM:

```
vserver peer create
```

8. Create the SnapMirror relationship between the source volume and destination volume:

```
snapmirror create
```

You must run the `snapmirror create` command from the SVM that is hosting the DP volume. The rehosted volume can be the source or destination of the SnapMirror relationship.

9. Resynchronize the SnapMirror relationship.

Related information

- [set](#)
- [snapmirror](#)
- [volume rehost](#)
- [volume unmount](#)
- [vserver peer create](#)

Features not supported with a volume rehost in ONTAP

There are several ONTAP features that do not support volume rehost. You should be aware of these features before attempting a rehost operation.

The following features are not supported with a volume rehost:

- SVM DR
- MetroCluster configurations



Cloning a volume as a FlexClone volume on a different SVM is also not supported on MetroCluster configurations.

- SnapLock volumes
- NetApp Volume Encryption (NVE) volumes (in versions of ONTAP before 9.8)

In ONTAP releases prior to 9.8, you must unencrypt the volume before rehosting it. Volume encryption keys depend on SVM keys. If a volume is moved to another SVM and if multitenant key configuration is enabled on either the source or destination SVM, the volume and the SVM keys will not match.

Beginning with ONTAP 9.8, you can rehost a volume with NVE.

- FlexGroup volumes
- Clone volumes

Recommended volume and file or LUN configuration combinations

Overview of recommended volume and file or LUN configuration combinations

There are specific combinations of FlexVol volume and file or LUN configurations you can use, depending on your application and administration requirements. Understanding the benefits and costs of these combinations can help you determine the right configuration for your environment.

The following volume and LUN configuration combinations are recommended:

- Space-reserved files or LUNs with thick volume provisioning
- Non-space-reserved files or LUNs with thin volume provisioning
- Space-reserved files or LUNs with semi-thick volume provisioning

You can use SCSI thin provisioning on your LUNs in conjunction with any of these configuration combinations.

Space-reserved files or LUNs with thick volume provisioning

Benefits:

- All write operations within space-reserved files are guaranteed; they will not fail due to insufficient space.
- There are no restrictions on storage efficiency and data protection technologies on the volume.

Costs and limitations:

- Enough space must be set aside from the aggregate up front to support the thickly provisioned volume.
- Space equal to twice the size of the LUN is allocated from the volume at LUN creation time.

Non-space-reserved files or LUNs with thin volume provisioning

Benefits:

- There are no restrictions on storage efficiency and data protection technologies on the volume.
- Space is allocated only as it is used.

Costs and restrictions:

- Write operations are not guaranteed; they can fail if the volume runs out of free space.
- You must manage the free space in the aggregate effectively to prevent the aggregate from running out of free space.

Space-reserved files or LUNs with semi-thick volume provisioning

Benefits:

Less space is reserved up front than for thick volume provisioning, and a best-effort write guarantee is still provided.

Costs and restrictions:

- Write operations can fail with this option.

You can mitigate this risk by properly balancing free space in the volume against data volatility.

- You cannot rely on retention of data protection objects such as snapshots and FlexClone files and LUNs.
- You cannot use ONTAP block-sharing storage efficiency capabilities that cannot be automatically deleted, including deduplication, compression, and ODX/Copy Offload.

Determine the correct volume and LUN configuration for your needs

Answering a few basic questions about your environment can help you determine the best FlexVol volume and LUN configuration for your environment.

About this task


You can optimize your LUN and volume configurations for maximum storage utilization or for the security of write guarantees. Based on your requirements for storage utilization and your ability to monitor and replenish free space quickly, you must determine the FlexVol volume and LUN volumes appropriate for your installation.



You do not need a separate volume for each LUN.

Step

1. Use the following decision tree to determine the best volume and LUN configuration combination for your environment:



Configuration settings for space-reserved files or LUNs with thick-provisioned volumes

There are several configuration combinations of FlexVol volume and file or LUN configurations you can use.

This combination based on thick-provisioned volumes provides the ability to use storage efficiency technologies and does not require you to actively monitor your free space because sufficient space is allocated up front.

The following settings are required to configure a space-reserved file or LUN in a volume using thick provisioning:

Volume setting	Value
Guarantee	Volume
Fractional reserve	100
Snapshot reserve	Any
Snapshot autodelete	Optional
Autogrow	Optional; if enabled, aggregate free space must be actively monitored.

File or LUN setting	Value
Space reservation	Enabled

Related information

- [Recommended volume and file or LUN configuration combinations overview](#)

Settings for non-space-reserved files or LUNs with thin-provisioned volumes

This FlexVol volume and file or LUN configuration combination requires the smallest amount of storage to be allocated up front, but requires active free space management to prevent errors due to lack of space.

The following settings are required to configure a non-space-reserved files or LUN in a thin-provisioned volume:

Volume setting	Value
Guarantee	None
Fractional reserve	0
Snapshot reserve	Any
Snapshot autodelete	Optional
Autogrow	Optional

File or LUN setting	Value
Space reservation	Disabled

Additional considerations

When the volume or aggregate runs out of space, write operations to the file or LUN can fail.

If you do not want to actively monitor free space for both the volume and the aggregate, you should enable Autogrow for the volume and set the maximum size for the volume to the size of the aggregate. In this configuration, you must monitor aggregate free space actively, but you do not need to monitor the free space in the volume.

Configuration settings for space-reserved files or LUNs with semi-thick volume provisioning

There are several configuration combinations of FlexVol volume and file or LUN configurations you can use. This combination based on semi-thick volume provisioning requires less storage to be allocated up front than the fully provisioned combination. But it places restrictions on the efficiency technologies you can use for the volume. Overwrites are fulfilled on a best-effort basis for this configuration combination.

The following settings are required to configure a space-reserved LUN in a volume using semi-thick provisioning:

Volume setting	Value
Guarantee	Volume
Fractional reserve	0
Snapshot reserve	0
Snapshot autodelete	On, with a commitment level of destroy, a destroy list that includes all objects, the trigger set to volume, and all FlexClone LUNs and FlexClone files enabled for automatic deletion.
Autogrow	Optional; if enabled, aggregate free space must be actively monitored.

File or LUN setting	Value
Space reservation	Enabled

Technology restrictions

You cannot use the following volume storage efficiency technologies for this configuration combination:

- Compression
- Deduplication
- ODX and FlexClone Copy Offload
- FlexClone LUNs and FlexClone files not marked for automatic deletion (active clones)
- FlexClone subfiles
- ODX/Copy Offload

Additional considerations

The following facts must be considered when employing this configuration combination:

- When the volume that supports that LUN runs low on space, protection data (FlexClone LUNs and files, snapshots) is destroyed.
- Write operations can time out and fail when the volume runs out of free space.

Compression is enabled by default for AFF platforms. You must explicitly disable compression for any volume for which you want to use semi-thick provisioning on an AFF platform.

Related information

- [Recommended volume and file or LUN configuration combinations overview](#)

Cautions and considerations for changing file or directory capacity

The default and maximum number of files allowed for FlexVol volumes in ONTAP

FlexVol volumes have a default and maximum number of files they can contain. If your data requires a large number of files, you can increase the number of user visible files allowed on a volume up to a maximum value. You should understand the limitations and caveats before proceeding.

The number of user visible files a volume can contain is determined by the available inode capacity for the volume. An inode is a data structure that contains information about files.

ONTAP automatically sets the default and maximum number of available inodes for a newly created volume as follows based on the size of the volume.

Default number of inodes	Maximum number of inodes
1 per 32 KB of volume size	1 per 4 KB of volume size

When the size of a volume is increased, either manually by an administrator or automatically by ONTAP's autosize feature, ONTAP also increases (if necessary) the number of available inodes so that there is at least 1 inode per 32 KB of volume size, until the volume reaches approximately 680 GB in size.

In ONTAP 9.12.1 and earlier, creating a new volume or resizing an existing volume greater than 680 GB in size does not automatically result in additional inode capacity. If you need more files than the default number for any size volume, you can use the `volume modify` command to increase the available number of inodes for the volume up to the maximum.

Beginning with ONTAP 9.13.1, creating a new volume or resizing an existing volume sets the default number of available inodes to 1 inode per 32 KB of volume space even if the volume is larger than 680 GB. This ratio persists until the volume reaches the absolute inode maximum of 2,040,109,451.

You can also decrease the available number of inodes. This does not change the amount of space allocated to inodes, but it does lower the maximum amount of space the public inode file can consume. After space has been allocated for inodes, it is never returned to the volume. Therefore, it is not possible to lower the maximum number of inodes below the number of inodes currently allocated.

More information

- [Determine file and inode usage for a volume](#)
- [NetApp Knowledge Base: FAQ - ONTAP default and maximum number of files \(inodes\)](#)

Maximum directory size for FlexVol volumes

You can increase the default maximum directory size for a specific FlexVol volume by using the `-maxdir-size` option of the `volume modify` command, but doing so could impact system performance. See the [NetApp Knowledge Base: What is maxdirsize?](#).

To learn more about the model-dependent maximum directory sizes for FlexVol volumes, visit the [NetApp Hardware Universe](#).

Learn more about `volume modify` in the [ONTAP command reference](#).

Restrictions on node root volumes and root aggregates

You should be aware of the restrictions governing a node's root volume and root aggregate.



A node's root volume contains special directories and files for the node. The root volume is contained in the root aggregate.

A node's root volume is a FlexVol volume that is installed at the factory or by setup software. It is reserved for system files, log files, and core files. The directory name is `/mroot`, which is accessible only through the systemshell by technical support. The minimum size for a node's root volume depends on the platform model.

- The following rules govern the node's root volume:
 - Unless technical support instructs you to do so, do not modify the configuration or content of the root volume.
 - Do not store user data in the root volume.

Storing user data in the root volume increases the storage giveback time between nodes in an HA pair.

- You can move the root volume to another aggregate.

[Relocating root volumes to new aggregates](#)

- The root aggregate is dedicated to the node's root volume only.

ONTAP prevents you from creating other volumes in the root aggregate.

[NetApp Hardware Universe](#)

Relocate a root volume to new aggregates

The root replacement procedure migrates the current root aggregate to another set of disks without disruption. You might need to perform this as part of a disk replacement or preventative maintenance process.

About this task

You can change the location of the root volume to a new aggregate in the following scenarios:

- When the root aggregates are not on the disk you prefer
- When you want to rearrange the disks connected to the node
- When you are performing a shelf replacement of the EOS disk shelves

Steps

1. Relocate the root aggregate:

```
system node migrate-root -node node_name -disklist disk_list -raid-type
raid_type
```

- **-node**

Specifies the node that owns the root aggregate that you want to migrate.

- **-disklist**

Specifies the list of disks on which the new root aggregate will be created. All disks must be spares and owned by the same node. The minimum number of disks required is dependent on the RAID type.

- **-raid-type**

Specifies the RAID type of the root aggregate. The default value is `raid-dp`. This is the only type supported in advanced mode.

2. Monitor the progress of the job:

```
job show -id jobid -instance
```

Results

If all of the pre-checks are successful, the command starts a root volume replacement job and exits.

Features supported by FlexClone files and FlexClone LUNs

Features supported by FlexClone files and FlexClone LUNs

FlexClone files and FlexClone LUNs work with different ONTAP features, such as deduplication, snapshots, quotas, and volume SnapMirror.

The following features are supported by FlexClone files and FlexClone LUNs:

- Deduplication
- Snapshots
- Access control lists
- Quotas
- FlexClone volumes
- NDMP
- Volume SnapMirror
- The `volume move` command
- Space reservation
- HA configuration

Deduplication with FlexClone files and FlexClone LUNs

You can efficiently use the physical storage space of the data blocks by creating a FlexClone file or FlexClone LUN of the parent file and parent LUN in a deduplication-enabled volume.

The block-sharing mechanism used by FlexClone files and LUNs is also used by deduplication. You can maximize the space savings in a FlexVol volume by enabling deduplication on the volume and then cloning the deduplication-enabled volume.



While executing the `sis undo` command on a deduplication-enabled volume, you cannot create FlexClone files and FlexClone LUNs of the parent files and parent LUNs residing in that volume.

Learn more about the commands described in this procedure in the [ONTAP command reference](#).

How snapshots work with FlexClone files and FlexClone LUNs

There is a synergy between snapshots and the FlexClone files and FlexClone LUNs. If you work with these technologies, you should be aware of what is possible as well as the relevant restrictions.

Creating FlexClone files and LUNs

You can create a FlexClone file or FlexClone LUN from an existing snapshot. The copy is based on the parent files and parent LUNs contained in a FlexVol volume.

Deleting a snapshot

You cannot manually delete a snapshot from which FlexClone files or FlexClone LUNs are currently being created. The snapshot remains locked until the background block-sharing process is completed. If you try to delete a locked snapshot, the system displays a message asking you to retry the operation after some amount of time. In this case, you need to continue retrying the deletion operation. You'll be able to delete the snapshot after the block sharing is done.

Inheritance of access control lists by FlexClone files and FlexClone LUNs

FlexClone files and FlexClone LUNs inherit the access control lists of their parent files and LUNs.

If the parent files contain Windows NT streams, the FlexClone files also inherit the stream information. However, parent files containing more than six streams cannot be cloned.

How quotas work with FlexClone files and FlexClone LUNs

You should be familiar with how quotas work with FlexClone files and FlexClone LUNs before using them.

Quota limits are applied on the total logical size of the FlexClone files or FlexClone LUNs. Cloning operations do not fail block sharing even if it causes quotas to be exceeded.

When you create a FlexClone file or FlexClone LUN, quotas do not recognize any space savings. For example, if you create a FlexClone file of a parent file of 10 GB, you are only using 10 GB of physical space, but the quota utilization is recorded as 20 GB (10 GB for the parent and 10 GB for the FlexClone file).

If the creation of a FlexClone file or LUN results in the group or user quota's being exceeded, the clone operation succeeds provided the FlexVol volume has enough space to hold the metadata for the clone. However, the quota for that user or group is oversubscribed.

FlexClone volumes and associated FlexClone files and FlexClone LUNs

You can create a FlexClone volume of a FlexVol volume that has both a FlexClone file and FlexClone LUN and its parent file or LUN in it.

FlexClone files or FlexClone LUNs and their parent files or LUNs that are present in the FlexClone volume continue to share blocks the same way they do in the parent FlexVol volume. In fact, all the FlexClone entities and their parents share the same underlying physical data blocks, minimizing physical disk space usage.

If the FlexClone volume is split from its parent volume, then the FlexClone files or FlexClone LUNs and their parent files or LUNs stop sharing the blocks in the clone of the FlexClone volume. Thereafter they exist as independent files or LUNs. This means that the clone of the volume uses more space than before the splitting operation.

How NDMP works with FlexClone files and FlexClone LUNs

NDMP works at the logical level with FlexClone files and FlexClone LUNs. All FlexClone files or LUNs are backed up as separate files or LUNs.

When you use NDMP services to back up a qtree or a FlexVol volume that contains FlexClone files or FlexClone LUNs, block sharing between parent and clone entities is not preserved, and clone entities are backed up to tape as separate files or LUNs. The space saving is lost. Therefore, the tape onto which you are backing up should have sufficient space to store the expanded amount of data. When you restore, all the FlexClone files and FlexClone LUNs are restored as separate physical files and LUNs. You can enable deduplication on the volume to restore the block-sharing benefits.

 When FlexClone files and FlexClone LUNs are being created from an existing snapshot of a FlexVol volume, you cannot back up the volume to tape until the block-sharing process, which happens in the background, is complete. If you use NDMP on the volume when the block-sharing process is in progress, the system displays a message asking you to retry the operation after some time. In such a situation, you must keep retrying the tape backup operation so that it succeeds after the block sharing is complete.

How volume SnapMirror works with FlexClone files and FlexClone LUNs

Using volume SnapMirror with FlexClone files and FlexClone LUNs helps in maintaining space savings because the cloned entities are replicated only once.

If a FlexVol volume is a volume SnapMirror source and contains FlexClone files or FlexClone LUNs, volume SnapMirror transfers only the shared physical block and a small amount of metadata to the volume SnapMirror destination. The destination stores only one copy of the physical block, and this block is shared between the parent and cloned entities. Therefore, the destination volume is an exact copy of the source volume and all the clone files or LUNs on the destination volume share the same physical block.

How space reservation works with FlexClone files and FlexClone LUNs

When using FlexClone files and FlexClone LUNs, you should understand how the space reservation attribute works.

By default, the FlexClone files and LUNs inherit the space reservation attribute from the parent file and parent LUN respectively. However, you can create FlexClone files and FlexClone LUNs with space reservation disabled if the FlexVol volume lacks space. This is possible even if the attribute in the respective parent is enabled.

Note that if the FlexVol volume does not contain enough space to create a FlexClone file or FlexClone LUN with the same space reservation as that of the parent, the cloning operation will fail.

How an HA configuration works with FlexClone files and FlexClone LUNs

FlexClone file and FlexClone LUN operations are supported in an HA configuration.

In an HA pair, you cannot create FlexClone files or FlexClone LUNs on the partner while the takeover or giveback operation is in progress. All the pending block sharing operations on the partner are resumed after the takeover or giveback operation is complete.

FlexGroup volumes management

Learn about ONTAP FlexGroup volumes management with the CLI

You can set up, manage, and protect FlexGroup volumes for scalability and performance. A FlexGroup volume is a scale-out volume that provides high performance along with automatic load distribution.

You can configure FlexGroup volumes if the following are true:

- You want to use best practices, not explore every available option.
- You have cluster administrator privileges, not SVM administrator privileges.



Beginning with ONTAP 9.5, FlexGroup volumes replace Infinite Volumes, which are not supported in ONTAP 9.5 or later releases.


See the [supported and unsupported configurations for FlexGroup volumes](#) for more information.

Related information

Conceptual information about FlexVol volumes is applicable to FlexGroup volumes. Information about FlexVol volumes and ONTAP technology is available in the ONTAP Reference Library and in Technical Reports (TRs).

Learn about ONTAP FlexGroup volumes

A FlexGroup volume is a scale-out NAS container that provides high performance along with automatic load distribution and scalability. A FlexGroup volume contains several member volumes (constituents) that automatically and transparently share the traffic. *Member volumes* are the underlying FlexVol volumes that make up a FlexGroup volume.



FlexGroup volumes provide the following benefits:

- High scalability

Multiple FlexGroup volumes can be provisioned on a cluster as long as the number of member volumes does not exceed the node or cluster limits.

Beginning with ONTAP 9.12.1P2, the maximum capacity for a single FlexGroup volume is 60PB, with 400 billion files on a 10-node cluster when [large volume support is enabled](#). Without large volume support, the maximum capacity for a single FlexGroup volume is 20PB.



Although the maximum capacity of a single FlexGroup volume is 60PB (200 member volumes x 300TB = 60PB), best performance is achieved when the used capacity of member volumes remains below 80% (200 member volumes x 240TB = 48PB).

- High performance

FlexGroup volumes can use the resources of the cluster to serve workloads that have high throughput and low latency.

- Simplified management

A FlexGroup volume is a single namespace container that can be managed in a similar way as FlexVol volumes.

Supported and unsupported configurations for ONTAP FlexGroup volumes

You should be aware of the ONTAP features that are supported and not supported with FlexGroup volumes in ONTAP 9.

Features supported beginning with ONTAP 9.18.1

- [Nested QoS policies](#) are supported for the following object pairs:
 - SVMs and FlexGroup volumes contained by the SVM
 - FlexGroup volumes and qtrees within the volumes

Features supported beginning with ONTAP 9.16.1

- [Advanced capacity balancing](#)

Features supported beginning with ONTAP 9.15.1

- [Automatic provisioning enhancements](#)

Features supported beginning with ONTAP 9.14.1

- Snapshot tagging: Support for creating, modifying and deleting snapshot tags (SnapMirror labels and comments) for snapshots on FlexGroup volumes using the `volume snapshot` command.

Features supported beginning with ONTAP 9.13.1

- [Autonomous Ransomware Protection \(ARP\)](#) for FlexGroup volumes, including the following supported functionality:
 - FlexGroup expand operations: A new member volume inherits Autonomous Ransomware Protection attributes.
 - FlexVol to FlexGroup conversions: Conversions of FlexVols with active Autonomous Ransomware Protection is possible.
 - FlexGroup rebalancing: Autonomous Ransomware Protection is supported during disruptive and non-disruptive rebalancing operations.
- Schedule a single FlexGroup rebalancing operation.
- [SnapMirror fanout](#) relationships with SVM DR on FlexGroup volumes. Supports fanout to eight sites.

Features supported beginning with ONTAP 9.12.1

- [FlexGroup rebalancing](#)
- SnapLock for SnapVault
- [SnapMirror Cloud](#)
- FabricPool, FlexGroup, and SVM DR working in conjunction. (In releases earlier than ONTAP 9.12.1, any two of these features worked together, but not all three in conjunction.)
- [Large volume support](#) increases FlexGroup volume member size from a maximum of 100TB to a maximum of 300TB.

Features supported beginning with ONTAP 9.11.1

- [SnapLock volumes](#)

SnapLock does not support the following features with FlexGroup volumes:

- Legal-hold
- Event-based retention
- SnapLock for SnapVault

You configure SnapLock at the FlexGroup level. You cannot configure SnapLock at the member volume level.

- [Client asynchronous directory delete](#)

Features supported beginning with ONTAP 9.10.1

- Convert a FlexVol volume to a FlexGroup volume within an SVM DR relationship
- SVM DR FlexClone support for FlexGroup volumes

Features supported beginning with ONTAP 9.9.1

- SVM disaster recovery

Cloning a FlexGroup volume that is part of an SVM DR relationship is not supported.

- SnapMirror fanout relationships of 2 or more (A to B, A to C), with a maximum of 8 fanout legs.

[Considerations for creating SnapMirror cascade and fanout relationships for FlexGroup volumes](#)

- SnapMirror cascading relationships up to two levels (A to B to C)

[Considerations for creating SnapMirror cascade and fanout relationships for FlexGroup volumes](#)

Features supported beginning with ONTAP 9.8

- Restoring a single file from a FlexGroup SnapMirror vault or from a UDP destination
 - Restore can be from a FlexGroup volume of any geometry to FlexGroup volume of any geometry
 - Only one file per restore operation is supported
- Converting volumes transitioned from 7-Mode systems to FlexGroup volumes

For more information, see the [NetApp Knowledge Base: How To Convert a Transitioned FlexVol to FlexGroup](#).

- NFSv4.2
- [Asynchronous delete of files and directories](#)
- [Files System Analytics \(FSA\)](#)
- FlexGroup as a VMware vSphere datastore
- Additional support for tape backup and restore using NDMP, including the following features:
 - NDMP restartable backup extension (RBE) and Snapshot Management Extension (SSME)
 - Environment variables EXCLUDE and MULTI_SUBTREE_NAMES support FlexGroup backups
 - Introduction of IGNORE_CTIME_MTIME environment variable for FlexGroup backups
 - Individual file recovery in a FlexGroup using the NDMP_SNAP_RECOVER message, which is part of extension 0x2050

Dump and restore sessions are aborted during an upgrade or revert.

Features supported beginning with ONTAP 9.7

- [FlexClone volume](#)
- NFSv4 and NFSv4.1
- pNFS
- [Tape backup and restore by using NDMP](#)

You must be aware of the following points for NDMP support on FlexGroup volumes:

- The NDMP_SNAP_RECOVER message in the extension class 0x2050 can be used only for recovering an entire FlexGroup volume.

Individual files in a FlexGroup volume cannot be recovered.

- NDMP restartable backup extension (RBE) is not supported for FlexGroup volumes.
- Environment variables EXCLUDE and MULTI_SUBTREE_NAMES are not supported for FlexGroup volumes.
- The `ndmpcopy` command is supported for data transfer between FlexVol and FlexGroup volumes.

If you revert from Data ONTAP 9.7 to an earlier version, the incremental transfer information of the previous transfers is not retained and therefore, you must perform a baseline copy after reverting.

- VMware vStorage APIs for Array Integration (VAAI)
- Conversion of a FlexVol volume to a FlexGroup volume
- FlexGroup volumes as FlexCache origin volumes

Features supported beginning with ONTAP 9.6

- Continuously available SMB shares
- [MetroCluster configurations](#)
- Renaming a FlexGroup volume (`volume rename` command)
- Shrinking or reducing the size of a FlexGroup volume (`volume size` command)
- Elastic sizing
- NetApp aggregate encryption (NAE)
- Cloud Volumes ONTAP

Features supported beginning with ONTAP 9.5

- ODX copy offload
- Storage-Level Access Guard
- Enhancements to change notifications for SMB shares

Change notifications are sent for changes to the parent directory on which the `changenotify` property is set and for changes to all of the subdirectories in that parent directory.

- FabricPool
- Quota enforcement
- Qtree statistics
- Adaptive QoS for files in FlexGroup volumes
- FlexCache (cache only; FlexGroup as origin supported in ONTAP 9.7)

Features supported beginning with ONTAP 9.4

- FPolicy

- File auditing
- Throughput floor (QoS Min) and adaptive QoS for FlexGroup volumes
- Throughput ceiling (QoS Max) and throughput floor (QoS Min) for files in FlexGroup volumes

You use the `volume file modify` command to manage the QoS policy group that is associated with a file.

- Relaxed SnapMirror limits
- SMB 3.x multichannel

Features supported in ONTAP 9.3 and earlier

- Antivirus configuration
- Change notifications for SMB shares

Notifications are sent only for changes to the parent directory on which the `changenotify` property is set. Change notifications are not sent for changes to subdirectories in the parent directory.

- Qtrees
- Throughput ceiling (QoS Max)
- Expand the source FlexGroup volume and destination FlexGroup volume in a SnapMirror relationship
- SnapVault backup and restore
- Unified data protection relationships
- Autogrow option and autoshrink option
- Inode count factored to ingest
- Volume encryption
- Aggregate inline deduplication (cross-volume deduplication)
- [NetApp volume encryption \(NVE\)](#)
- SnapMirror technology
- Snapshots
- Digital Advisor
- Inline adaptive compression
- Inline deduplication
- Inline data compaction
- AFF
- Quota reporting
- NetApp Snapshot technology
- SnapRestore software (FlexGroup level)
- Hybrid aggregates
- Constituent or member volume move
- Postprocess deduplication
- NetApp RAID-TEC technology

- Per-aggregate consistency point
- Sharing FlexGroup with FlexVol volume in the same SVM


Unsupported FlexGroup volume configurations in ONTAP 9

Unsupported protocols	Unsupported data protection features	Other unsupported ONTAP features
<ul style="list-style-type: none"> • pNFS (ONTAP 9.6 and earlier) • SMB 1.0 • SMB transparent failover (ONTAP 9.5 and earlier) • SAN 	<ul style="list-style-type: none"> • SnapLock volumes (ONTAP 9.10.1 and earlier) • SMTape • SnapMirror synchronous • SVM DR with FlexGroup volumes containing FabricPools (ONTAP 9.11.1 and earlier) 	<ul style="list-style-type: none"> • Remote Volume Shadow Copy Service (VSS) • SVM data mobility

FlexGroup volume setup

ONTAP FlexGroup volume setup workflow

You can either provision a FlexGroup volume where ONTAP automatically selects the aggregates based on the best practices for optimum performance, or create a FlexGroup volume by manually selecting the aggregates and configuring it for data access.



Before you begin

You must have created the SVM with NFS and SMB added to the list of allowed protocols for the SVM.

About this task

You can automatically provision a FlexGroup volume only on clusters with four nodes or less. On clusters with more than four nodes, you must create a FlexGroup volume manually.

Enable 64-bit NFSv3 identifiers on ONTAP SVMs with FlexGroups

To support the high file count of FlexGroup volumes and to avoid file ID collisions, you should enable 64-bit file identifiers on the SVM on which the FlexGroup volume must be created.

Steps

1. Log in to the advanced privilege level: `set -privilege advanced`
2. Modify the SVM to use 64-bit NFSv3 FSIDs and file IDs: `vserver nfs modify -vserver svm_name -v3-64bit-identifiers enabled`

```
cluster1::>*> vserver nfs modify -vserver vs0 -v3-64bit-identifiers
enabled
```

```
Warning: You are attempting to increase the number of bits used for
NFSv3
```

```
FSIDs and File IDs from 32 to 64 on Vserver "vs0". This could
result in older client software no longer working with the
volumes
```

```
owned by Vserver "vs0".
```

```
Do you want to continue? {y|n}: y
```

```
Warning: Based on the changes you are making to the NFS server on
Vserver
```

```
"vs0", it is highly recommended that you remount all NFSv3
clients
```

```
connected to it after the command completes.
```

```
Do you want to continue? {y|n}: y
```

After you finish

All of the clients must be remounted. This is required because the file system IDs change, and the clients might receive stale file handle messages when attempting NFS operations.

Provision an ONTAP FlexGroup volume automatically

When you create a FlexGroup volume, you can choose to have ONTAP automatically provision the FlexGroup volume by selecting the underlying local tiers (aggregates). Local tiers are selected based on the best practices for optimum performance and capacity.

Before you begin

Each node in the cluster must have at least one local tier.



When creating a FlexGroup volume that will tier inactive data, each node must have at least one local tier with FabricPool enabled.

About this task

ONTAP selects two local tiers with the largest amount of usable space on each node to create the FlexGroup volume. If two local tiers are not available, ONTAP selects one local tier per node to create the FlexGroup volume.

Beginning with ONTAP 9.15.1, when you automatically provision a FlexGroup volume, ONTAP uses balanced placement (BP) to choose the local tiers and FlexGroup member (constituent) volumes layout. One aspect of BP is how it limits over-provisioning local tiers when creating 'none' guaranteed (thin-provisioned) FlexGroup volumes. The size of the overall FlexGroup volume is limited by the amount of free space on the local tier, although the limit is higher than it is for 'volume' guaranteed (thick-provisioned) FlexGroup volumes. When you create a FlexGroup volume using REST APIs or auto-provision-as with the ONTAP CLI, provisioning might fail because of insufficient space due to this limit. You can avoid this by creating smaller FlexGroup volumes, or by [creating a FlexGroup volume and selecting the local tiers manually](#) using the aggr-list parameter.

Steps

1. Provision the FlexGroup volume:

```
volume create -vserver svm_name -volume fg_vol_name -auto-provision-as
flexgroup -size fg_size [-encrypt true] [-qos-policy-group
qos_policy_group_name] [-support-tiering true] [-granular-data advanced]
```

Beginning with ONTAP 9.16.1, you can enable [advanced capacity balancing](#) (-granular-data advanced in the CLI) to write data across multiple FlexGroup member volumes when files are larger than 10GB.

Beginning with ONTAP 9.5, you can create FlexGroup volumes on local tiers with FabricPool enabled. To automatically provision a FlexGroup volume on local tiers with FabricPool enabled, you must set the -support-tiering parameter to true. The volume guarantee must be always set to none for FabricPool. You can also specify the tiering policy and tiering minimum cooling period for the FlexGroup volume.

Disk and aggregate management

You can specify a throughput ceiling (QoS Max) for FlexGroup volumes. This limits the performance resources that the FlexGroup volume can consume. Beginning with ONTAP 9.4, you can specify throughput floors (QoS Min) and adaptive QoS for FlexGroup volumes.

Performance management

You can set the -encrypt parameter to true if you want to enable encryption on the FlexGroup volume. For creating an encrypted volume, you must have installed the volume encryption license and the key manager.



You must enable encryption on FlexGroup volumes at the time of creation. You cannot enable encryption on existing FlexGroup volumes.

Encryption of data at rest

The size parameter specifies the size of the FlexGroup volume in KB, MB, GB, TB, or PB.

The following example shows how to provision a FlexGroup volume of size 400 TB:

```
cluster-1::> volume create -vserver vs0 -volume fg -auto-provision-as
flexgroup -size 400TB
Warning: The FlexGroup "fg" will be created with the following number of
constituents of size 25TB: 16.
The constituents will be created on the following aggregates:
aggr1,aggr2
Do you want to continue? {y|n}: y
[Job 34] Job succeeded: Successful
```

The following example shows how to create a QoS policy group for throughput ceiling and how to apply it to a FlexGroup volume:

```
cluster1::> qos policy-group create -policy group pg-vs1 -vserver vs1
-max-throughput 5000iops
```

```
cluster-1::> volume create -vserver vs0 -volume fg -auto-provision-as
flexgroup -size 400TB -qos-policy-group pg-vs1
Warning: The FlexGroup "fg" will be created with the following number of
constituents of size 25TB: 16.
The constituents will be created on the following aggregates:
aggr1,aggr2
Do you want to continue? {y|n}: y
[Job 34] Job succeeded: Successful
```

The following example shows how to provision a FlexGroup volume of size 400 TB on local tiers with FabricPool enabled:

```
cluster-1::> volume create -vserver vs0 -volume fg -auto-provision-as
flexgroup -size 400TB -support-tiering true -tiering-policy auto
Warning: The FlexGroup "fg" will be created with the following number of
constituents of size 25TB: 16.
The constituents will be created on the following aggregates:
aggr1,aggr2
Do you want to continue? {y|n}: y
[Job 34] Job succeeded: Successful
```

The FlexGroup volume is created with eight member volumes on each node in the cluster. The member volumes are distributed equally between the two largest local tiers on each node.

By default, the FlexGroup volume is created with the `volume` space guarantee setting except on AFF systems. For AFF systems, by default the FlexGroup volume is created with the `none` space guarantee.

2. Mount the FlexGroup volume with a junction path:

```
volume mount -vserver vserver_name -volume vol_name -junction-path
junction_path
```

```
cluster1::> volume mount -vserver vs0 -volume fg2 -junction-path /fg2
```

After you finish

You should mount the FlexGroup volume from the client.

If you are running ONTAP 9.6 or earlier and if the storage virtual machine (SVM) has both NFSv3 and NFSv4 configured, mounting the FlexGroup volume from the client might fail. In such cases, you must explicitly specify the NFS version when mounting the FlexGroup volume from the client.

```
# mount -t nfs -o vers=3 192.53.19.64:/fg2 /mnt/fg2
# ls /mnt/fg2
file1  file2
```

Related information

- [qos policy-group create](#)

Create ONTAP FlexGroup volumes

You can create a FlexGroup volume by manually selecting the local tiers (aggregates) on which the FlexGroup volume must be created, and then specifying the number of member volumes (constituents) on each local tier.

Alternatively, you can choose to have ONTAP [automatically provision](#) the FlexGroup volume by selecting the local tiers and letting ONTAP set the number of member volumes based on the best practices for optimum performance and capacity.

About this task

You must be aware of the space required in the local tiers for creating a FlexGroup volume.

You must consider the following guidelines when creating a FlexGroup volume for obtaining the best performance results with a FlexGroup volume:

- A FlexGroup volume should use local tiers that are on identical hardware systems.

The use of identical hardware systems helps in providing predictable performance across the FlexGroup volume. Note: C-Series r1 and C-Series r2 systems are not identical systems. For example, the AFF C80 r1 and AFF C80 r2 are not identical.

- A FlexGroup volume should span local tiers using the same disk type and RAID group configurations.

For consistent performance, you must ensure that all of the local tiers are made of all SSDs, all HDDs, or all Flash Pool (hybrid) local tiers. Additionally, the local tiers should have the same number of drives and RAID groups across the FlexGroup volume.

- A FlexGroup volume can span parts of a cluster.

A FlexGroup volume does not have to be configured to span the entire cluster, but doing so can take greater advantage of the hardware resources that are available.

- When creating a FlexGroup volume, it is best if the local tiers on which the FlexGroup volume is deployed have the following characteristics:
 - Approximately the same amount of free space should be available across multiple local tier, especially when using thin provisioning.
 - Approximately 3 percent of the free space should be reserved for local tier metadata after creation of the FlexGroup volume.
- For FAS systems, it is best to have two local tiers per node and for AFF systems, you must have one local tier per node for the FlexGroup volume.
- For each FlexGroup volume, you should create at least eight member volumes that are distributed over two or more local tiers on FAS systems, and over one or more local tiers on AFF systems.
- Beginning with ONTAP 9.9.1, SnapMirror fanout relationships of two or more FlexGroup volumes are supported, with a maximum of eight fanout legs. System Manager does not support SnapMirror cascading FlexGroup volume relationships.
- When you use System Manager to create a FlexGroup volume, ONTAP automatically selects the local tiers required for creating the FlexGroup volume.
- Beginning with ONTAP 9.8, when you provision storage, QoS is enabled by default. You can disable QoS, or choose a custom QoS policy during the provisioning process or at a later time.

Before you begin

- Beginning with 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](#). Learn more about `volume create` in the [ONTAP command reference](#).

System Manager

Using System Manager, you can create a FlexGroup volume.

Steps

1. Navigate to **Storage > Volumes** and select **+ Add**.
2. In the **Add volume** window, enter a volume name and size, then select **More options**.
3. In the **Storage and optimization** section, select **Distribute volume data across the cluster (FlexGroup)**.



If you are running ONTAP 9.8 or later and you want to disable QoS or choose a custom QoS policy, click **More Options**, and then under **Storage and Optimization**, select **Performance Service Level**.

4. Complete the remaining information for the volume and select **Save**.

CLI

1. Create the FlexGroup volume:

```
volume create -vserver <svm_name> -volume <flexgroup_name> -aggr
-list aggr1,aggr2,... -aggr-list-multiplier <constituents_per_aggr>
-size <fg_size> [-encrypt true] [-qos-policy-group
qos_policy_group_name] [-granular-data advanced]
```

- The **-aggr-list** parameter specifies the list of local tiers to be used for FlexGroup member volumes.

For consistent performance across the FlexGroup volume, all of the local tiers must use the same disk type and RAID group configurations.

- The **-aggr-list-multiplier** parameter specifies the number of member volumes that will be created on each local tier listed with the **-aggr-list** parameter.

The default value of the **-aggr-list-multiplier** parameter is 4.

- The **size** parameter specifies the size of the FlexGroup volume in KB, MB, GB, TB, or PB.

◦

Beginning with ONTAP 9.16.1, you can enable [advanced capacity balancing](#) (**-granular-data advanced** in the CLI) to write data across multiple FlexGroup member volumes when files are larger than 10GB.

- Beginning with ONTAP 9.5, you can create FlexGroup volumes using local tiers with FabricPool enabled.

To create a FlexGroup volume for FabricPool, all the local tiers specified with the **-aggr-list** parameter must have FabricPool enabled. The volume guarantee must be always set to none when using FabricPool. You can also specify the tiering policy and tiering minimum cooling period for the FlexGroup volume.

[Disk and aggregate management](#)

- Beginning with ONTAP 9.4, you can specify throughput floors (QoS Min) and adaptive QoS for FlexGroup volumes.

Performance management

- You can specify a throughput ceiling (QoS Max) for FlexGroup volumes, which limits the performance resources that the FlexGroup volume can consume.
- You can set the `-encrypt` parameter to `true` if you want to enable encryption on the FlexGroup volume.

For creating an encrypted volume, you must have installed the volume encryption license and the key manager.



You must enable encryption on FlexGroup volumes at the time of creation. You cannot enable encryption on existing FlexGroup volumes.

Encryption of data at rest

```
cluster-1::> volume create -vserver vs0 -volume fg2 -aggr-list
aggr1,aggr2,aggr3,aggr1 -aggr-list-multiplier 2 -size 500TB
```

Warning: A FlexGroup "fg2" will be created with the following number of constituents of size 62.50TB: 8.

Do you want to continue? {y|n}: y

[Job 43] Job succeeded: Successful

In the previous example, if you want to create the FlexGroup volume for FabricPool, all local tiers (aggr1, aggr2, and aggr3) must have FabricPool enabled. Mount the FlexGroup volume with a junction path:

```
volume mount -vserver vserver_name -volume vol_name -junction-path
junction_path
```

```
cluster1::> volume mount -vserver vs0 -volume fg2 -junction-path /fg
```

After you finish

You should mount the FlexGroup volume from the client.

If you are running ONTAP 9.6 or earlier and if the storage virtual machine (SVM) has both NFSv3 and NFSv4 configured, mounting the FlexGroup volume from the client might fail. In such cases, you must explicitly specify the NFS version when you are mounting the FlexGroup volume from the client.

```
# mount -t nfs -o vers=3 192.53.19.64:/fg /mnt/fg2
# ls /mnt/fg2
file1  file2
```

Related information

[NetApp Technical Report 4571: NetApp FlexGroup Best Practices and Implementation Guide](#)

Manage FlexGroup volumes

Monitor the space usage of ONTAP FlexGroup volumes

You can view a FlexGroup volume and its constituents, and monitor the space used by the FlexGroup volume.

About this task

Beginning with ONTAP 9.6, elastic sizing is supported. ONTAP automatically grows a constituent of a FlexGroup volume if it is running out of space by shrinking any other constituent in the FlexGroup volume that has free space by an equivalent amount. Elastic sizing avoids any out-of-space errors that are generated because of one or more FlexGroup constituent volumes running out of space.



Beginning with ONTAP 9.9.1, logical space reporting and enforcement is also available for FlexGroup volumes. For more information, see [Logical space reporting and enforcement for volumes](#).

Step

1. View the space used by the FlexGroup volume and its constituents: `volume show -vserver vserver_name -volume-style-extended [flexgroup | flexgroup-constituent]`

```
cluster-2::> volume show -vserver vs1 -volume-style-extended flexgroup
Vserver      Volume      Aggregate      State      Type      Size
Available    Used%
-----  -----
-----  -----
vs1          fg1          -            online    RW      500GB
207.5GB     56%
```

```
ccluster-2::> volume show -vserver vs1 -volume-style-extended flexgroup-constituent
Vserver      Volume          Aggregate      State      Type      Size
Available    Used%
-----  -----
vs1          fg1_0001        aggr3          online     RW       31.25GB
12.97GB      56%
vs1          fg1_0002        aggr1          online     RW       31.25GB
12.98GB      56%
vs1          fg1_0003        aggr1          online     RW       31.25GB
13.00GB      56%
vs1          fg1_0004        aggr3          online     RW       31.25GB
12.88GB      56%
vs1          fg1_0005        aggr1          online     RW       31.25GB
13.00GB      56%
vs1          fg1_0006        aggr3          online     RW       31.25GB
12.97GB      56%
vs1          fg1_0007        aggr1          online     RW       31.25GB
13.01GB      56%
vs1          fg1_0008        aggr1          online     RW       31.25GB
13.01GB      56%
vs1          fg1_0009        aggr3          online     RW       31.25GB
12.88GB      56%
vs1          fg1_0010        aggr1          online     RW       31.25GB
13.01GB      56%
vs1          fg1_0011        aggr3          online     RW       31.25GB
12.97GB      56%
vs1          fg1_0012        aggr1          online     RW       31.25GB
13.01GB      56%
vs1          fg1_0013        aggr3          online     RW       31.25GB
12.95GB      56%
vs1          fg1_0014        aggr3          online     RW       31.25GB
12.97GB      56%
vs1          fg1_0015        aggr3          online     RW       31.25GB
12.88GB      56%
vs1          fg1_0016        aggr1          online     RW       31.25GB
13.01GB      56%
16 entries were displayed.
```

You can use the available space and percentage space used to monitor the space usage of the FlexGroup volume.

Increase the size of ONTAP FlexGroup volumes

You can increase the size of a FlexGroup volume either by adding more capacity to all existing member volumes (constituents) of the FlexGroup volume or by expanding the FlexGroup volume with new member volumes. A FlexGroup volume cannot have more than 200 member volumes.

You can also increase the size of an individual volume within a FlexGroup volume if needed.

Before you begin

Sufficient space must be available in the aggregates.

About this task

If you want to add more space, you can increase the collective size of the FlexGroup volume. Increasing the size of a FlexGroup volume resizes the existing member volumes of the FlexGroup volume.

If you want to improve performance, you can expand the FlexGroup volume. You might want to expand a FlexGroup volume and add new member volumes in the following situations:

- New nodes have been added to the cluster.
- New local tiers (aggregates) have been created on the existing nodes.
- The existing member volumes of the FlexGroup volume have reached the maximum FlexVol size for the hardware (100TB or 300TB if [large volume support](#) has been enabled), and therefore the FlexGroup volume cannot be resized without adding additional member volumes.

If you modify a FlexGroup volume to include more members, previously created snapshots are considered "partial" and are only available for access by clients from the `.snapshot` directory or the **Previous Versions** tab.



If a snapshot is considered "partial", it cannot be used in SnapRestore operations. However, partial snapshots can be used to restore individual files from `.snapshot` directories or the **Previous Versions** tab.

In releases earlier than ONTAP 9.3, do not expand FlexGroup volumes after a SnapMirror relationship is established. If you expand the source FlexGroup volume after breaking the SnapMirror relationship in releases earlier than ONTAP 9.3, you must perform a baseline transfer to the destination FlexGroup volume once again. Beginning with ONTAP 9.3, you can expand FlexGroup volumes that are in a SnapMirror relationship.

Steps

1. Increase the size of the FlexGroup volume by increasing the capacity or performance of the FlexGroup volume, as required:

If you want to increase the...	Then do this...
Capacity of the FlexGroup volume	Resize all the member volumes of the FlexGroup volume: <code>volume modify -vserver <svm_name> -volume <fg_name> -size <new_size></code>

<p>Performance to the FlexGroup volume</p>	<p>Expand the FlexGroup volume by adding new member volumes (constituents):</p> <pre>volume expand -vserver vserver_name -volume fg_name -aggr-list aggregate name,... [-aggr-list-multiplier constituents_per_aggr]</pre> <p>The default value of the <code>-aggr-list</code> <code>-multiplier</code> parameter is 1.</p> <p>When expanding a FlexGroup volume using FabricPool, all local tiers (aggregates) must be attached to the same cloud tier.</p>
--	--

Assuming existing aggregates (local tiers) or member volumes have not reached their maximum capacities (100/300TB or two billion files each), it is preferable to increase the overall size of the FlexGroup volume rather than adding additional member volumes.

Use `volume expand` only if increasing the existing volume size or file count is not an option or if the FlexGroup is being expanded across new hardware. The same number of member volumes should be added to all nodes in order to ensure consistent performance. For example, if an existing FlexGroup volume has 8 member volumes with four member volumes per node, adding two members per node will result in 12 member volumes, six member volumes per node.

When adding new members to new nodes, try to maintain a consistent number of member volumes per node as in the existing nodes. For example, if an existing FlexGroup volume has 8 member volumes with four member volumes per node, if the FlexGroup volumes is expanded to the new node, four member volumes should be added, resulting in a 12 member FlexGroup volume.

Adding new members to a FlexGroup volume changes the ingest heuristics to favor the new, empty, member volumes and can affect overall system performance for new data ingest until the new member volumes become balanced with pre-existing member volumes.

Examples

Example of increasing the capacity of the existing member volumes

The following example shows how to add 20 TB space to a FlexGroup volume `volX`:

```
cluster1::> volume modify -vserver svm1 -volume volX -size +20TB
```

If the FlexGroup volume has 16 member volumes, the space of each member volumes is increased by 1.25 TB.

Example of improving performance by adding new member volumes

The following example shows how to add four additional member volumes, two to each of the underlying local tiers (aggregates) to FlexGroup volume `fg1`:

```
cluster1::> volume expand -vserver svm1 -volume fg1 -aggr-list aggr1,aggr2  
-aggr-list-multiplier 2
```

The size of the new member volumes is the same as that of the existing member volumes.

Increase the size of an individual volume

If you want to increase the size of an individual member volume in a FlexGroup volume, you can use the `volume resize` command.

Step

1. Increase the size of a single FlexGroup member volume:

```
volume size -volume <volume_name> -vserver <svm1> -new-size <new_size>
```

The following example increases the size of FlexGroup member volume FG_0003 to 3.7GB:

```
volume size -volume FG_0003 -vserver svm1 -new-size 3.7GB  
vol size: Volume "svm1:FG_0003" size set to 3.70g.
```

Reduce the size of ONTAP FlexGroup volumes

Beginning with ONTAP 9.6, you can resize a FlexGroup volume to a value lower than its current size to free up the unused space from the volume. When you reduce the size of a FlexGroup volume, ONTAP automatically resizes all of the FlexGroup constituents.

Step

1. Check the current FlexGroup volume size: `'volume size -vserver vserver_name -volume fg_name'`
2. Reduce the size of the FlexGroup volume: `volume size -vserver vserver_name -volume fg_name new_size`

When you specify the new size, you can specify either a lower value than the current size or a negative value using the minus sign (-) by which the current size of the FlexGroup volume is reduced.



If automatic shrinking is enabled for the volume (`volume autosize` command), the minimum autosize is set to the new size of the volume.

The following example displays the current volume size for the FlexGroup volume named volX and resizes the volume to 10TB:

```
cluster1::> volume size -vserver svm1 -volume volX
(volume size)
vol size: FlexGroup volume 'svm1:volX' has size 15TB.

cluster1::> volume size -vserver svm1 -volume volX 10TB
(volume size)
vol size: FlexGroup volume 'svm1:volX' size set to 10TB.
```

The following example displays the current volume size for the FlexGroup volume named volX and reduces the size of the volume by 5TB:

```
cluster1::> volume size -vserver svm1 -volume volX
(volume size)
vol size: FlexGroup volume 'svm1:volX' has size 15TB.

cluster1::> volume size -vserver svm1 -volume volX -5TB
(volume size)
vol size: FlexGroup volume 'svm1:volX' size set to 10TB.
```

Configure ONTAP FlexGroup volumes to automatically grow and shrink their size

Beginning with ONTAP 9.3, you can configure FlexGroup volumes to automatically grow and shrink according to how much space they currently require.

Before you begin

The FlexGroup volume must be online.

About this task

You can autosize FlexGroup volumes in two modes:

- Increase the size of the volume automatically (grow mode)

Automatic growing helps prevent a FlexGroup volume from running out of space, if the aggregate can supply more space. You can configure the 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.

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.

- Shrink the size of the volume automatically (grow_shrink mode)

Automatic shrinking prevents a volume from being larger than needed, freeing space in the aggregate for use by other volumes.

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.

Steps

1. Configure the volume to grow and shrink its size automatically: `volume autosize -vserver vserver_name -volume vol_name -mode [grow | grow_shrink]`

You can also specify the maximum size, minimum size, and thresholds for growing or shrinking the volume.

The following command enables automatic size changes for a volume called `fg1`. The volume is configured to grow to a maximum size of 5 TB when it is 70% full.

```
cluster1::> volume autosize -volume fg1 -mode grow -maximum-size 5TB  
-grow-threshold-percent 70  
vol autosize: volume "vs_src:fg1" autosize settings UPDATED.
```

Delete directories asynchronously from ONTAP FlexGroup volumes

Beginning with ONTAP 9.8, you can delete directories from Linux and Windows client shares asynchronously (that is, in the background). Cluster and SVM administrators can perform asynchronous delete operations on both FlexVol and FlexGroup volumes.

About this task

You must be a cluster administrator or an SVM administrator using the advanced privilege mode.

Beginning with ONTAP 9.8, you can use asynchronous delete functionality using the ONTAP CLI. Beginning with ONTAP 9.9.1, you can use this functionality with System Manager. For more information about this process, see [Take corrective action based on ONTAP analytics in FSA](#).

Beginning with ONTAP 9.11.1, a storage administrator can grant rights on a volume to allow NFS and SMB clients to perform asynchronous delete operations. For more information, see [Manage client rights to delete directories asynchronously](#).

You can use the `volume file async-delete show` command to check the status of in-progress asynchronous delete jobs, and, beginning with ONTAP 9.17.1, the status of asynchronous delete jobs issued from clients is also displayed.

Delete directories asynchronously

You can use System Manager or the ONTAP CLI to delete directories asynchronously.

System Manager

Beginning with ONTAP 9.10.1	In ONTAP 9.9.1
<ol style="list-style-type: none">1. Select Storage > Volumes and select the desired volume name.2. In the individual volume page, select the File system tab, and then select the Explorer tab.3. In the Explorer view, select the desired directory.4. To delete, hover over a file or folder, and the delete  option appears. <p>You can only delete one object at a time.</p> <p> When directories and files are deleted, the new storage capacity values are not displayed immediately.</p>	<ol style="list-style-type: none">1. Select Storage > Volumes.2. Select the desired volume, then select Explorer.3. In the Explorer view, select the desired directory.4. To delete, hover over a file or folder, and the delete  option appears.

CLI

Use the CLI to perform an asynchronous delete

1. Enter advanced privilege mode:

```
set -privilege advanced
```

2. Delete directories on a FlexVol or FlexGroup volume:

```
volume file async-delete start -vserver <SVM_name> -volume <volume_name>
-path <file_path> -throttle <throttle>
```

The minimum throttle value is 10, the maximum is 100,000, and the default is 5000. Lower throttle values use less resources, which can result in a slower deletion rate, while higher throttle values use more resources, but can result in a faster deletion rate.

The following example deletes the directory named d2, which is located in the directory named d1.

```
cluster::>*> volume file async-delete start -vserver vs1 -volume vol1
-path d1/d2
```

3. (Optional) Check the status of the in-progress async delete jobs:

```
volume file async-delete show
```

4. Verify that the directory was deleted:

```
event log show
```

The following example shows output for the event log when the directory is successfully deleted.

```
cluster::>*> event log show

Time           Node           Severity      Event
-----
-----
7/7/2025 09:04:04  cluster-vsimg    NOTICE
asyncDelete.message.success: Async delete job on path d1/d2 of
volume (MSID: 2162149232) was completed. Number of files deleted: 7,
Number of directories deleted: 5. Total number of bytes deleted:
135168.
```

Learn more about `event log show` in the [ONTAP command reference](#).

Cancel a directory delete job

1. Enter advanced privilege mode:

```
set -privilege advanced
```

2. Verify that the directory delete is in progress:

```
volume file async-delete show
```

If the SVM, volume, JobID, and path of your directory is displayed, you can cancel the job.

3. Cancel the directory delete:

```
volume file async-delete cancel -vserver <SVM_name> -volume <volume_name>
-jobid <job_id>
```

Manage client rights to delete ONTAP directories asynchronously with FlexGroups

Beginning with ONTAP 9.11.1, storage administrators can grant rights on a volume to allow NFS and SMB clients to perform asynchronous delete operations. When asynchronous delete is enabled on the cluster, Linux client users can use the `mv` command and Windows client users can use the `rename` command to delete a directory on the specified volume by moving it to a hidden directory that by default is named `.ontaptrashbin`.

Rights are granted per volume. NFS client users should have root access on the NFS client and superuser access to the NFS export.

You can move directories only. You cannot move files to the `.ontaptrashbin` directory.

[Learn about using ONTAP to delete directories asynchronously from FlexGroup volumes.](#)

Enable client asynchronous directory delete

Steps

1. From the cluster CLI, enter advanced privilege mode: `-privilege advance`
2. Enable client asynchronous delete at the mountpoint of a volume and, if desired, provide an alternate name for the trashbin directory:

```
volume file async-delete client enable volume volname vserver vserverName  
trashbinname name
```

Example using the default trashbin name:

```
cluster1::*> volume file async-delete client enable -volume v1 -vserver  
vs0  
  
Info: Async directory delete from the client has been enabled on volume  
"v1" in  
Vserver "vs0".
```

Example specifying an alternate trashbin name:

```
cluster1::*> volume file async-delete client enable -volume test  
-trashbin .ntaptrash -vserver vs1  
  
Success: Async directory delete from the client is enabled on volume  
"v1" in  
Vserver "vs0".
```

3. Verify client asynchronous delete is enabled:

```
volume file async-delete client show
```

Example:

```
cluster1::*> volume file async-delete client show  
  
Vserver Volume      async-delete client TrashBinName  
-----  
vs1      vol1        Enabled        .ntaptrash  
vs2      vol2        Disabled      -  
  
2 entries were displayed.
```

Disable client asynchronous directory delete

Steps

1. From the cluster CLI, disable client asynchronous directory delete:

```
volume file async-delete client disable volume volname vserver vserverName
```

Example:

```
cluster1::*> volume file async-delete client disable -volume vol1  
-vserver vs1
```

```
Success: Asynchronous directory delete client disabled  
successfully on volume.
```

2. Verify client asynchronous delete is disabled:

```
volume file async-delete client show
```

Example:

```
cluster1::*> volume file async-delete client show
```

Vserver	Volume	async-delete client	TrashBinName
vs1	vol1	Disabled	-
vs2	vol2	Disabled	-

```
2 entries were displayed.
```

Create qtrees with ONTAP FlexGroup volumes

Beginning with ONTAP 9.3, you can create qtrees with FlexGroup volumes. Qtrees enable you to partition your FlexGroup volumes into smaller segments that you can manage individually.

About this task

- If the source FlexGroup volume has qtrees in a SnapMirror relationship, the destination cluster must be running ONTAP 9.3 or later (a version of ONTAP software that supports qtrees).
- Beginning with ONTAP 9.5, qtree statistics are supported for FlexGroup volumes.

Steps

1. Create a qtree in the FlexGroup volume:

```
volume qtree create -vserver <vserver_name> -volume <volume_name> -qtree <qtree_name>
```

You can optionally specify the security style, SMB oplocks, UNIX permissions, and export policy for the qtree.

```
cluster1::> volume qtree create -vserver vs0 -volume fg1 -qtree qtree1  
-security-style mixed
```

Related information

[Logical storage management](#)

Use quotas for ONTAP FlexGroup volumes

In ONTAP 9.4 and earlier, you can apply quota rules to FlexGroup volumes only for reporting purposes, but not for enforcing quota limits. Beginning with ONTAP 9.5, you can enforce limits on quota rules that are applied to FlexGroup volumes.

About this task

- Beginning with ONTAP 9.5, you can specify hard, soft, and threshold limit quotas for FlexGroup volumes.

You can specify these limits to constrain the amount of space, the number of files that a specific user, group, or qtree can create, or both. Quota limits generate warning messages in the following scenarios:

- When usage exceeds a configured soft limit, ONTAP issues a warning message, but further traffic is still allowed.

If usage later drops below the configured soft limit again, an all-clear message is issued.

- When usage exceeds a configured threshold limit, ONTAP issues a second warning message.

No all-clear administrative message is issued when usage later drops below a configured threshold limit.

- If usage reaches a configured hard limit, ONTAP prevents further resource consumption by rejecting traffic.

- In ONTAP 9.5, quota rules cannot be created or activated on the destination FlexGroup volume of a SnapMirror relationship.
- During quota initialization, quotas are not enforced, and there are no notifications of breached quotas following quota initialization.

To check if quotas were breached during quota initialization, you can use the `volume quota report` command.

Quota targets and types

Quotas have a type: they can be either user, group, or tree. Quota targets specify the user, group, or qtree for which the quota limits are applied.

The following table lists the kinds of quota targets, what types of quotas each quota target is associated with, and how each quota target is represented:

Quota target	Quota type	How target is represented	Notes
user	user quota	UNIX user name UNIX UID Windows user name in pre-Windows 2000 format Windows SID	User quotas can be applied for a specific volume or qtree.
group	group quota	UNIX group name UNIX GID	Group quotas can be applied for a specific volume or qtree.  ONTAP does not apply group quotas based on Windows IDs.
qtree	tree quota	qtree name	Tree quotas are applied to a particular volume and do not affect qtrees in other volumes.
""	user quota group quota tree quota	Double quotation marks ("")	A quota target of "" denotes a <i>default quota</i> . For default quotas, the quota type is determined by the value of the type field.

Behavior of FlexGroup volumes when quota limits are exceeded

Beginning with ONTAP 9.5, quota limits are supported on FlexGroup volumes. There are some differences in the way quota limits are enforced on a FlexGroup volume when compared to a FlexVol volume.

FlexGroup volumes might show the following behaviors when the quota limits are exceeded:

- The space and file usage in a FlexGroup volume might reach up to 5 percent higher than the configured hard limit before the quota limit is enforced by rejecting further traffic.

To provide the best performance, ONTAP might allow the space consumption to exceed the configured hard limit by a small margin before the quota enforcement begins. This additional space consumption does not exceed 5 percent of the configured hard limits, 1 GB, or 65536 files, whichever is lower.

- After the quota limit is reached, if a user or administrator deletes some files or directories such that the quota usage is now below the limit, the subsequent quota-consuming file operation might resume with a delay (might take up to 5 seconds to resume).
- When the total space and file usage of a FlexGroup volume exceed the configured quota limits, there might be a slight delay in logging an event log message.
- You might get “no space” errors if some constituents of the FlexGroup volume get full, but the quota limits are not reached.
- Operations, such as renaming a file or directory or moving files between qtrees, on quota targets, for which quota hard limits are configured, might take longer when compared to similar operations on FlexVol volumes.

Examples of quota enforcement for FlexGroup volumes

You can use the examples to understand how to configure quotas with limits in ONTAP 9.5 and later.

Example 1: Enforcing a quota rule with disk limits

1. You should create a quota policy rule of type `user` with both an achievable soft disk limit and hard disk limit.

```
cluster1::> volume quota policy rule create -vserver vs0 -policy-name
default -volume FG -type user -target "" -qtree "" -disk-limit 1T -soft
-disk-limit 800G
```

2. You can view the quota policy rule:

```
cluster1::> volume quota policy rule show -vserver vs0 -policy-name
default -volume FG
```

			Vserver: vs0		Policy: default		Volume: FG	
Type	Target	Qtree	User	Disk	Disk	Files	Files	
Threshold			Mapping	Limit	Limit	Limit	Limit	
-----	-----	-----	-----	-----	-----	-----	-----	
-----	-----	-----	-----	-----	-----	-----	-----	
user	""	""	off	1TB	800GB	-	-	
-	-	-	-	-	-	-	-	

3. To activate the new quota rule, you initialize quotas on the volume:

```
cluster1::> volume quota on -vserver vs0 -volume FG -foreground true
[Job 49] Job succeeded: Successful
```

4. You can view the disk usage and file usage information of the FlexGroup volume by using the quota report.

```
cluster1::> volume quota report -vserver vs0 -volume FG
Vserver: vs0

-----Disk----- -----Files----- Quota
Volume   Tree     Type     ID      Used   Limit   Used   Limit
Specifier
-----  -----  -----  -----  -----  -----  -----  -----
-----  -----
FG          user    root    50GB    -      1      -      -
FG          user    *      800GB   1TB    0      -      *
2 entries were displayed.
```

After the hard disk limit is reached, the quota policy rule target (user, in this case) is blocked from writing more data to the files.

Example 2: Enforcing a quota rule for multiple users

1. You should create a quota policy rule of type `user`, where multiple users are specified in the quota target (UNIX users, SMB users, or a combination of both) and where the rule has both an achievable soft disk limit and hard disk limit.

```
cluster1::> quota policy rule create -vserver vs0 -policy-name default
-volume FG -type user -target "rdavis,ABCCORP\RobertDavis" -qtree ""
-disk-limit 1TB -soft-disk-limit 800GB
```

2. You can view the quota policy rule:

```
cluster1::> quota policy rule show -vserver vs0 -policy-name default
-volume FG

Vserver: vs0          Policy: default          Volume: FG

                                         Soft          Soft
                                         User          Disk          Disk          Files          Files
                                         Mapping
Type   Target   Qtree   Mapping   Limit   Limit   Limit   Limit
Threshold
-----  -----  -----  -----  -----  -----  -----  -----
-----  -----
user   "rdavis,ABCCORP\RobertDavis"  ""  off   1TB   800GB   -   -
```

3. To activate the new quota rule, you initialize quotas on the volume:

```
cluster1::> volume quota on -vserver vs0 -volume FG -foreground true
[Job 49] Job succeeded: Successful
```

4. You can verify that the quota state is active:

```
cluster1::> volume quota show -vserver vs0 -volume FG
  Vserver Name: vs0
  Volume Name: FG
  Quota State: on
  Scan Status: -
  Logging Messages: on
  Logging Interval: 1h
  Sub Quota Status: none
  Last Quota Error Message: -
  Collection of Quota Errors: -
```

5. You can view the disk usage and file usage information of the FlexGroup volume by using the quota report.

```
cluster1::> quota report -vserver vs0 -volume FG
Vserver: vs0

                                         ----Disk----  ----Files----  Quota
Volume   Tree      Type     ID          Used    Limit     Used    Limit
Specifier
-----  -----  -----  -----  -----  -----  -----  -----
-----  -----
FG           user     rdavis,ABCCORP\RobertDavis  0B   1TB   0   -
rdavis,ABCCORP\RobertDavis
```

The quota limit is shared among all users listed in the quota target.

After the hard disk limit is reached, users listed in the quota target are blocked from writing more data to the files.

Example 3: Enforcing quota with user mapping enabled

1. You should create a quota policy rule of type `user`, specify a UNIX user or a Windows user as the quota target with `user-mapping` set to `on`, and create the rule with both an achievable soft disk limit and hard disk limit.

The mapping between UNIX and Windows users must be configured earlier by using the `vserver name-mapping create` command.

```
cluster1::> quota policy rule create -vserver vs0 -policy-name default  
-volume FG -type user -target rdavis -qtree "" -disk-limit 1TB -soft  
-disk-limit 800GB -user-mapping on
```

2. You can view the quota policy rule:

```
cluster1::> quota policy rule show -vserver vs0 -policy-name default  
-volume FG
```

Vserver: vs0			Policy: default			Volume: FG		
Type	Target	Qtree	User Mapping	Disk Limit	Disk Limit	Files Limit	Files Limit	
Threshold			Soft			Soft		
user	rdavis	""	on	1TB	800GB	-	-	
-			-			-		

3. To activate the new quota rule, you initialize quotas on the volume:

```
cluster1::> volume quota on -vserver vs0 -volume FG -foreground true  
[Job 49] Job succeeded: Successful
```

4. You can verify that the quota state is active:

```
cluster1::> volume quota show -vserver vs0 -volume FG  
Vserver Name: vs0  
Volume Name: FG  
Quota State: on  
Scan Status: -  
Logging Messages: on  
Logging Interval: 1h  
Sub Quota Status: none  
Last Quota Error Message: -  
Collection of Quota Errors: -
```

5. You can view the disk usage and file usage information of the FlexGroup volume by using the quota report.

```

cluster1::> quota report -vserver vs0 -volume FG
Vserver: vs0

-----Disk----- -----Files----- Quota
Volume   Tree      Type     ID      Used   Limit   Used   Limit
Specifier
-----  -----  -----  -----  -----  -----  -----  -----
-----  -----
FG          user    rdavis,ABCCORP\RobertDavis  0B   1TB   0   -
rdavis

```

The quota limit is shared between the user listed in the quota target and its corresponding Windows or UNIX user.

After the hard disk limit is reached, both the user listed in the quota target and its corresponding Windows or UNIX user is blocked from writing more data to the files.

Example 4: Verifying the qtree size when quota is enabled

1. You should create a quota policy rule of type `tree` and where the rule has both an achievable soft disk limit and hard disk limit.

```

cluster1::> quota policy rule create -vserver vs0 -policy-name default
-volume FG -type tree -target tree_4118314302 -qtree "" -disk-limit 48GB
-soft-disk-limit 30GB

```

2. You can view the quota policy rule:

```

cluster1::> quota policy rule show -vserver vs0

Vserver: vs0          Policy: default          Volume: FG

                                         Soft          Soft
                                         User          Disk          Disk          Files          Files
                                         Mapping       Limit       Limit       Limit       Limit
Type   Target   Qtree   Mapping
Threshold
-----  -----  -----  -----  -----  -----  -----  -----  -----
-----  -----
tree   tree_4118314302  ""  -          48GB      -          20      -

```

3. To activate the new quota rule, you initialize quotas on the volume:

```

cluster1::> volume quota on -vserver vs0 -volume FG -foreground true
[Job 49] Job succeeded: Successful

```

a. You can view the disk usage and file usage information of the FlexGroup volume by using the quota report.

```
cluster1::> quota report -vserver vs0
Vserver: vs0
----Disk---- ----Files---- Quota
Volume Tree Type ID Used Limit Used Limit Specifier
----- -----
FG tree_4118314302 tree 1 30.35GB 48GB 14 20 tree_4118314302
```

The quota limit is shared between the user listed in the quota target and its corresponding Windows or UNIX user.

4. From an NFS client, use the `df` command to view the total space usage, available space, and the used space.

```
scsps0472342001# df -m /t/10.53.2.189/FG-3/tree_4118314302
Filesystem 1M-blocks Used Available Use% Mounted on
10.53.2.189/FG-3 49152 31078 18074 63% /t/10.53.2.189/FG-3
```

With hard limit, the space usage is calculated from an NFS client as follows:

- Total space usage = hard limit for tree
- Free space = Hard limit minus qtree space usage

Without hard limit, the space usage is calculated from an NFS client as follows:

- Space usage = quota usage
- Total space = Sum of quota usage and physical free space in the volume

5. From the SMB share, use Windows Explorer to view the total space usage, available space, and the used space.

From an SMB share, you should be aware of the following considerations for calculating the space usage:

- The user quota hard limit for the user and group is taken into consideration for calculating the total available space.
- The minimum value among the free space of the tree quota rule, the user quota rule, and the group quota rule is considered as the free space for the SMB share.
- The total space usage is variable for SMB and depends on the hard limit that corresponds to the minimum free space among the tree, user, and group.

Apply rules and limits on the FlexGroup volume

Steps

1. Create quota rules for targets : `volume quota policy rule create -vserver vs0 -policy -name quota_policy_of_the_rule -volume flexgroup_vol -type {tree|user|group} -target target_for_rule -qtree qtree_name [-disk-limit hard_disk_limit_size] [-file-limit hard_limit_number_of_files] [-threshold`

```
threshold_disk_limit_size] [-soft-disk-limit soft_disk_limit_size] [-soft-file-limit soft_limit_number_of_files]
```

- The quota target type can be user, group, or tree for FlexGroup volumes.
- A path is not supported as the target when creating quota rules for FlexGroup volumes.
- Beginning with ONTAP 9.5, you can specify hard disk limit, hard file limit, soft disk limit, soft file limit, and threshold limit quotas for FlexGroup volumes.

In ONTAP 9.4 and earlier, you cannot specify the disk limit, file limit, threshold for disk limit, soft disk limit, or soft file limit when you create quota rules for FlexGroup volumes.

The following example shows a default quota rule being created for the user target type:

```
cluster1::> volume quota policy rule create -vserver vs0 -policy-name
quota_policy_vs0_1 -volume fg1 -type user -target "" -qtree ""
```

The following example shows a tree quota rule being created for the qtree named qtree1:

```
cluster1::> volume quota policy rule create -policy-name default -vserver
vs0 -volume fg1 -type tree -target "qtree1"
```

1. Activate the quotas for the specified FlexGroup volume: `volume quota on -vserver svm_name -volume flexgroup_vol -foreground true`

```
cluster1::> volume quota on -vserver vs0 -volume fg1 -foreground true
```

1. Monitor the state of quota initialization: `volume quota show -vserver svm_name`

FlexGroup volumes might show the `mixed` state, which indicates that all of the constituent volumes are not in the same state yet.

```
cluster1::> volume quota show -vserver vs0
                                         Scan
Vserver      Volume      State      Status
-----  -----  -----
vs0          fg1        initializing  95%
vs0          vol1        off        -
2 entries were displayed.
```

1. View the quota report for the FlexGroup volume with active quotas: `volume quota report -vserver svm_name -volume flexgroup_vol`

You cannot specify a path with the `volume quota report` command for FlexGroup volumes.

The following example shows the user quota for the FlexGroup volume fg1:

```
cluster1::> volume quota report -vserver vs0 -volume fg1
Vserver: vs0
          ----Disk----  ----Files-----
Quota
  Volume  Tree      Type     ID      Used  Limit     Used  Limit
Specifier
  -----  -----  -----  -----  -----  -----  -----  -----
  -----
  fg1          user    *      0B      -      0      -      *
  fg1          user    root   1GB      -      1      -      *
2 entries were displayed.
```

The following example shows the tree quota for the FlexGroup volume fg1:

```
cluster1::> volume quota report -vserver vs0 -volume fg1
Vserver: vs0
          ----Disk----  ----Files-----  Quota
  Volume  Tree      Type     ID      Used  Limit     Used  Limit
Specifier
  -----  -----  -----  -----  -----  -----  -----  -----
  -----
  fg1      qtree1  tree     1      68KB      -      18      -
  qtree1
  fg1          tree    *      0B      -      0      -      *
2 entries were displayed.
```

Results

The quota rules and limits are applied on the FlexGroup volume.

The usage might reach up to 5 percent higher than a configured hard limit before ONTAP enforces the quota by rejecting further traffic.

Related information

- [ONTAP command reference](#)

Enable storage efficiency on ONTAP FlexGroup volumes

You can run deduplication and data compression together or independently on a FlexGroup volume to achieve optimal space savings.

Before you begin

The FlexGroup volume must be online.

Steps

1. Enable storage efficiency on the FlexGroup volume: `volume efficiency on -vserver svm_name -volume volume_name`

Storage efficiency operations are enabled on all the constituents of the FlexGroup volume.

If a FlexGroup volume is expanded after storage efficiency is enabled on the volume, storage efficiency is automatically enabled on the new constituents.

2. Enable the required storage efficiency operation on the FlexGroup volume by using the `volume efficiency modify` command.

You can enable inline deduplication, postprocess deduplication, inline compression, and postprocess compression on FlexGroup volumes. You can also set the type of compression (secondary or adaptive) and specify a schedule or efficiency policy for the FlexGroup volume.

3. If you are not using schedules or efficiency policies for running the storage efficiency operations, start the efficiency operation: `volume efficiency start -vserver svm_name -volume volume_name`

If deduplication and data compression are enabled on a volume, data compression is run initially followed by deduplication. This command fails if any efficiency operation is already active on the FlexGroup volume.

4. Verify the efficiency operations that are enabled on the FlexGroup volume: `volume efficiency show -vserver svm_name -volume volume_name`

```
cluster1::> volume efficiency show -vserver vs1 -volume fg1
    Vserver Name: vs1
    Volume Name: fg1
    Volume Path: /vol/fg1
    State: Enabled
    Status: Idle
    Progress: Idle for 17:07:25
    Type: Regular
    Schedule: sun-sat@0

    ...
    Compression: true
    Inline Compression: true
    Incompressible Data Detection: false
    Constituent Volume: false
    Compression Quick Check File Size: 524288000
    Inline Dedupe: true
    Data Compaction: false
```

Protect ONTAP FlexGroup volumes using snapshots

You can create snapshot policies that automatically manage the creation of snapshots or

you can manually create snapshots for FlexGroup volumes. A valid snapshot is created for a FlexGroup volume only after ONTAP can successfully create a snapshot for each constituent of the FlexGroup volume.

About this task

- If you have multiple FlexGroup volumes associated with a snapshot policy, you should ensure that the FlexGroup volumes schedules do not overlap.
- Beginning with ONTAP 9.8, the maximum number of snapshots supported on a FlexGroup volume is 1023.



Beginning with ONTAP 9.8, the `volume snapshot show` command for FlexGroup volumes reports snapshot size using logical blocks, rather than calculating the youngest owned blocks. This new size calculation method might make the snapshot size appear larger than calculations in earlier versions of ONTAP.

Steps

1. Create a snapshot policy or manually create a snapshot:

If you want to create a...	Enter this command...
Snapshot policy	<pre>volume snapshot policy create</pre> <p> The schedules that are associated with the snapshot policy of a FlexGroup volume must have an interval greater than 30 minutes.</p> <p>When you create a FlexGroup volume, the default snapshot policy is applied to the FlexGroup volume.</p>
Snapshot manually	<pre>volume snapshot create</pre> <p> After you create a snapshot for a FlexGroup volume, you cannot modify the attributes of the snapshot. If you want to modify the attributes, you must delete and then re-create the snapshot.</p>

Client access to the FlexGroup volume is briefly quiesced when a snapshot is created.

1. Verify that a valid snapshot is created for the FlexGroup volume: `volume snapshot show -volume volume_name -fields state`

```
cluster1::> volume snapshot show -volume fg -fields state
vserver volume snapshot state
-----
fg_vs fg hourly.2016-08-23_0505 valid
```

2. View the snapshots for the constituents of the FlexGroup volume: `volume snapshot show -is-constituent true`

```
cluster1::> volume snapshot show -is-constituent true

---Blocks---
Vserver Volume Snapshot Size Total%
Used%
-----
-----
fg_vs fg__0001 hourly.2016-08-23_0505 72MB 0%
27%
fg__0002 hourly.2016-08-23_0505 72MB 0%
27%
fg__0003 hourly.2016-08-23_0505 72MB 0%
27%
...
fg__0016 hourly.2016-08-23_0505 72MB 0%
27%
```

Move constituents from ONTAP FlexGroup volumes

You can move the constituents of a FlexGroup volume from one aggregate to another to balance the load when certain constituents experience more traffic. Moving constituents also helps in freeing up space on an aggregate for resizing the existing constituents.

Before you begin

To move a FlexGroup volume constituent that is in a SnapMirror relationship, you must have initialized the SnapMirror relationship.

About this task

You cannot perform a volume move operation while the constituents of the FlexGroup volume are being expanded.

Steps

1. Identify the FlexGroup volume constituent that you want to move:

```
volume show -vserver svm_name -is-constituent true
```

```
cluster1::> volume show -vserver vs2 -is-constituent true
Vserver      Volume      Aggregate      State      Type      Size
Available    Used%
-----
-----
```

Vserver	Volume	Aggregate	State	Type	Size
vs2	fg1	-	online	RW	400TB
15.12TB	62%				
vs2	fg1_0001	aggr1	online	RW	25TB
8.12MB	59%				
vs2	fg1_0002	aggr2	online	RW	25TB
2.50TB	90%				
...					

2. Identify an aggregate to which you can move the FlexGroup volume constituent:

```
volume move target-aggr show -vserver svm_name -volume vol_constituent_name
```

The available space in the aggregate that you select must be greater than the size of the FlexGroup volume constituent that you are moving.

```
cluster1::> volume move target-aggr show -vserver vs2 -volume fg1_0002
Aggregate Name      Available Size      Storage Type
-----
-----
```

Aggregate Name	Available Size	Storage Type
aggr2	467.9TB	hdd
node12a_aggr3	100.34TB	hdd
node12a_aggr2	100.36TB	hdd
node12a_aggr1	100.36TB	hdd
node12a_aggr4	100.36TB	hdd
5 entries were displayed.		

3. Verify that the FlexGroup volume constituent can be moved to the intended aggregate:

```
volume move start -vserver svm_name -volume vol_constituent_name -destination
-aggregate aggr_name -perform-validation-only true
```

```
cluster1::> volume move start -vserver vs2 -volume fg1_0002 -destination
-aggregate node12a_aggr3 -perform-validation-only true
Validation succeeded.
```

4. Move the FlexGroup volume constituent:

```
volume move start -vserver svm_name -volume vol_constituent_name -destination
-aggregate aggr_name [-allow-mixed-aggr-types {true|false}]
```

The volume move operation runs as a background process.

Beginning with ONTAP 9.5, you can move FlexGroup volume constituents from a Fabric Pool to a non-Fabric Pool, or vice versa by setting the `-allow-mixed-aggr-types` parameter to `true`. By default, the `-allow-mixed-aggr-types` option is set to `false`.



You cannot use the `volume move` command for enabling encryption on FlexGroup volumes.

```
cluster1::> volume move start -vserver vs2 -volume fg1_002 -destination  
-aggregate node12a_aggr3
```



If the volume move operation fails due to an active SnapMirror operation, you should abort the SnapMirror operation by using the `snapmirror abort -h` command. In some cases, the SnapMirror abort operation might also fail. In such situations, you should abort the volume move operation and retry later. Learn more about `snapmirror abort` in the [ONTAP command reference](#).

5. Verify the state of the volume move operation:

```
volume move show -volume vol_constituent_name
```

The following example shows the state of a FlexGroup constituent volume that completed the replication phase and is in the cutover phase of the volume move operation:

```
cluster1::> volume move show -volume fg1_002  
Vserver      Volume      State      Move Phase  Percent-Complete Time-To-  
Complete  
-----  
-----  
vs2          fg1_002     healthy    cutover      -          -
```

Use aggregates in FabricPool for existing ONTAP FlexGroup volumes

Beginning with ONTAP 9.5, FabricPool is supported for FlexGroup volumes. If you want to use aggregates in FabricPool for your existing FlexGroup volumes, you can either convert the aggregates on which the FlexGroup volume resides to aggregates in FabricPool or migrate the FlexGroup volume constituents to aggregates in FabricPool.

Before you begin

- The FlexGroup volume must have `space-guarantee` set to `none`.
- If you want to convert the aggregates on which the FlexGroup volume resides to aggregates in FabricPool, the aggregates must be using all SSD disks.

About this task

If an existing FlexGroup volume resides on non-SSD aggregates, you must migrate the FlexGroup volume

constituents to aggregates in FabricPool.

Choices

- To convert the aggregates on which the FlexGroup volume resides to aggregates in FabricPool, perform the following steps:

- a. Set the tiering policy on the existing FlexGroup volume: `volume modify -volume flexgroup_name -tiering-policy [auto|snapshot|none|backup]`

```
cluster-2::> volume modify -volume fg1 -tiering-policy auto
```

- b. Identify the aggregates on which the FlexGroup volume resides: `volume show -volume flexgroup_name -fields aggr-list`

```
cluster-2::> volume show -volume fg1 -fields aggr-list
vserver volume aggr-list
-----
vs1      fg1      aggr1,aggr3
```

- c. Attach an object store to each aggregate listed in the aggregate list: `storage aggregate object-store attach -aggregate aggregate_name -name object-store-name -allow-flexgroup true`

You must attach all of the aggregates to an object store.

```
cluster-2::> storage aggregate object-store attach -aggregate aggr1
-object-store-name Amazon01B1
```

- To migrate the FlexGroup volume constituents to aggregates in FabricPool, perform the following steps:

- a. Set the tiering policy on the existing FlexGroup volume: `volume modify -volume flexgroup_name -tiering-policy [auto|snapshot|none|backup]`

```
cluster-2::> volume modify -volume fg1 -tiering-policy auto
```

- b. Move each constituent of the FlexGroup volume to an aggregate in FabricPool in the same cluster: `volume move start -volume constituent-volume -destination-aggregate FabricPool_aggregate -allow-mixed-aggr-types true`

You must move all FlexGroup volume constituents to aggregates in FabricPool (in case the FlexGroup volume constituents are on mixed aggregate types) and ensure that all the constituents are balanced across the nodes in the cluster.

```
cluster-2::> volume move start -volume fg1_001 -destination-aggregate
FP_aggr1 -allow-mixed-aggr-types true
```

Related information

- [Disk and aggregate management](#)
- [storage aggregate object-store attach](#)

Balance ONTAP FlexGroup volumes by redistributing file data

Beginning with ONTAP 9.16.1, you can enable advanced capacity balancing to enable data distribution between FlexGroup member volumes when very large files grow and consume space on one member volume.

Advanced capacity balancing expands the granular data functionality introduced in ONTAP 9.12.1, which enables ONTAP to [rebalance FlexGroup volumes](#) by moving files to other members. Beginning with ONTAP 9.16.1, when advanced capacity balancing is enabled with the `-granular-data` advanced option, both the "basic" file rebalancing capabilities as well as the advanced capacity capabilities are activated.

 Both file rebalancing and advanced capacity balancing are disabled by default. After these features are enabled they cannot be disabled. If you need to disable capacity balancing, you must restore from a snapshot that was created before advanced capacity balancing was enabled.

Advanced capacity balancing is triggered by new writes reaching to 10GB or 1% of a volume's free space.

How files are distributed

If a file is created or grows large enough to trigger advanced capacity balancing, the file is distributed in stripes between 1GB and 10GB across the member FlexGroup volumes.

When advanced capacity balancing is enabled, ONTAP will not retroactively stripe existing large files. If an existing large file continues to grow after advanced capacity balancing is enabled, new content in existing large files might be striped across member FlexGroup volumes depending on the file's size and available space.

One way advanced capacity balancing determines stripe width is by using the amount of free space available on the member volume. Advanced capacity balancing creates a file stripe that is 1% of the available free space available. This means that stripes can start out larger if more space is available, and they become smaller as the FlexGroup fills up.

Supported protocols

Advanced capacity balancing is supported with the following protocols:

- NFSv3, NFSv4, NFSv4.1
- pNFS
- SMB

Enable advanced capacity balancing

Advanced capacity balancing is disabled by default. You must enable advanced capacity balancing to automatically balance FlexGroup capacity. Keep in mind that you cannot disable this feature once you enable it, but you can restore from a snapshot created before advanced capacity balancing was enabled.

Before you begin

- All nodes in the cluster must be running ONTAP 9.16.1 or later.

- You cannot revert to a release earlier than ONTAP 9.16.1 if advanced capacity balancing is enabled. If you need to revert, you must first restore from a snapshot created before advanced capacity balancing was enabled.
- If NFS copy offload has been enabled (`vserver nfs -vstorage enabled`) on an SVM, you cannot enable advanced capacity balancing on a FlexGroup volume. Similarly, if you have advanced capacity balancing enabled on any FlexGroup volume in an SVM, you cannot enable NFS copy offload.
- Advanced capacity balancing is not supported with FlexCache write-back.
- SnapMirror transfers are not supported with ONTAP versions earlier than ONTAP 9.16.1 when advanced capacity balancing is enabled on volumes in clusters running ONTAP 9.16.1 or later.
- Disable SMB Multichannel before enabling advanced capacity balancing. Using SMB Multichannel with advanced capacity rebalancing can cause high latency. For more information, see [ONTAP-400433: High read/write latency when using FlexGroup Rebalancing/GDD over SMB Multichannel enabled clients](#).

About this task

During creation of DP destination volumes using either of the granular data options (basic or advanced), the destination displays the setting as "disabled" until the SnapMirror transfer completes. After the transfer completes, the DP destination displays granular data as "enabled".

Enable advanced capacity balancing during FlexGroup creation

Steps

You can use System Manager or the ONTAP CLI to enable advanced capacity balancing when you create a new FlexGroup volume.

System Manager

1. Navigate to **Storage > Volumes** and click  .
2. In the **Add volume** window, enter the volume name and size. Then click **More Options**.
3. Under **Storage and optimization**, select **Distribute volume data across the cluster (FlexGroup)**.
4. Select **Advanced capacity balancing**.
5. Finish configuring the volume and click **Save**.

CLI

1. Create a volume with advanced capacity balancing enabled:

```
volume create -vserver <svm name> -volume <volume name> -size <volume size> -auto-provision-as flexgroup -junction-path /<path> -granular -data advanced
```

Example:

```
volume create -vserver vs0 -volume newvol -size 1TB -auto-provision -as flexgroup -junction-path /newvol -granular-data advanced
```

Enable advanced capacity balancing on existing FlexGroup volumes

Steps

You can use System Manager or the ONTAP CLI to enable advanced capacity balancing.

System Manager

1. Navigate to **Storage > Volumes**, click , and choose **Edit > Volume**.
2. In the **Edit volume** window, under **Storage and optimization**, select **Advanced capacity balancing**.
3. Click **Save**.

CLI

1. Modify an existing FlexGroup volume to enable advanced capacity balancing:

```
volume modify -vserver <svm name> -volume <volume name> -granular  
-data advanced
```

Example:

```
volume modify -vserver vs0 -volume newvol -granular-data advanced
```

Rebalance ONTAP FlexGroup volumes by moving files

Beginning with ONTAP 9.12.1, you can rebalance FlexGroup volumes by non-disruptively moving files from one constituent in a FlexGroup to another constituent.

FlexGroup rebalancing helps redistribute capacity when imbalances develop over time due to the addition of new files and file growth. After you manually start the rebalance operation, ONTAP selects the files and moves them automatically and non-disruptively.

 You should be aware that FlexGroup rebalancing degrades system performance when large numbers of files are moved as part of a single rebalancing event or over multiple rebalancing events because of the creation of multi-part inodes. Every file moved as part of a rebalancing event has 2 multi-part inodes associated with that file. The larger the number of files with multi-part inodes as a percentage of the total number of files in a FlexGroup, the greater the performance impact. Certain use cases, such as a FlexVol to FlexGroup conversion, can result in a significant amount of multi-part inode creation.

Rebalancing is available only when all nodes in the cluster are running ONTAP 9.12.1 or later releases. You must enable granular data functionality on any FlexGroup volume that runs the rebalancing operation. Once that functionality is enabled, you cannot revert to ONTAP 9.11.1 and earlier versions unless you delete this volume or restore from a snapshot that was created before the setting was enabled.

Beginning with ONTAP 9.14.1, ONTAP introduces an algorithm to non-disruptively and proactively move files in volumes that have granular data enabled without user interaction. The algorithm operates in very specific, targeted scenarios to alleviate performance bottlenecks. Scenarios where this algorithm might act include very heavy write load on a particular set of files on one node in the cluster or a continually growing file in a very hot

parent directory.

Beginning with ONTAP 9.16.1, you can also enable [advanced capacity balancing](#) to redistribute a large file's data between FlexGroup member volumes.

FlexGroup rebalancing considerations

You should be aware of how FlexGroup rebalancing works and how it interacts with other ONTAP features.

- FlexVol to FlexGroup conversion

It is recommended that you *not* use automatic FlexGroup rebalancing after a FlexVol to FlexGroup conversion. Instead, you can redistribute existing files by using the `volume rebalance file-move start` command, available in ONTAP 9.10.1 and later. This operation is non-disruptive by default (`-is-disruptive false`). If some busy files cannot be moved, you can rerun the command in disruptive mode (`-is-disruptive true`) during a planned maintenance window. Learn more about `volume rebalance file-move start` in the [ONTAP command reference](#).

Rebalancing with the automatic FlexGroup rebalancing feature can degrade performance when moving large numbers of files, like when you perform a FlexVol to FlexGroup conversion, and as much as 50 to 85% of the data on the FlexVol volume is moved to a new constituent.

- Minimum and maximum file size

File selection for automatic rebalancing is based on blocks saved. The minimum file size considered for rebalancing is 100 MB by default (can be configured as low as 20 MB using the `min-file-size` parameter shown below) and the maximum file size is 100 GB.

- Files in snapshots

You can configure FlexGroup rebalancing to only consider files to be moved which are not currently present in any snapshots. When rebalancing is started, a notification displays if a snapshot operation is scheduled anytime during a rebalancing operation.

Snapshots are restricted if a file is being moved and is undergoing framing at the destination. A snapshot restore operation is not allowed while file rebalancing is in progress.

Any snapshot created after the `granular-data` option is enabled cannot be replicated to a system running ONTAP 9.11.1 and earlier versions because ONTAP 9.11.1 and earlier versions do not support multi-part inodes.

- SnapMirror operations

FlexGroup rebalancing should take place between scheduled SnapMirror operations. A SnapMirror operation might fail if a file is being relocated before a SnapMirror operation begins if that file move does not complete within the 24-minute SnapMirror retry period. Any new file relocation that begins after a SnapMirror transfer has started will not fail.

- File-based compression storage efficiency

With file-based compression storage efficiency, the file is decompressed before it's moved to the destination, so the compression savings is lost. The compression savings is regained after a manually initiated background scanner runs on the FlexGroup volume after rebalancing. However, if any file is associated with a snapshot on any volume, the file will be ignored for compression.

- Deduplication

Moving deduplicated files can cause increased overall usage for the FlexGroup volume. During file rebalancing, only unique blocks are moved to the destination, freeing that capacity on the source. Shared blocks remain on the source and are copied to the destination. While this achieves the goal of reducing the used capacity on a nearly full source constituent, it can also lead to increased overall usage on the FlexGroup volume due to copies of shared blocks on the new destinations. This is also possible when files that are part of a snapshot are moved. The space savings is not fully recognized until the snapshot schedule recycles and there are no longer copies of the files in snapshots.

- FlexClone volumes

If file rebalancing is in progress when a FlexClone volume is created, the rebalancing will not be performed on the FlexClone volume. Rebalancing on the FlexClone volume should be performed after it is created.

- File move

When a file is moved during a FlexGroup rebalancing operation, the file size is reported as part of quota accounting on both the source and destination constituents. Once the move is completed, quota accounting returns to normal, and the file size is only reported on the new destination.

- Autonomous Ransomware Protection

Beginning with ONTAP 9.13.1, Autonomous Ransomware Protection is supported during disruptive and non-disruptive rebalance operations.

- Object store volumes

Volume capacity rebalancing is not supported on object store volumes, such as S3 buckets.

Enable FlexGroup rebalancing

Beginning with ONTAP 9.12.1, you can enable automatic non-disruptive FlexGroup volume rebalancing to redistribute files between FlexGroup constituents.

Beginning with ONTAP 9.13.1, you can schedule a single FlexGroup rebalancing operation to begin at a date and time in the future.

Before you begin

You must have enabled the `granular-data` option on the FlexGroup volume before enabling FlexGroup rebalancing. You can enable it by using one of these methods:

- When you create FlexGroup volume using the `volume create` command
- By modifying an existing FlexGroup volume to enable the setting using the `volume modify` command
- Setting it automatically when FlexGroup rebalancing is initiated using the `volume rebalance` command



If you are using ONTAP 9.16.1 or later and [FlexGroup advanced capacity balancing](#) is enabled using either the `granular-data` advanced option in the ONTAP CLI or using System Manager, FlexGroup rebalancing is also enabled.

Steps

You can manage FlexGroup rebalancing by using ONTAP System Manager or the ONTAP CLI.

System Manager

1. Navigate to **Storage > Volumes** and locate the FlexGroup volume to rebalance.
2. Select  to view the volume details.
3. Under **FlexGroup Balance Status** select **Rebalance**.



The **Rebalance** option is only available when the FlexGroup status is out of balance.

4. In the **Rebalance Volume** window, change the default settings as needed.
5. To schedule the rebalancing operation, select **Rebalance Later** and enter the date and time.

CLI

1. Start automatic rebalancing:

```
volume rebalance start -vserver <SVM name> -volume <volume name>
```

Optionally, you can specify the following options:

[-max-runtime] <time interval>] Maximum Runtime
[-max-threshold <percent>] Maximum Imbalance Threshold per Constituent
[-min-threshold <percent>] Minimum Imbalance Threshold per Constituent
[-max-file-moves <integer>] Maximum Concurrent File Moves per Constituent
[-min-file-size {<integer>[KB|MB|GB|TB|PB]}] Minimum file size
[-start-time <mm/dd/yyyy-00:00:00>] Schedule rebalance start date and time
[-exclude-snapshots {true|false}] Exclude files stuck in snapshots

Example:

```
volume rebalance start -vserver vs0 -volume fg1
```

Modify FlexGroup rebalance configurations

You can change a FlexGroup rebalancing configuration to update the imbalance threshold, number of concurrent files moves minimum file size, maximum runtime, and to include or exclude snapshots. Options to modify your FlexGroup rebalancing schedule are available beginning with ONTAP 9.13.1.

System Manager

1. Navigate to **Storage > Volumes** and locate the FlexGroup volume to rebalance.
2. Select  to view the volume details.
3. Under **FlexGroup Balance Status** select **Rebalance**.



The **Rebalance** option is only available when the FlexGroup status is out of balance.

4. In the **Rebalance Volume** window, change the default settings as needed.

CLI

1. Modify automatic rebalancing:

```
volume rebalance modify -vserver <SVM name> -volume <volume name>
```

You can specify one or more of the following options:

[-max-runtime] <time interval> Maximum Runtime
[-max-threshold <percent>] Maximum Imbalance Threshold per Constituent
[-min-threshold <percent>] Minimum Imbalance Threshold per Constituent
[-max-file-moves <integer>] Maximum Concurrent File Moves per Constituent
[-min-file-size {<integer>[KB|MB|GB|TB|PB]}] Minimum file size
[-start-time <mm/dd/yyyy-00:00:00>] Schedule rebalance start date and time
[-exclude-snapshots {true|false}] Exclude files stuck in snapshots

Stop FlexGroup rebalance

After FlexGroup rebalancing is enabled or scheduled, you can stop it at any time.

System Manager

1. Navigate to **Storage > Volumes** and locate the FlexGroup volume.
2. Select  to view the volume details.
3. Select **Stop Rebalance**.

CLI

1. Stop FlexGroup rebalancing:

```
volume rebalance stop -vserver <SVM name> -volume <volume name>
```

View FlexGroup rebalance status

You can display the status about a FlexGroup rebalance operation, the FlexGroup rebalance configuration, the rebalance operation time, and the rebalance instance details.

System Manager

1. Navigate to **Storage > Volumes** and locate the FlexGroup volume.
2. Select  to view the FlexGroup details.
3. **FlexGroup Balance Status** is displayed near the bottom of the details pane.
4. To view information about the last rebalance operation, select **Last Volume Rebalance Status**.

CLI

1. View the status of a FlexGroup rebalance operation:

```
volume rebalance show
```

Example of rebalance state:

```
> volume rebalance show
Vserver: vs0

Target
Imbalance
Volume      State          Total      Used      Used
Size        %
-----
-----
fg1        idle          4GB      115.3MB      -
8KB        0%
```

Example of rebalance configuration details:

```
> volume rebalance show -config
Vserver: vs0

Max          Threshold          Max
Min          Exclude
Volume      Runtime          Min      Max      File Moves
File Size   Snapshot
-----
-----
fg1          6h0m0s          5%      20%      25
4KB          true
```

Example of rebalance time details:

```

> volume rebalance show -time
Vserver: vs0
Volume           Start Time           Runtime
Max Runtime
-----
-----
fg1             Wed Jul 20 16:06:11 2022 0h1m16s
6h0m0s

```

Example of rebalance instance details:

```

> volume rebalance show -instance
Vserver Name: vs0
Volume Name: fg1
Is Constituent: false
Rebalance State: idle
Rebalance Notice Messages: -
Total Size: 4GB
AFS Used Size: 115.3MB
Constituent Target Used Size: -
Imbalance Size: 8KB
Imbalance Percentage: 0%
Moved Data Size: -
Maximum Constituent Imbalance Percentage: 1%
Rebalance Start Time: Wed Jul 20 16:06:11 2022
Rebalance Stop Time: -
Rebalance Runtime: 0h1m32s
Rebalance Maximum Runtime: 6h0m0s
Maximum Imbalance Threshold per Constituent: 20%
Minimum Imbalance Threshold per Constituent: 5%
Maximum Concurrent File Moves per Constituent: 25
Minimum File Size: 4KB
Exclude Files Stuck in snapshots: true

```

Data protection for FlexGroup volumes

Data protection for ONTAP FlexGroup volumes workflow summary

You can create SnapMirror disaster recovery (DR) relationships for FlexGroup volumes. You can also backup and restore FlexGroup volumes by using SnapVault technology, and you can create a unified data protection relationship that uses the same destination for backup and DR.

About this task

The SnapMirror relationship type is always XDP for FlexGroup volumes. The type of data protection that is provided by a SnapMirror relationship is determined by the replication policy that you use. You can use either the default policy or a custom policy of the required type for the replication relationship that you want to create.

1

Peer the clusters and SVMs

If the clusters and SVMs are not already peered, create the [cluster peers](#) and the [SVM peers](#).

2

Create a job schedule

You must [create a job schedule](#) to determine when SnapMirror updates will take place.

3

Depending on the type of data protection, follow one of these paths:

- **If SnapMirror DR:**

[Create a SnapMirror relationship](#). When you create the relationship, you can select the default policy `MirrorAllSnapshots` or a custom policy of type `async-mirror`.

- **If SnapMirror vault:**

[Create a SnapMirror vault relationship](#). When you create the relationship, you can select the default policy `XDPDefault` or a custom policy of type `vault`.

- **If unified data protection:**

[Create a unified data protection relationship](#). When you create the relationship, you can select the default policy `MirrorAndVault` or a custom policy of type `mirror-vault`.

Create SnapMirror relationships for ONTAP FlexGroup volumes

You can create a SnapMirror relationship between the source FlexGroup volume and the destination FlexGroup volume on a peered SVM for replicating data for disaster recovery. You can use the mirror copies of the FlexGroup volume to recover data when a disaster occurs.

Before you begin

You must have created the cluster peering relationship and SVM peering relationship.

[Cluster and SVM peering](#)

About this task

- Beginning with ONTAP 9.9.1, you can use the ONTAP CLI to create SnapMirror cascade and fanout relationships for FlexGroup volumes.

For details, see [Considerations for creating SnapMirror cascade and fanout relationships for FlexGroup volumes](#).

- You can create both intercluster SnapMirror relationships and intracluster SnapMirror relationships for FlexGroup volumes.

- Beginning with ONTAP 9.3, you can expand FlexGroup volumes that are in a SnapMirror relationship.

If you are using a version of ONTAP earlier than ONTAP 9.3, do not expand FlexGroup volumes after a SnapMirror relationship is established; however, you can increase the capacity of FlexGroup volumes after establishing a SnapMirror relationship. If you expand the source FlexGroup volume after breaking the SnapMirror relationship in releases earlier than ONTAP 9.3, you must perform a baseline transfer to the destination FlexGroup volume.

Steps

- Create a destination FlexGroup volume of type DP that has the same number of constituents as that of the source FlexGroup volume:

- From the source cluster, determine the number of constituents in the source FlexGroup volume:
`volume show -volume volume_name* -is-constituent true`

```
cluster1::> volume show -volume srcFG* -is-constituent true
Vserver    Volume          Aggregate      State       Type       Size
Available  Used%
-----  -----
vss        srcFG           -             online      RW        400TB
172.86GB  56%
vss        srcFG__0001      Aggr_cmode   online      RW        25GB
10.86TB   56%
vss        srcFG__0002      aggr1        online      RW        25TB
10.86TB   56%
vss        srcFG__0003      Aggr_cmode   online      RW        25TB
10.72TB   57%
vss        srcFG__0004      aggr1        online      RW        25TB
10.73TB   57%
vss        srcFG__0005      Aggr_cmode   online      RW        25TB
10.67TB   57%
vss        srcFG__0006      aggr1        online      RW        25TB
10.64TB   57%
vss        srcFG__0007      Aggr_cmode   online      RW        25TB
10.63TB   57%
...
...
```

- From the destination cluster, create a destination FlexGroup volume of type DP with the same number of constituents as that of the source FlexGroup volume.

```
cluster2::> volume create -vserver vsd -aggr-list aggr1,aggr2 -aggr
-list-multiplier 8 -size 400TB -type DP dstFG

Warning: The FlexGroup volume "dstFG" will be created with the
following number of constituents of size 25TB: 16.

Do you want to continue? {y|n}: y
[Job 766] Job succeeded: Successful
```

c. From the destination cluster, verify the number of constituents in the destination FlexGroup volume:

```
volume show -volume volume_name* -is-constituent true
```

Vserver	Volume	Aggregate	State	Type	Size
Available	Used%				
vsd	dstFG	-	online	DP	400TB
172.86GB	56%				
vsd	dstFG__0001	Aggr_cmode	online	DP	25GB
10.86TB	56%				
vsd	dstFG__0002	aggr1	online	DP	25TB
10.86TB	56%				
vsd	dstFG__0003	Aggr_cmode	online	DP	25TB
10.72TB	57%				
vsd	dstFG__0004	aggr1	online	DP	25TB
10.73TB	57%				
vsd	dstFG__0005	Aggr_cmode	online	DP	25TB
10.67TB	57%				
vsd	dstFG__0006	aggr1	online	DP	25TB
10.64TB	57%				
vsd	dstFG__0007	Aggr_cmode	online	DP	25TB
10.63TB	57%				
...					

2. Create a job schedule: `job schedule cron create -name job_name -month month -dayofweek day_of_week -day day_of_month -hour hour -minute minute`

For the `-month`, `-dayofweek`, and `-hour` options, you can specify `all` to run the job every month, every day of the week, and every hour, respectively.

The following example creates a job schedule named `my_weekly` that runs on Saturdays at 3:00 a.m.:

```
cluster1::> job schedule cron create -name my_weekly -dayofweek
"Saturday" -hour 3 -minute 0
```

3. Create a custom policy of type `async-mirror` for the SnapMirror relationship: `snapmirror policy create -vserver SVM -policy snapmirror_policy -type async-mirror`

If you do not create a custom policy, you should specify the `MirrorAllSnapshots` policy for SnapMirror relationships.

4. From the destination cluster, create a SnapMirror relationship between the source FlexGroup volume and the destination FlexGroup volume: `snapmirror create -source-path src_svm:src_flexgroup -destination-path dest_svm:dest_flexgroup -type XDP -policy snapmirror_policy -schedule sched_name`

SnapMirror relationships for FlexGroup volumes must be of type `XDP`.

If you specify a throttle value for the SnapMirror relationship for the FlexGroup volume, each constituent uses the same throttle value. The throttle value is not divided among the constituents.



You cannot use SnapMirror labels of snapshots for FlexGroup volumes.

In ONTAP 9.4 and earlier, if the policy is not specified with the `snapmirror create` command, the `MirrorAllSnapshots` policy is used by default. In ONTAP 9.5, if the policy is not specified with the `snapmirror create` command, the `MirrorAndVault` policy is used by default.

```
cluster2::> snapmirror create -source-path vss:srcFG -destination-path
vsd:dstFG -type XDP -policy MirrorAllSnapshots -schedule hourly
Operation succeeded: snapmirror create for the relationship with
destination "vsd:dstFG".
```

5. From the destination cluster, initialize the SnapMirror relationship by performing a baseline transfer: `snapmirror initialize -destination-path dest_svm:dest_flexgroup`

After the baseline transfer is completed, the destination FlexGroup volume is updated periodically based on the schedule of the SnapMirror relationship.

```
cluster2::> snapmirror initialize -destination-path vsd:dstFG
Operation is queued: snapmirror initialize of destination "vsd:dstFG".
```



If you have created any SnapMirror relationship between FlexGroup volumes with the source cluster running ONTAP 9.3 and the destination cluster running ONTAP 9.2 or earlier, and if you create any qtrees in the source FlexGroup volume, the SnapMirror updates fail. To recover from this situation, you must delete all of the non-default qtrees in the FlexGroup volume, disable the qtree functionality on the FlexGroup volume, and then delete all of the snapshots that are enabled with the qtree functionality.

After you finish

You should set up the destination SVM for data access by setting up required configurations such as LIFs and export policies.

Related information

- [snapmirror create](#)
- [snapmirror initialize](#)
- [snapmirror policy create](#)
- [snapmirror update](#)

Create SnapVault relationships for ONTAP FlexGroup volumes

You can configure a SnapVault relationship and assign a SnapVault policy to the relationship to create a SnapVault backup.

Before you begin

You must be aware of the considerations for creating a SnapVault relationship for FlexGroup volumes.

Steps

1. Create a destination FlexGroup volume of type DP that has the same number of constituents as that of the source FlexGroup volume:

- a. From the source cluster, determine the number of constituents in the source FlexGroup volume:
`volume show -volume volume_name* -is-constituent true`

```
cluster1::> volume show -volume src* -is-constituent true
Vserver    Volume        Aggregate    State      Type       Size
Available  Used%
-----  -----  -----  -----  -----  -----
-----  -----
vss        src          -           online    RW        400TB
172.86GB  56%
vss        src_0001     Aggr_cmode  online    RW        25GB
10.86TB   56%
vss        src_0002     aggr1      online    RW        25TB
10.86TB   56%
vss        src_0003     Aggr_cmode  online    RW        25TB
10.72TB   57%
vss        src_0004     aggr1      online    RW        25TB
10.73TB   57%
vss        src_0005     Aggr_cmode  online    RW        25TB
10.67TB   57%
vss        src_0006     aggr1      online    RW        25TB
10.64TB   57%
vss        src_0007     Aggr_cmode  online    RW        25TB
10.63TB   57%
...
...
```

- b. From the destination cluster, create a destination FlexGroup volume of type DP with the same number of constituents as that of the source FlexGroup volume.

```
cluster2::> volume create -vserver vsd -aggr-list aggr1,aggr2 -aggr
-list-multiplier 8 -size 400TB -type DP dst

Warning: The FlexGroup volume "dst" will be created with the
following number of constituents of size 25TB: 16.
Do you want to continue? {y|n}: y
[Job 766] Job succeeded: Successful
```

c. From the destination cluster, verify the number of constituents in the destination FlexGroup volume:
`volume show -volume volume_name* -is-constituent true`

Vserver	Volume	Aggregate	State	Type	Size
Available	Used%				
vsd	dst	-	online	RW	400TB
172.86GB	56%				
vsd	dst_0001	Aggr_cmode	online	RW	25GB
10.86TB	56%				
vsd	dst_0002	aggr1	online	RW	25TB
10.86TB	56%				
vsd	dst_0003	Aggr_cmode	online	RW	25TB
10.72TB	57%				
vsd	dst_0004	aggr1	online	RW	25TB
10.73TB	57%				
vsd	dst_0005	Aggr_cmode	online	RW	25TB
10.67TB	57%				
vsd	dst_0006	aggr1	online	RW	25TB
10.64TB	57%				
vsd	dst_0007	Aggr_cmode	online	RW	25TB
10.63TB	57%				
...					

2. Create a job schedule: `job schedule cron create -name job_name -month month -dayofweek day_of_week -day day_of_month -hour hour -minute minute`

For `-month`, `-dayofweek`, and `-hour`, you can specify `all` to run the job every month, day of the week, and hour, respectively.

The following example creates a job schedule named `my_weekly` that runs on Saturdays at 3:00 a.m.:

```
cluster1::> job schedule cron create -name my_weekly -dayofweek
"Saturday" -hour 3 -minute 0
```

3. Create a SnapVault policy, and then define a rule for the SnapVault policy:

- a. Create a custom policy of type `vault` for the SnapVault relationship: `snapmirror policy create -vserver svm_name -policy policy_name -type vault`
- b. Define a rule for the SnapVault policy that determines which snapshots are transferred during initialization and update operations: `snapmirror policy add-rule -vserver svm_name -policy policy_for_rule - snapmirror-label snapmirror-label -keep retention_count -schedule schedule`

If you do not create a custom policy, you should specify the `XDPDefault` policy for SnapVault relationships.

4. Create a SnapVault relationship: `snapmirror create -source-path src_svm:src_flexgroup -destination-path dest_svm:dest_flexgroup -type XDP -schedule schedule_name -policy XDPDefault`

In ONTAP 9.4 and earlier, if the policy is not specified with the `snapmirror create` command, the `MirrorAllSnapshots` policy is used by default. In ONTAP 9.5, if the policy is not specified with the `snapmirror create` command, the `MirrorAndVault` policy is used by default.

```
cluster2::> snapmirror create -source-path vss:srcFG -destination-path vsd:dstFG -type XDP -schedule Daily -policy XDPDefault
```

Learn more about `snapmirror create` in the [ONTAP command reference](#).

5. From the destination cluster, initialize the SnapVault relationship by performing a baseline transfer: `snapmirror initialize -destination-path dest_svm:dest_flexgroup`

```
cluster2::> snapmirror initialize -destination-path vsd:dst
Operation is queued: snapmirror initialize of destination "vsd:dst".
```

Related information

- [snapmirror create](#)
- [snapmirror initialize](#)
- [snapmirror policy add-rule](#)
- [snapmirror policy create](#)

Create unified data protection relationships for ONTAP FlexGroup volumes

Beginning with ONTAP 9.3, you can create and configure SnapMirror unified data protection relationships to configure disaster recovery and archiving on the same destination volume.

Before you begin

You must be aware of the considerations for creating unified data protection relationships for FlexGroup volumes.

Considerations for creating a SnapVault backup relationship and a unified data protection relationship for FlexGroup volumes

Steps

1. Create a destination FlexGroup volume of type DP that has the same number of constituents as that of the source FlexGroup volume:

- a. From the source cluster, determine the number of constituents in the source FlexGroup volume:
`volume show -volume volume_name* -is-constituent true`

```
cluster1::> volume show -volume srcFG* -is-constituent true
Vserver      Volume          Aggregate      State      Type      Size
Available    Used%
-----  -----  -----  -----  -----  -----
-----  -----
vss          srcFG           -            online    RW       400TB
172.86GB    56%
vss          srcFG__0001      Aggr_cmode   online    RW       25GB
10.86TB     56%
vss          srcFG__0002      aggr1        online    RW       25TB
10.86TB     56%
vss          srcFG__0003      Aggr_cmode   online    RW       25TB
10.72TB     57%
vss          srcFG__0004      aggr1        online    RW       25TB
10.73TB     57%
vss          srcFG__0005      Aggr_cmode   online    RW       25TB
10.67TB     57%
vss          srcFG__0006      aggr1        online    RW       25TB
10.64TB     57%
vss          srcFG__0007      Aggr_cmode   online    RW       25TB
10.63TB     57%
...
...
```

- b. From the destination cluster, create a destination FlexGroup volume of type DP with the same number of constituents as that of the source FlexGroup volume.

```
cluster2::> volume create -vserver vsd -aggr-list aggr1,aggr2 -aggr
-list-multiplier 8 -size 400TB -type DP dstFG
```

Warning: The FlexGroup volume "dstFG" will be created with the following number of constituents of size 25TB: 16.

Do you want to continue? {y|n}: y
[Job 766] Job succeeded: Successful

- c. From the destination cluster, verify the number of constituents in the destination FlexGroup volume:
`volume show -volume volume_name* -is-constituent true`

```
cluster2::> volume show -volume dstFG* -is-constituent true
Vserver      Volume      Aggregate      State      Type      Size
Available    Used%
-----  -----
vsd          dstFG       -            online    RW       400TB
172.86GB    56%
vsd          dstFG_0001  Aggr_cmode   online    RW       25GB
10.86TB     56%
vsd          dstFG_0002  aggr1       online    RW       25TB
10.86TB     56%
vsd          dstFG_0003  Aggr_cmode   online    RW       25TB
10.72TB     57%
vsd          dstFG_0004  aggr1       online    RW       25TB
10.73TB     57%
vsd          dstFG_0005  Aggr_cmode   online    RW       25TB
10.67TB     57%
vsd          dstFG_0006  aggr1       online    RW       25TB
10.64TB     57%
vsd          dstFG_0007  Aggr_cmode   online    RW       25TB
10.63TB     57%
...
...
```

2. Create a job schedule: `job schedule cron create -name job_name -month month -dayofweek day_of_week -day day_of_month -hour hour -minute minute`

For the `-month`, `-dayofweek`, and `-hour` options, you can specify `all` to run the job every month, every day of the week, and every hour, respectively.

The following example creates a job schedule named `my_weekly` that runs on Saturdays at 3:00 a.m.:

```
cluster1::> job schedule cron create -name my_weekly -dayofweek
"Saturday" -hour 3 -minute 0
```

3. Create a custom policy of type `mirror-vault`, and then define a rule for the mirror and vault policy:

a. Create a custom policy of type `mirror-vault` for the unified data protection relationship:

```
snapmirror policy create -vserver svm_name -policy policy_name -type mirror-
vault
```

b. Define a rule for the mirror and vault policy that determines which snapshots are transferred during initialization and update operations: `snapmirror policy add-rule -vserver svm_name -policy policy_for_rule - snapmirror-label snapmirror-label -keep retention_count -schedule schedule`

If you do not specify a custom policy, the `MirrorAndVault` policy is used for unified data protection relationships.

4. Create a unified data protection relationship:

```
snapmirror create -source-path src_svm:src_flexgroup -destination-path dest_svm:dest_flexgroup -type XDP -schedule schedule_name -policy MirrorAndVault
```

In ONTAP 9.4 and earlier, if the policy is not specified with the `snapmirror create` command, the `MirrorAllSnapshots` policy is used by default. In ONTAP 9.5, if the policy is not specified with the `snapmirror create` command, the `MirrorAndVault` policy is used by default.

```
cluster2::> snapmirror create -source-path vss:srcFG -destination-path vsd:dstFG -type XDP -schedule Daily -policy MirrorAndVault
```

Learn more about `snapmirror create` in the [ONTAP command reference](#).

5. From the destination cluster, initialize the unified data protection relationship by performing a baseline transfer: `snapmirror initialize -destination-path dest_svm:dest_flexgroup`

```
cluster2::> snapmirror initialize -destination-path vsd:dstFG
Operation is queued: snapmirror initialize of destination "vsd:dstFG".
```

Related information

- [snapmirror create](#)
- [snapmirror initialize](#)
- [snapmirror policy add-rule](#)
- [snapmirror policy create](#)

Create SVM disaster recovery relationships for ONTAP FlexGroup volumes

Beginning with ONTAP 9.9.1, you can create SVM disaster recovery (SVM DR) relationships using FlexGroup volumes. An SVM DR relationship provides redundancy and the ability to recover FlexGroup volumes in the event of a disaster by synchronizing and replicating the SVM configuration and its data. A SnapMirror license is required for SVM DR.

Before you begin

You *cannot* create a FlexGroup SVM DR relationship with the following applies.

- A FlexClone FlexGroup configuration exists
- The FlexGroup volume is part of a cascading relationship
- The FlexGroup volume is part of a fanout relationship, and your cluster is running an ONTAP version earlier than ONTAP 9.12.1. (Beginning with ONTAP 9.13.1, fanout relationships are supported.)

About this task

- All nodes in both clusters must be running the same ONTAP version as the node on which SVM DR support was added (ONTAP 9.9.1 or later).
- The SVM DR relationship between the primary and secondary sites should be healthy and should have

enough space on both the primary and secondary SVMs to support the FlexGroup volumes.

- Beginning with ONTAP 9.12.1, FabricPool, FlexGroup, and SVM DR can work in conjunction. In releases earlier than ONTAP 9.12.1, any two of these features worked together, but not all three in conjunction.
- When you create a FlexGroup SVM DR relationship in which the FlexGroup volume is part of a fanout relationship, you should be aware of the following requirements:
 - The source and destination cluster must be running ONTAP 9.13.1 or later.
 - SVM DR with FlexGroup volumes supports SnapMirror fanout relationships to eight sites.

For information about creating an SVM DR relationship, see [Manage SnapMirror SVM replication](#).

Steps

1. Create an SVM DR relationship, or use an existing relationship.

[Replicate an entire SVM configuration](#)

2. Create a FlexGroup volume on the primary site with the required number of constituents.

[Creating a FlexGroup volume.](#)

Wait until FlexGroup and all of its constituents are created before proceeding.

3. To replicate the FlexGroup volume, update the SVM at the secondary site: `snapmirror update -destination-path destination_svm_name: -source-path source_svm_name:`

You can also check if a scheduled SnapMirror update already exists by entering `snapmirror show -fields schedule`

4. From the secondary site, verify that the SnapMirror relationship is healthy: `snapmirror show`

```
cluster2::> snapmirror show

Progress
Source          Destination Mirror  Relationship  Total
Last
Path           Type   Path        State   Status      Progress  Healthy
Updated
-----  -----  -----  -----  -----  -----  -----
-----  -----
vs1:          XDP    vs1_dst:    Snapmirrored
                           Idle      -        true      -
```

5. From the secondary site, verify that the new FlexGroup volume and its constituents exist: `snapmirror show -expand`

```

cluster2::> snapmirror show -expand

Progress
Source          Destination Mirror  Relationship  Total
Last
Path           Type   Path        State   Status      Progress  Healthy
Updated

-----
-----
```

Source	Destination	Mirror	Relationship	Total
vs1:	XDP	vs1_dst:	Snapmirrored	
			Idle	-
vs1:fg_src	XDP	vs1_dst:fg_src	Snapmirrored	
			Idle	-
vs1:fg_src_0001	XDP	vs1_dst:fg_src_0001	Snapmirrored	
			Idle	-
vs1:fg_src_0002	XDP	vs1_dst:fg_src_0002	Snapmirrored	
			Idle	-
vs1:fg_src_0003	XDP	vs1_dst:fg_src_0003	Snapmirrored	
			Idle	-
vs1:fg_src_0004	XDP	vs1_dst:fg_src_0004	Snapmirrored	
			Idle	-
6 entries were displayed.				

Related information

- [snapmirror show](#)
- [snapmirror update](#)

Transition ONTAP FlexGroup SnapMirror relationships to SVM DR

You can create a FlexGroup SVM DR relationship by transitioning an existing FlexGroup volume SnapMirror relationship.

Before you begin

- The FlexGroup volume SnapMirror relationship is in a healthy state.
- The source and destination FlexGroup volumes have the same name.

Steps

1. From the SnapMirror destination, resynchronize the FlexGroup level SnapMirror relationship: `snapmirror resync`
2. Create the FlexGroup SVM DR SnapMirror relationship. Use the same SnapMirror policy which is configured on the FlexGroup volume SnapMirror relationships: `snapmirror create -destination-path dest_svm: -source-path src_svm: -identity-preserve true -policy MirrorAllSnapshots`



You must use the `-identity-preserve true` option of the `snapmirror create` command when you create your replication relationship.

Learn more about `snapmirror create` in the [ONTAP command reference](#).

3. Verify the relationship is broken off: `snapmirror show -destination-path dest_svm: -source-path src_svm:`

```
snapmirror show -destination-path fg_vs_renamed: -source-path fg_vs:  
  
Progress  
Source          Destination Mirror  Relationship  Total  
Last  
Path           Type   Path       State   Status      Progress  Healthy  
Updated  
-----  
-----  
fg_vs:         XDP    fg_vs1_renamed:      Broken-off  
                           Idle           -           true     -
```

4. Stop the destination SVM: `vserver stop -vserver vs_name`

```
vserver stop -vserver fg_vs_renamed  
[Job 245] Job is queued: Vserver Stop fg_vs_renamed.  
[Job 245] Done
```

5. Resynchronize the SVM SnapMirror relationship: `snapmirror resync -destination-path dest_svm: -source-path src_svm:`

```
snapmirror resync -destination-path fg_vs_renamed: -source-path fg_vs:  
Warning: This Vserver has volumes which are the destination of FlexVol  
or FlexGroup SnapMirror relationships. A resync on the Vserver  
SnapMirror relationship will cause disruptions in data access
```

6. Verify that the SVM DR level SnapMirror relationship reaches a healthy idle state: `snapmirror show`

-expand

7. Verify that the FlexGroup SnapMirror relationship is in a healthy state: `snapmirror show`

Related information

- [snapmirror create](#)
- [snapmirror resync](#)
- [snapmirror show](#)

Convert ONTAP FlexVol volumes to FlexGroup volumes within an SVM-DR relationship

Beginning with ONTAP 9.10.1, you can convert a FlexVol volume to a FlexGroup volume on an SVM-DR source.

Before you begin

- The FlexVol volume that is being converted must be online.
- The operations and configurations on the FlexVol volume must be compatible with the conversion process.

An error message is generated if the FlexVol volume has any incompatibility, and the volume conversion is cancelled. You can take corrective actions and retry the conversion.

For more details, see [Considerations for converting FlexVol volumes to FlexGroup volumes](#)

Steps

1. Login using advance privilege mode: `set -privilege advanced`
2. From the destination, update the SVM-DR relationship:

```
snapmirror update -destination-path <destination_svm_name>: -source-path  
<source_svm_name>:
```



You must enter a colon (:) after the SVM name in the `-destination-path` option.

3. Ensure that the SVM-DR relationship is in a SnapMirrored state and is not broken-off:

```
snapmirror show
```

4. From the destination SVM, verify that the FlexVol volume is ready for conversion:

```
volume conversion start -vserver <svm_name> -volume <vol_name> -check  
-only true
```

If this command generates any errors other than "This is a destination SVMDR volume," you can take the appropriate corrective action, run the command again, and continue the conversion.

5. From the destination, disable transfers on the SVM-DR relationship:

```
snapmirror quiesce -destination-path <dest_svm>:
```



You must enter a colon (:) after the SVM name in the -destination-path option.

6. From the source cluster, start the conversion:

```
volume conversion start -vserver <svm_name> -volume <vol_name>
```

7. Verify that the conversion is successful:

```
volume show <vol_name> -fields volume-style-extended,state
```

```
cluster-1::*> volume show my_volume -fields volume-style-extended,state
```

vserver	volume	state	volume-style-extended
-----	-----	-----	-----
vs0	my_volume	online	flexgroup

8. From the destination cluster, resume transfers for the relationship:

```
snapmirror resume -destination-path <dest_svm>:
```



You must enter a colon (:) after the SVM name in the -destination-path option.

9. From the destination cluster, perform an update to propagate the conversion to the destination:

```
snapmirror update -destination-path <dest_svm>:
```



You must enter a colon (:) after the SVM name in the -destination-path option.

10. Ensure that the SVM-DR relationship is in a SnapMirrored state and is not broken off:

```
snapmirror show
```

11. Ensure the conversion occurred on the destination:

```
volume show <vol_name> -fields volume-style-extended,state
```

```
cluster-2::*> volume show my_volume -fields volume-style-extended,state

vserver    volume      state    volume-style-extended
-----  -----  -----
vs0_dst    my_volume  online   flexgroup
```

Related information

- [snapmirror resume](#)
- [snapmirror quiesce](#)
- [snapmirror show](#)
- [snapmirror update](#)

Considerations for creating SnapMirror cascade and fanout relationships for ONTAP FlexGroup volumes

There are support considerations and limitations you should keep in mind when creating SnapMirror cascade and fanout relationships for FlexGroup volumes.

Considerations for creating cascading relationships

- Each relationship can be either an inter cluster or intra cluster relationship.
- All asynchronous policy types, including async-mirror, mirror-vault, and vault, are supported for both relationships.
- Only "MirrorAllSnapshots," not "MirrorLatest" async-mirror policies are supported.
- Long-term retention snapshots are not supported.

Learn more about [long-term retention snapshots](#).

- Concurrent updates of cascaded XDP relationships is supported.
- Supports removing A to B and B to C and resync A to C or resync C to A.
- A and B FlexGroup volumes also support fanout when all nodes are running ONTAP 9.9.1 or later.
- Restore operations from B or C FlexGroup volumes are supported.
- Transfers on FlexGroup relationships are not support while the destination is the source of a restore relationship.
- The destination of a FlexGroup restore cannot be the destination of any other FlexGroup relationship.
- FlexGroup file restore operations have the same restrictions as regular FlexGroup restore operations.
- All nodes in the cluster where the B and C FlexGroup volumes reside must be running ONTAP 9.9.1 or later.
- All expand and auto expand functionality is supported.
- In a cascade configuration such as A to B to C, if A to B and B to C have different numbers of constituent SnapMirror relationships, then an abort operation from the source is not supported for the B to C SnapMirror relationship.
- System Manager does not support cascading relationships regardless of the ONTAP version.

- When converting an A to B to C set of FlexVol relationship to a FlexGroup relationship, you must convert the B to C hop first.
- All FlexGroup cascade configurations for relationships with policy types supported by REST are also supported by REST APIs in cascading FlexGroup configurations.
- As with FlexVol relationships, FlexGroup cascading is not supported by the `snapmirror protect` command.

Considerations for creating fanout relationships

- Two or more FlexGroup fanout relationships are supported; for example, A to B, A to C, with a maximum of 8 fanout legs.
- Each relationship can be either intercluster or intracluster.
- Concurrent updates are supported for the two relationships.
- All expand and auto expand functionality is supported.
- If the fanout legs of the relationship have different numbers of constituent SnapMirror relationships, then an abort operation from the source is not supported for the A to B and A to C relationships.
- All nodes in the cluster where the source and destination FlexGroup volumes reside must be running ONTAP 9.9.1 or later.
- All asynchronous policy types currently supported for FlexGroup SnapMirror are supported in fanout relationships.
- You can perform restore operations from B to C FlexGroup volumes.
- All fanout configurations with policy types supported by REST are also supported for REST APIs in FlexGroup fanout configurations.

Related information

- [snapmirror protect](#)

Considerations for creating SnapVault backup relationships and unified data protection relationships for ONTAP FlexGroup volumes

You must be aware of the considerations for creating a SnapVault backup relationship and unified data protection relationship for FlexGroup volumes.

- You can resynchronize a SnapVault backup relationship and a unified data protection relationship by using the `-preserve` option that enables you to preserve snapshots on the destination volume that are newer than the latest common snapshot.
- Long-term retention is not supported with FlexGroup volumes.

Long-term retention enables creating snapshots directly on the destination volume without requiring to store the snapshots on the source volume.

- The `snapshot` command `expiry-time` option is not supported for FlexGroup volumes.
- Storage efficiency cannot be configured on the destination FlexGroup volume of a SnapVault backup relationship and unified data protection relationship.
- You cannot rename snapshots of a SnapVault backup relationship and unified data protection relationship for FlexGroup volumes.
- A FlexGroup volume can be the source volume of only one backup relationship or restore relationship.

A FlexGroup volume cannot be the source of two SnapVault relationships, two restore relationships, or a SnapVault backup relationship and a restore relationship.

- If you delete a snapshot on the source FlexGroup volume and re-create a snapshot with the same name, the next update transfer to the destination FlexGroup volume fails if the destination volume has a snapshot of the same name.

This is because snapshots cannot be renamed for FlexGroup volumes.

Monitor SnapMirror data transfers for ONTAP FlexGroup volumes

You should periodically monitor the status of the FlexGroup volume SnapMirror relationships to verify that the destination FlexGroup volume is updated periodically as per the specified schedule.

About this task

You must perform this task from the destination cluster.

Steps

1. View the SnapMirror relationship status of all FlexGroup volume relationships: `snapmirror show -relationship-group-type flexgroup`

```
cluster2::> snapmirror show -relationship-group-type flexgroup

Progress
Source          Destination Mirror  Relationship  Total
Last
Path           Type   Path        State   Status      Progress  Healthy
Updated
-----
-----
vss:s          XDP   vsd:d      Snapmirrored
                  Idle
vss:s2         XDP   vsd:d2     Uninitialized
                  Idle
2 entries were displayed.
```

Related information

- [snapmirror show](#)


Manage data protection operations for FlexGroup volumes

Disaster recovery for FlexGroup volumes

Disaster recovery workflow for ONTAP FlexGroup volumes

When a disaster strikes on the source FlexGroup volume, you should activate the

destination FlexGroup volume and redirect client access. Depending on whether the source FlexGroup volume can be recovered, you should either reactivate the source FlexGroup volume or reverse the SnapMirror relationship.



About this task

Client access to the destination FlexGroup volume is blocked for a brief period when some SnapMirror operations, such as SnapMirror break and resynchronization, are running. If the SnapMirror operation fails, it is possible that some of the constituents remain in this state and access to the FlexGroup volume is denied. In such cases, you must retry the SnapMirror operation.

Activate the destination ONTAP FlexGroup volume

When the source FlexGroup volume is unable to serve data due to events such as data

corruption, accidental deletion or an offline state, you must activate the destination FlexGroup volume to provide data access until you recover the data on the source FlexGroup volume. Activation involves stopping future SnapMirror data transfers and breaking the SnapMirror relationship.

About this task

You must perform this task from the destination cluster.

Steps

1. Disable future transfers for the FlexGroup volume SnapMirror relationship: `snapmirror quiesce dest_svm:dest_flexgroup`

```
cluster2::> snapmirror quiesce -destination-path vsd:dst
```

2. Break the FlexGroup volume SnapMirror relationship: `snapmirror break dest_svm:dest_flexgroup`

```
cluster2::> snapmirror break -destination-path vsd:dst
```

3. View the status of the SnapMirror relationship: `snapmirror show -expand`

```

cluster2::> snapmirror show -expand

Progress
Source          Destination Mirror  Relationship  Total
Last
Path           Type   Path        State   Status
Updated

-----
vss:s          XDP   vsd:dst      Broken-off
                           Idle      -       true    -
vss:s_0001     XDP   vsd:dst_0001 Broken-off
                           Idle      -       true    -
vss:s_0002     XDP   vsd:dst_0002 Broken-off
                           Idle      -       true    -
vss:s_0003     XDP   vsd:dst_0003 Broken-off
                           Idle      -       true    -
vss:s_0004     XDP   vsd:dst_0004 Broken-off
                           Idle      -       true    -
vss:s_0005     XDP   vsd:dst_0005 Broken-off
                           Idle      -       true    -
vss:s_0006     XDP   vsd:dst_0006 Broken-off
                           Idle      -       true    -
vss:s_0007     XDP   vsd:dst_0007 Broken-off
                           Idle      -       true    -
vss:s_0008     XDP   vsd:dst_0008 Broken-off
                           Idle      -       true    -
...

```

The SnapMirror relationship status of each constituent is Broken-off.

4. Verify that the destination FlexGroup volume is read/write: `volume show -vserver svm_name`

```

cluster2::> volume show -vserver vsd
Vserver      Volume      Aggregate      State      Type      Size
Available    Used%
-----  -----
vsd          dst          -            online    **RW**    2GB
1.54GB      22%
vsd          d2           -            online    DP        2GB
1.55GB      22%
vsd          root_vs0    aggr1        online    RW        100MB
94.02MB     5%
3 entries were displayed.

```

5. Redirect clients to the destination FlexGroup volume.

Related information

- [snapmirror break](#)
- [snapmirror quiesce](#)
- [snapmirror show](#)

Reactivate the original source ONTAP FlexGroup volume after disaster

When the source FlexGroup volume becomes available, you can resynchronize the original source and original destination FlexGroup volumes. Any new data on the destination FlexGroup volume is lost.

About this task

Any active quota rules on the destination volume are deactivated and the quota rules are deleted before resynchronization is performed.

You can use the `volume quota policy rule create` and `volume quota modify` commands to create and reactivate quota rules after the resynchronization operation is complete.

Steps

1. From the destination cluster, resynchronize the FlexGroup volume SnapMirror relationship: `snapmirror resync -destination-path dst_svm:dest_flexgroup`
2. View the status of the SnapMirror relationship: `snapmirror show -expand`

```
cluster2::> snapmirror show -expand

Progress
Source          Destination Mirror  Relationship  Total
Last
Path           Type   Path        State   Status
Updated

vss:s          XDP   vsd:dst    Snapmirrored
                           Idle
vss:s_0001     XDP   vsd:dst_0001 Snapmirrored
                           Idle
vss:s_0002     XDP   vsd:dst_0002 Snapmirrored
                           Idle
vss:s_0003     XDP   vsd:dst_0003 Snapmirrored
                           Idle
vss:s_0004     XDP   vsd:dst_0004 Snapmirrored
                           Idle
vss:s_0005     XDP   vsd:dst_0005 Snapmirrored
                           Idle
vss:s_0006     XDP   vsd:dst_0006 Snapmirrored
                           Idle
vss:s_0007     XDP   vsd:dst_0007 Snapmirrored
                           Idle
vss:s_0008     XDP   vsd:dst_0008 Snapmirrored
                           Idle
...

```

The SnapMirror relationship status of each constituent is Snapmirrored.

Related information

- [snapmirror resync](#)
- [snapmirror show](#)

Reverse SnapMirror relationships between ONTAP FlexGroup volumes during disaster recovery

When a disaster disables the source FlexGroup volume of a SnapMirror relationship, you can use the destination FlexGroup volume to serve data while you repair or replace the source FlexGroup volume. After the source FlexGroup volume is online, you can make the original source FlexGroup volume a read-only destination and reverse the SnapMirror relationship.

About this task

Any active quota rules on the destination volume are deactivated and the quota rules are deleted before

resynchronization is performed.

You can use the `volume quota policy rule create` and `volume quota modify` commands to create and reactivate quota rules after the resynchronization operation is complete.

Steps

1. On the original destination FlexGroup volume, remove the data protection mirror relationship between the source FlexGroup volume and the destination FlexGroup volume: `snapmirror delete -destination-path svm_name:volume_name`

```
cluster2::> snapmirror delete -destination-path vsd:dst
```

2. On the original source FlexGroup volume, remove the relationship information from the source FlexGroup volume: `snapmirror release -destination-path svm_name:volume_name -relationship -info-only`

After deleting a SnapMirror relationship, you must remove the relationship information from the source FlexGroup volume before attempting a resynchronization operation.

```
cluster1::> snapmirror release -destination-path vsd:dst -relationship -info-only true
```

3. On the new destination FlexGroup volume, create the mirror relationship: `snapmirror create -source-path src_svm_name:volume_name -destination-path dst_svm_name:volume_name -type XDP -policy MirrorAllSnapshots`

```
cluster1::> snapmirror create -source-path vsd:dst -destination-path vss:src -type XDP -policy MirrorAllSnapshots
```

Learn more about `snapmirror create` in the [ONTAP command reference](#).

4. On the new destination FlexGroup volume, resynchronize the source FlexGroup: `snapmirror resync -source-path svm_name:volume_name`

```
cluster1::> snapmirror resync -source-path vsd:dst
```

5. Monitor the SnapMirror transfers: `snapmirror show -expand`

```
cluster2::> snapmirror show -expand

Progress
Source          Destination Mirror  Relationship  Total
Last
Path           Type   Path        State   Status
Updated

vsd:dst      XDP   vss:src      Snapmirrored
                           Idle      -       true    -
vss:dst_0001 XDP   vss:src_0001 Snapmirrored
                           Idle      -       true    -
vsd:dst_0002 XDP   vss:src_0002 Snapmirrored
                           Idle      -       true    -
vsd:dst_0003 XDP   vss:src_0003 Snapmirrored
                           Idle      -       true    -
vsd:dst_0004 XDP   vss:src_0004 Snapmirrored
                           Idle      -       true    -
vsd:dst_0005 XDP   vss:src_0005 Snapmirrored
                           Idle      -       true    -
vsd:dst_0006 XDP   vss:src_0006 Snapmirrored
                           Idle      -       true    -
vsd:dst_0007 XDP   vss:src_0007 Snapmirrored
                           Idle      -       true    -
vsd:dst_0008 XDP   vss:src_0008 Snapmirrored
                           Idle      -       true    -
...

```

The SnapMirror relationship status of each constituent shows as `Snapmirrored` that indicates that the resynchronization was successful.

Related information

- [snapmirror create](#)
- [snapmirror delete](#)
- [snapmirror release](#)
- [snapmirror resync](#)
- [snapmirror show](#)

Expand FlexGroup volumes in a SnapMirror relationship

Expand ONTAP FlexGroup volumes in a SnapMirror relationship

Beginning with ONTAP 9.3, you can expand the source FlexGroup volume and destination FlexGroup volume that are in a SnapMirror relationship by adding new

constituents to the volumes. You can expand the destination volumes either manually or automatically.

About this task

- This task is not applicable to SVM-DR relationships, which automatically manage the expansion of the FlexGroup volume.
- After expansion, the number of constituents in the source FlexGroup volume and destination FlexGroup volume of a SnapMirror relationship must match.

If the number of constituents in the volumes does not match, the SnapMirror transfers fail.

- You should not perform any SnapMirror operation when the expansion process is in progress.
- If a disaster strikes before the expansion process is complete, you must break the SnapMirror relationship and wait until the operation succeeds.



You should break the SnapMirror relationship when the expansion process is in progress only in the case of a disaster. In the case of a disaster, the break operation can take some time to complete. You should wait for the break operation to get completed successfully before performing a resync operation. If the break operation fails, you must retry the break operation. If the break operation fails, some of the new constituents might remain in the destination FlexGroup volume after the break operation. It is best to delete these constituents manually before proceeding further.

Expand the source ONTAP FlexGroup volume of a SnapMirror relationship

Beginning with ONTAP 9.3, you can expand the source FlexGroup volume of a SnapMirror relationship by adding new constituents to the source volume. You can expand the source volume in the same way that you expand a regular FlexGroup volume (read-write volume).

Steps

1. Expand the source FlexGroup volume: `volume expand -vserver vserver_name -volume fg_src -aggr-list aggregate_name,... [-aggr-list-multiplier constituents_per_aggr]`

```
cluster1::> volume expand -volume src_fg -aggr-list aggr1 -aggr-list -multiplier 2 -vserver vs_src
```

Warning: The following number of constituents of size 50GB will be added to FlexGroup "src_fg": 2.

Expanding the FlexGroup will cause the state of all Snapshot copies to be set to "partial".

Partial Snapshot copies cannot be restored.

Do you want to continue? {y|n}: Y

[Job 146] Job succeeded: Successful

The state of all of the snapshots that are taken before the volume is expanded changes to partial.

Expand the destination ONTAP FlexGroup volume of a SnapMirror relationship

You can expand the destination FlexGroup volume and reestablish the SnapMirror relationship either automatically or manually. By default, the SnapMirror relationship is set for automatic expansion, and the destination FlexGroup volume expands automatically if the source volume expands.

Before you begin

- The source FlexGroup volume must have been expanded.
- The SnapMirror relationship must be in the `SnapMirrored` state.

The SnapMirror relationship must not be broken or deleted.

About this task

- When the destination FlexGroup volume is created, the volume is set up for automatic expansion by default.

You can modify the destination FlexGroup volume for manual expansion, if required.



The best practice is to expand the destination FlexGroup volume automatically.

- All SnapMirror operations fail until both the source FlexGroup volume and destination FlexGroup volume have expanded and have the same number of constituents.
- If you expand the destination FlexGroup volume after the SnapMirror relationship is broken or deleted, you cannot resync the original relationship again.

If you intend to reuse the destination FlexGroup volume, do not expand the volume after deleting the SnapMirror relationship.

Choices

- Perform an update transfer to expand the destination FlexGroup volume automatically:
 - a. Perform a SnapMirror update transfer: `snapmirror update -destination-path svm:vol_name`
 - b. Verify that the status of the SnapMirror relationship is in the `SnapMirrored` state: `snapmirror show`

```

cluster2::> snapmirror show

Progress
Source          Destination Mirror  Relationship  Total
Last
Path           Type   Path        State   Status      Progress
Healthy Updated
-----
-----
vs_src:src_fg
    XDP  vs_dst:dst_fg
                    Snapmirrored
                    Idle
                    -
                    -

```

Based on the size and availability of aggregates, the aggregates are automatically selected, and new constituents that match the constituents of the source FlexGroup volume are added to the destination FlexGroup volume. After expansion, a resynchronization operation is automatically triggered.

- Expand the destination FlexGroup volume manually:

a. If the SnapMirror relationship is in the auto-expand mode, set the SnapMirror relationship to the manual expand mode: `snapmirror modify -destination-path svm:vol_name -is-auto-expand -enabled false`

```

cluster2::> snapmirror modify -destination-path vs_dst:dst_fg -is
           -auto-expand-enabled false
Operation succeeded: snapmirror modify for the relationship with
destination "vs_dst:dst_fg".

```

b. Quiesce the SnapMirror relationship: `snapmirror quiesce -destination-path svm:vol_name`

```

cluster2::> snapmirror quiesce -destination-path vs_dst:dst_fg
Operation succeeded: snapmirror quiesce for destination
"vs_dst:dst_fg".

```

c. Expand the destination FlexGroup volume: `volume expand -vserver vserver_name -volume fg_name -aggr-list aggregate_name,... [-aggr-list-multiplier constituents_per_aggr]`

```
cluster2::> volume expand -volume dst_fg -aggr-list aggr1 -aggr-list
-multiplier 2 -vserver vs_dst

Warning: The following number of constituents of size 50GB will be
added to FlexGroup "dst_fg": 2.
Do you want to continue? {y|n}: y
[Job 68] Job succeeded: Successful
```

d. Resynchronize the SnapMirror relationship: `snapmirror resync -destination-path svm:vol name`

```
cluster2::> snapmirror resync -destination-path vs_dst:dst_fg
Operation is queued: snapmirror resync to destination
"vs dst:dst fg".
```

e. Verify that the status of the SnapMirror relationship is SnapMirrored: snapmirror show

Related information

- snapmirror quiesce
- snapmirror resync
- snapmirror show

Perform a SnapMirror single file restore from an ONTAP FlexGroup volume

Beginning with ONTAP 9.8, you can restore a single file from a FlexGroup SnapMirror vault or from a UDP destination.

About this task

- You can restore from a FlexGroup volume of any geometry to FlexGroup volume of any geometry.
- Only one file per restore operation is supported.
- You can restore to either the original source FlexGroup volume or to a new FlexGroup volume.
- Remote fenced file lookup is not supported.

Single file restore fails if the source file is fenced.

- You can restart or clean up an aborted single file restore.
- You should clean up a failed single file restore transfer by using the `clean-up-failure` option of the `snapmirror restore` command.

Learn more about `snapmirror restore` in the [ONTAP command reference](#).

- Expansion of FlexGroup volumes is supported when a FlexGroup single file restore is in progress or in an aborted state.

Steps

1. Restore a file from a FlexGroup volume: `snapmirror restore -destination-path destination_path -source-path source_path -file-list /f1 -throttle throttle -source-snapshot snapshot`

The following is an example of a FlexGroup volume single file restore operation.

```
vserverA::> snapmirror restore -destination-path vs0:fg2 -source-path
vs0:fgd -file-list /f1 -throttle 5 -source-snapshot snapmirror.81072ce1-
d57b-11e9-94c0-005056a7e422_2159190496.2019-09-19_062631
[Job 135] Job is queued: snapmirror restore from source "vs0:fgd" for
the snapshot snapmirror.81072ce1-d57b-11e9-94c0-
005056a7e422_2159190496.2019-09-19_062631.
vserverA::> snapmirror show

Source          Destination Mirror  Relationship
Total  Last
Path      Type    Path          State   Status      Progress
Healthy Updated
-----  -----  -----  -----  -----
-----  -----  -----  -----
vs0:v1d  RST    vs0:v2        -      Transferring  Idle  83.12KB
true  09/19 11:38:42

vserverA::*> snapmirror show vs0:fg2

Source Path: vs0:fgd
Source Cluster: -
Source Vserver: vs0
Source Volume: fgd
Destination Path: vs0:fg2
```

```
Destination Cluster: -
Destination Vserver: vs0
Destination Volume: fg2
Relationship Type: RST
Relationship Group Type: none
Managing Vserver: vs0
SnapMirror Schedule: -
SnapMirror Policy Type: -
SnapMirror Policy: -
Tries Limit: -
Throttle (KB/sec): unlimited
Current Transfer Throttle (KB/sec): 2
Mirror State: -
Relationship Status: Transferring
File Restore File Count: 1
File Restore File List: f1
Transfer Snapshot: snapmirror.81072ce1-d57b-11e9-94c0-
005056a7e422_2159190496.2019-09-19_062631
Snapshot Progress: 2.87MB
Total Progress: 2.87MB
Network Compression Ratio: 1:1
Snapshot Checkpoint: 2.97KB
Newest Snapshot: -
Newest Snapshot Timestamp: -
Exported Snapshot: -
Exported Snapshot Timestamp: -
Healthy: true
Physical Replica: -
Relationship ID: e6081667-dacb-11e9-94c0-005056a7e422
Source Vserver UUID: 81072ce1-d57b-11e9-94c0-005056a7e422
Destination Vserver UUID: 81072ce1-d57b-11e9-94c0-005056a7e422
Current Operation ID: 138f12e6-dacc-11e9-94c0-005056a7e422
Transfer Type: cg_file_restore
Transfer Error: -
Last Transfer Type: -
Last Transfer Error: -
Last Transfer Error Codes: -
Last Transfer Size: -
Last Transfer Network Compression Ratio: -
Last Transfer Duration: -
Last Transfer From: -
Last Transfer End Timestamp: -
Unhealthy Reason: -
Progress Last Updated: 09/19 07:07:36
Relationship Capability: 8.2 and above
Lag Time: -
```

```
Current Transfer Priority: normal
SMTape Operation: -
Constituent Relationship: false
Destination Volume Node Name: vserverA
Identity Preserve Vserver DR: -
Number of Successful Updates: 0
Number of Failed Updates: 0
Number of Successful Resyncs: 0
Number of Failed Resyncs: 0
Number of Successful Breaks: 0
Number of Failed Breaks: 0
Total Transfer Bytes: 0
Total Transfer Time in Seconds: 0
Source Volume MSIDs Preserved: -
OpMask: ffffffffffffffff
Is Auto Expand Enabled: -
Source Endpoint UUID: -
Destination Endpoint UUID: -
Is Catalog Enabled: false
```

Related information

- [snapmirror show](#)

Restore ONTAP FlexGroup volumes from SnapVault backups

You can perform a full-volume restore operation of FlexGroup volumes from a snapshot in the SnapVault secondary volume. You can restore the FlexGroup volume either to the original source volume or to a new FlexGroup volume.

Before you begin

You must be aware of certain considerations when you restore from SnapVault backups for FlexGroup volumes.

- Only baseline restore is supported with partial snapshots from a SnapVault backup. The number of constituents in the destination volume must match the number of constituents in the source volume when the snapshot was taken.
- If a restore operation fails, no other operations are allowed until the restore operation is complete. You can either retry the restore operation or run the restore operation with the `cleanup` parameter.
- A FlexGroup volume can be the source volume of only one backup relationship or restore relationship. A FlexGroup volume cannot be the source of two SnapVault relationships, two restore relationships, or a SnapVault relationship and a restore relationship.
- SnapVault backup and restore operations cannot run in parallel. When either a baseline restore operation or an incremental restore operation is in progress, you should quiesce the backup operations.
- You must abort a restore operation of a partial snapshot from the destination FlexGroup volume. You cannot abort the restore operation of a partial snapshot from the source volume.

- If you abort a restore operation, you must restart the restore operation with the same snapshot that was used for the previous restore operation.

About this task

Any active quota rules on the destination FlexGroup volume are deactivated before the restore is performed.

You can use the `volume quota modify` command to reactivate quota rules after the restore operation is complete.

Steps

1. Restore the FlexGroup volume: `snapmirror restore -source-path src_svm:src_flexgroup -destination-path dest_svm:dest_flexgroup -snapshot snapshot_name`
`snapshot_name` is the snapshot that is to be restored from the source volume to the destination volume. If the snapshot is not specified, the destination volume is restored from the latest snapshot.

```
vserverA::> snapmirror restore -source-path vserverB:dstFG -destination -path vserverA:newFG -snapshot daily.2016-07-15_0010
Warning: This is a disruptive operation and the volume vserverA:newFG
will be read-only until the operation completes
Do you want to continue? {y|n}: y
```

Related information

- [snapmirror restore](#)

Disable SVM protection on ONTAP FlexGroup volumes

When the SVM DR flag is set to **protected** on a FlexGroup volume, you can set the flag to **unprotected** to disable SVM DR protection on a FlexGroup volume.

Before you begin

- The SVM DR relationship between the primary and secondary is healthy.
- SVM DR protection parameter is set to **protected**.

Steps

1. Disable protection by using the `volume modify` command to change the `vserver-dr-protection` parameter for the FlexGroup volume to **unprotected**.

```
cluster2::> volume modify -vserver vs1 -volume fg_src -vserver-dr -protection unprotected
[Job 5384] Job is queued: Modify fg_src.
[Job 5384] Steps completed: 4 of 4.
cluster2::>
```

2. Update the SVM at the secondary site: `snapmirror update -destination-path destination_svm_name: -source-path Source_svm_name:`

3. Verify that the SnapMirror relationship is healthy: `snapmirror show`
4. Verify that the FlexGroup SnapMirror relationship has been removed: `snapmirror show -expand`

Related information

- [snapmirror show](#)
- [snapmirror update](#)

Enable SVM protection on ONTAP FlexGroup volumes

When the SVM DR protection flag is set to unprotected on a FlexGroup volume, you can set the flag to protected to enable SVM DR protection.

Before you begin

- The SVM DR relationship between the primary and secondary is healthy.
- SVM DR protection parameter is set to unprotected.

Steps

1. Enable protection by using the `volume modify` to change the `vserver-dr-protection` parameter for the FlexGroup volume to protected.

```
cluster2::> volume modify -vserver vs1 -volume fg_src -vserver-dr-protection protected
[Job 5384] Job is queued: Modify fg_src.
[Job 5384] Steps completed: 4 of 4.
cluster2::>
```

2. Update the SVM at the secondary site: `snapmirror update -destination-path destination_svm_name -source-path source_svm_name`

```
snapmirror update -destination-path vs1_dst: -source-path vs1:
```

3. Verify that the SnapMirror relationship is healthy: `snapmirror show`

```
cluster2::> snapmirror show

Progress
Source          Destination Mirror  Relationship  Total
Last
Path           Type   Path       State   Status      Progress  Healthy
Updated

-----
vs1:          XDP   vs1_dst:    Snapmirrored
                           Idle      -         true      -
```

4. Verify that the FlexGroup SnapMirror relationship is healthy: `snapmirror show -expand`

```

cluster2::> snapmirror show -expand

Progress
Source          Destination Mirror  Relationship  Total
Last
Path           Type   Path        State   Status      Progress  Healthy
Updated

-----
-----
```

Source	Destination	Mirror	Relationship	Total
vs1:	XDP	vs1_dst:	Snapmirrored	
			Idle	-
vs1:fg_src	XDP	vs1_dst:fg_src	Snapmirrored	
			Idle	-
vs1:fg_src_0001	XDP	vs1_dst:fg_src_0001	Snapmirrored	
			Idle	-
vs1:fg_src_0002	XDP	vs1_dst:fg_src_0002	Snapmirrored	
			Idle	-
vs1:fg_src_0003	XDP	vs1_dst:fg_src_0003	Snapmirrored	
			Idle	-
vs1:fg_src_0004	XDP	vs1_dst:fg_src_0004	Snapmirrored	
			Idle	-

6 entries were displayed.

Related information

- [snapmirror show](#)

Convert FlexVol volumes to FlexGroup volumes

Learn about converting ONTAP FlexVol volumes to FlexGroup volumes

If you want to expand a FlexVol volume beyond its space limit, you can convert the FlexVol volume to a FlexGroup volume. Beginning with ONTAP 9.7, you can convert standalone FlexVol volumes or FlexVol volumes that are in a SnapMirror relationship to FlexGroup volumes.

Considerations for converting FlexVol volumes to FlexGroup volumes

You should be aware of the [features and operations that are supported](#) before you decide to convert FlexVol volumes to FlexGroup volumes.

Operations not supported during conversion

The following operations are not allowed when volume conversion is in progress:

- Volume move
- Aggregate relocation
- Planned takeover and giveback in a high-availability configuration
- Manual and automatic giveback in an high-availability configuration
- Cluster upgrade and revert
- FlexClone volume split
- Volume rehost
- Volume modify and autosize
- Volume rename
- Attaching an object store to an aggregate
- Negotiated switchover in MetroCluster configuration
- SnapMirror operations
- Restoring from a snapshot
- Quota operations
- Storage efficiency operations

You can perform these operations on the FlexGroup volume after successful conversion.

Configurations that are not supported with FlexGroup volumes

- Offline or restricted volume
- SVM root volume
- SAN
- SMB 1.0
- NVMe namespaces
- Remote Volume Shadow Copy Service (VSS)

Convert ONTAP FlexVol volumes to ONTAP FlexGroup volumes

Beginning with ONTAP 9.7, you can perform an in-place conversion of a FlexVol volume to a FlexGroup volume without requiring a data copy or additional disk space.

Before you begin

- Transitioned volumes can be converted to FlexGroup volumes beginning with ONTAP 9.8.
- The FlexVol volume that is being converted must be online.

- The operations and configurations on the FlexVol volume must be compatible with the conversion process.

Check for the following conditions that can prevent the conversion from succeeding:

- A FlexVol volume was transitioned from 7-Mode using 7MTT (ONTAP 9.7).

Transitioned volumes can be converted beginning with ONTAP 9.8.

- Something is enabled on the volume that is not yet supported with FlexGroup volume; for example, SAN LUNs, Windows NFS, SMB1, snapshot naming/autodelete, vmalign set, SnapLock with releases earlier than ONTAP 9.11.1 (SnapLock is supported beginning with ONTAP 9.11.1), space SLO, or logical space enforcement/reporting. For more information see [Supported and unsupported configurations for FlexGroup volumes](#).
- The SVM where the FlexVol volume to be converted is located is currently using SVM DR.
- NetApp FlexClone volumes are present, and the FlexVol volume is the parent volume. The volume being converted cannot be a parent or a clone.
- The volume is a NetApp FlexCache origin volume.
- For ONTAP 9.7 and earlier, NetApp snapshots must not exceed 255. For ONTAP 9.8 and later, 1023 snapshots are supported.
- Storage efficiencies are enabled. These must be disabled and can be reenabled after conversion.
- The volume is a source of a SnapMirror relationship, and the destination has not yet been converted.
- The volume is part of an active (not quiesced) SnapMirror relationship.
- Autonomous Ransomware Protection (ARP) has been disabled on the volume. You should not enable it again until the conversion is complete.
- Quotas are enabled. These must be disabled and can be reenabled after conversion.
- Volume names are longer than 197 characters.
- The volume is associated with an application.

This is applicable to ONTAP 9.7 only. The limitation is removed in ONTAP 9.8.

- ONTAP processes are running, such as mirroring, jobs, wafliron, NDMP backup, and inode conversion in process.
- The volume is an SVM root volume.
- The volume is too full.

If any of these incompatibilities exist, an error message is generated if the FlexVol volume, and the volume conversion is aborted. You can take corrective actions and retry the conversion.

- If a FlexVol volume is currently at 80% or greater maximum capacity, consider copying the data to a newly created FlexGroup volume instead of performing an in-place conversion. Although FlexGroup member volumes will naturally rebalance over time, converting a high-capacity FlexVol volume to a FlexGroup volume may create performance or balance issues that will not quickly be rebalanced across member volumes.



Converting a very large FlexGroup volume results in a very full FlexGroup volume member constituent, which can create performance issues. For more information, see the section called "When not to create a FlexGroup volume" in the TR [FlexGroup volumes - Best Practices and Implementation Guide](#).

Steps

1. Verify that the FlexVol volume is online: `volume show -fields vol_name volume-style-extended, state`

```
cluster-1::> volume show my_volume -fields volume-style-extended, state
vserver volume      state  volume-style-extended
-----
vs0      my_volume online flexvol
```

2. Verify whether the FlexVol volume can be converted without issues:

- a. Log in to the advance privilege mode: `set -privilege advanced`
- b. Verify the conversion process: `volume conversion start -vserver vs1 -volume flexvol -check-only true`

You must rectify all errors before converting the volume.



You cannot convert a FlexGroup volume back to a FlexVol volume.

3. Start the conversion: `volume conversion start -vserver svm_name -volume vol_name`

```
cluster-1::*> volume conversion start -vserver vs0 -volume my_volume

Warning: Converting flexible volume "my_volume" in Vserver "vs0" to a
FlexGroup
      will cause the state of all Snapshot copies from the volume to
be set
      to "pre-conversion". Pre-conversion Snapshot copies cannot be
restored.

Do you want to continue? {y|n}: y
[Job 57] Job succeeded: success
```

4. Verify that the conversion is successful: `volume show vol_name -fields volume-style-extended, state`

```
cluster-1::*> volume show my_volume -fields volume-style-extended, state
vserver volume      state  volume-style-extended
-----
vs0      my_volume online flexgroup
```

Results

The FlexVol volume is converted to a single-member FlexGroup volume.

After you finish

You can expand the FlexGroup volume, as required.

Convert ONTAP FlexVol volume SnapMirror relationships to ONTAP FlexGroup volume SnapMirror relationships

To convert a FlexVol volume SnapMirror relationship to a FlexGroup volume SnapMirror relationship in ONTAP, you must first convert the destination FlexVol volume followed by the source FlexVol volume.

About this task

- Mixing FlexGroups with FlexVol volumes in SnapMirror/SnapVault relationships is not supported beyond the conversion process.
- FlexGroup conversion is supported only for SnapMirror asynchronous relationships.
- FlexGroup conversion is not supported in SnapMirror cloud relationships.
- Conversion time depends on several variables. Some of the variables include:
 - CPU of the controller
 - Utilization of CPU by other applications
 - Amount of data in the initial snapshot
 - Network bandwidth
 - Bandwidth used by other applications

Before you begin

- The FlexVol volume that is being converted must be online.
- The source FlexVol volume in the SnapMirror relationship must not be the source volume for multiple SnapMirror relationships.

Beginning with ONTAP 9.9.1, fanout SnapMirror relationships are supported for FlexGroup volumes. For more information, see [Considerations for creating SnapMirror cascade and fanout relationships for FlexGroup volumes](#).

- The operations and configurations on the FlexVol volume must be compatible with the conversion process.

An error message is generated if the FlexVol volume has any incompatibility and the volume conversion is aborted. You can take corrective actions and retry the conversion.

Steps

1. Verify that the SnapMirror relationship is healthy:

```
snapmirror show
```

Only XDP type mirror relationships can be converted.

Example:

```

cluster2::> snapmirror show

Progress
Source          Destination Mirror  Relationship  Total
Last
Path           Type   Path       State   Status      Progress  Healthy
Updated

-----
----- vs0:src_dpv  DP    vs2:dst_dpv  Snapmirrored
                               Idle          -        true     -
----- vs0:src_xdp  XDP   vs2:dst_xdp  Snapmirrored
                               Idle          -        true     -

```

2. Verify whether the source volume is compatible for conversion:

- Log in to the advance privilege mode:

```
set -privilege advanced
```

- Verify the conversion process:

```
volume conversion start -vserver <src_svm_name> -volume <src_vol>
-check-only true
```

Example:

```
volume conversion start -vserver vs1 -volume src_vol -check-only true
```

You must rectify all errors before converting the volume.

3. Convert the destination FlexVol volume to FlexGroup volume.

- Quiesce the FlexVol SnapMirror relationship:

```
snapmirror quiesce -destination-path <dest_svm:dest_volume>
```

Example:

```
cluster2::> snapmirror quiesce -destination-path vs2:dst_xdp
```

- Start the conversion:

```
volume conversion start -vserver <dest_svm> -volume <dest_volume>
```

Example:

```
cluster-1::> volume conversion start -vserver vs2 -volume dst_xdp

Warning: After the volume is converted to a FlexGroup, it will not be
possible
to change it back to a flexible volume.
Do you want to continue? {y|n}: y

[Job 510] Job succeeded: SnapMirror destination volume "dst_xdp" has
been successfully converted to a FlexGroup volume.
You must now convert the relationship's source volume, "vs0:src_xdp",
to a FlexGroup.
Then, re-establish the SnapMirror relationship using the "snapmirror
resync" command.
```

4. Convert the source FlexVol volume to FlexGroup volume:

```
volume conversion start -vserver <src_svm_name> -volume <src_vol_name>
```

Example:

```
cluster-1::> volume conversion start -vserver vs0 -volume src_xdp

Warning: Converting flexible volume "src_xdp" in Vserver "vs0" to a
FlexGroup
      will cause the state of all Snapshot copies from the volume to
be set
      to "pre-conversion". Pre-conversion snapshots cannot be
restored.
Do you want to continue? {y|n}: y
[Job 57] Job succeeded: success
```

5. Resync the relationship:

```
snapmirror resync -destination-path dest_svm_name:dest_volume
```

Example:

```
cluster2::> snapmirror resync -destination-path vs2:dst_xdp
```

After you finish

You must ensure that when the source FlexGroup volume is expanded to include more constituents, the destination volume is also expanded.

Related information

- [snapmirror quiesce](#)
- [snapmirror resync](#)
- [snapmirror show](#)

FlexCache volumes management

Learn about ONTAP FlexCache volumes

NetApp FlexCache technology accelerates data access, reduces WAN latency and lowers WAN bandwidth costs for read-intensive workloads, especially where clients need to access the same data repeatedly. When you create a FlexCache volume, you create a remote cache of an already existing (origin) volume that contains only the actively accessed data (hot data) of the origin volume.

When a FlexCache volume receives a read request of the hot data it contains, it can respond faster than the origin volume because the data does not need to travel as far to reach the client. If a FlexCache volume receives a read request for infrequently read data (cold data), it retrieves the needed data from the origin volume and then stores the data before serving the client request. Subsequent read requests for that data are then served directly from the FlexCache volume. After the first request, the data no longer needs to travel across the network, or be served from a heavily loaded system. For example, suppose you are experiencing bottlenecks within your cluster at a singular access point for frequently requested data. You can use FlexCache volumes within the cluster to provide multiple mount points to the hot data, thereby reducing the bottlenecks and increasing performance. As another example, suppose you need to decrease network traffic to a volume that is accessed from multiple clusters. You can use FlexCache volumes to distribute hot data from the origin volume across the clusters within your network. This reduces WAN traffic by giving users closer access points.

You can also use FlexCache technology to improve performance in cloud and hybrid cloud environments. A FlexCache volume can help you transition workloads to the hybrid cloud by caching data from an on-premises data center to cloud. You can also use FlexCache volumes to remove cloud silos by caching data from one cloud provider to another or between two regions of the same cloud provider.

Beginning with ONTAP 9.10.1, you can [enable global file locking](#) across all FlexCache volumes. Global file locking prevents a user from accessing a file that is already opened by another user. Updates to the origin volume are then distributed to all FlexCache volumes simultaneously.


Beginning with ONTAP 9.9.1, FlexCache volumes maintain a list of files not found. This helps reduce network traffic by removing the need to send multiple calls to the origin when clients search for non-existent files.

A list of additional [features supported for FlexCache volumes and their origin volumes](#), including a list of supported protocols by ONTAP version, is also available.


You can learn more about the architecture of ONTAP FlexCache technology in [TR-4743: FlexCache in ONTAP](#).

Videos

How FlexCache can reduce WAN latency and read times for global data



Learn about the performance benefits of ONTAP FlexCache!



Supported and unsupported features for ONTAP FlexCache volumes

Beginning with ONTAP 9.5, you can configure FlexCache volumes. FlexVol volumes are supported as origin volumes and FlexGroup volumes are supported as FlexCache volumes. Beginning with ONTAP 9.7 both FlexVol volumes and FlexGroup volumes are supported as origin volumes. The supported features and protocols for the origin volume and the FlexCache volume vary.



Cache volumes and origin volumes can interoperate as long as both are running on a supported version of ONTAP. Keep in mind that features are supported only when both the cache and the origin are running at least the ONTAP version where support was introduced or a later ONTAP version.

ONTAP version support between FlexCache volumes and origin volumes

The recommended ONTAP version supported between origin volume and the cache volume is no more than four versions earlier or four versions later. For example, if the cache is using ONTAP 9.14.1, the earliest version the origin can be using is ONTAP 9.10.1.

Supported protocols

Protocol	Supported at the origin volume?	Supported at the FlexCache volume?
NFSv3	Yes	Yes
NFSv4	Yes To access cache volumes using NFSv4.x protocol, both the origin and cache clusters must be using ONTAP 9.10.1 or later. The origin cluster and FlexCache cluster can have different ONTAP versions, but both should be ONTAP 9.10.1 and later versions, for example, the origin can have ONTAP 9.10.1, and the cache can have ONTAP 9.11.1.	Yes Supported beginning with ONTAP 9.10.1. To access cache volumes using NFSv4.x protocol, both the origin and cache clusters must be using ONTAP 9.10.1 or later. The origin cluster and FlexCache cluster can have different ONTAP versions, but both should be ONTAP 9.10.1 and later versions, for example, the origin can have ONTAP 9.10.1, and the cache can have ONTAP 9.11.1.
NFSv4.2	Yes	No
SMB	Yes	Yes Supported beginning with ONTAP 9.8.

Supported features

Feature	Supported at the origin volume?	Supported at the FlexCache volume?
Autonomous ransomware protection	<p>Yes</p> <p>Supported for FlexVol origin volumes beginning with ONTAP 9.10.1, and supported for FlexGroup origin volumes beginning with ONTAP 9.13.1. See Autonomous Ransomware Protection use cases and considerations.</p>	No
Antivirus	<p>Yes</p> <p>Supported beginning with ONTAP 9.7.</p>	<p>Not applicable</p> <p>If you configure antivirus scanning at the origin, it is not required on the cache. The origin antivirus scanning detects files infected with viruses before writes are committed, regardless of the write source. For more information about using antivirus scanning with FlexCache, see the FlexCache with ONTAP technical report.</p>
Auditing	<p>Yes</p> <p>Supported beginning with ONTAP 9.7.</p> <p>You can audit NFS file access events in FlexCache relationships using native ONTAP auditing. For more information, see Considerations for auditing FlexCache volumes</p>	<p>Yes</p> <p>Supported beginning with ONTAP 9.7.</p> <p>You can audit NFS file access events in FlexCache relationships using native ONTAP auditing. For more information, see Considerations for auditing FlexCache volumes</p>
Cloud Volumes ONTAP	<p>Yes</p> <p>Supported beginning with ONTAP 9.6</p>	<p>Yes</p> <p>Supported beginning with ONTAP 9.6</p>
Compaction	<p>Yes</p> <p>Supported beginning with ONTAP 9.6</p>	<p>Yes</p> <p>Supported beginning with ONTAP 9.7</p>

Compression	Yes Supported beginning with ONTAP 9.6	Yes Supported beginning with ONTAP 9.6
Deduplication	Yes	Yes Inline deduplication is supported on FlexCache volumes beginning with ONTAP 9.6. Cross-volume deduplication is supported on FlexCache volumes beginning with ONTAP 9.7.
FabricPool	Yes	Yes Supported beginning with ONTAP 9.7  You can create a FlexCache volume as a cache for an origin volume that has FabricPool tiering enabled, but the FlexCache volume itself cannot be tiered.
FlexCache DR	Yes	Yes Supported beginning with ONTAP 9.9.1, with NFSv3 protocol, only. FlexCache volumes must be in separate SVMs or in separate clusters.
FlexGroup volume	Yes Supported beginning with ONTAP 9.7	Yes
FlexVol volume	Yes	No
FPolicy	Yes Supported beginning with ONTAP 9.7	Yes Supported for NFS beginning with ONTAP 9.7. Supported for SMB beginning with ONTAP 9.14.1.

MetroCluster configuration	Yes Supported beginning with ONTAP 9.7	Yes Supported beginning with ONTAP 9.7
Microsoft Offloaded Data Transfer (ODX)	Yes	No
NetApp Aggregate Encryption (NAE)	Yes Supported beginning with ONTAP 9.6	Yes Supported beginning with ONTAP 9.6
NetApp Volume Encryption (NVE)	Yes Supported beginning with ONTAP 9.6	Yes Supported beginning with ONTAP 9.6
ONTAP S3 NAS bucket	Yes Supported beginning with ONTAP 9.12.1	Yes Supported beginning with ONTAP 9.18.1
QoS	Yes	Yes  File-level QoS is not supported for FlexCache volumes.
Qtrees	Yes Beginning with ONTAP 9.6, you can create and modify qtrees. Qtrees created on the source can be accessed on the cache.	No
Quotas	Yes Beginning with ONTAP 9.6, quota enforcement on FlexCache origin volumes is supported for users, groups, and qtrees.	No  Beginning with ONTAP 9.6, remote quota (rquota) is supported at FlexCache volumes. With FlexCache writearound mode (the default mode), writes on the cache are forwarded to the origin volume. Quotas are enforced at the origin.

SMB Change Notify	Yes	Yes Beginning with ONTAP 9.14.1, SMB Change Notify is supported at the cache.
SnapLock volumes	No	No
SnapMirror asynchronous relationships*	Yes	No
*FlexCache origins:		
		<ul style="list-style-type: none"> • You can have a FlexCache volume from an origin FlexVol • You can have a FlexCache volume from an origin FlexGroup • You can have a FlexCache volume from an origin primary volume in SnapMirror relationship. • Beginning with ONTAP 9.8, a SnapMirror secondary volume can be a FlexCache origin volume. The SnapMirror secondary volume must be idle with no active SnapMirror updates; otherwise, FlexCache creation fails.
SnapMirror synchronous relationships	No	No
SnapRestore	Yes	No
Snapshots	Yes	No
SVM DR configuration	<p>Yes</p> <p>Supported beginning with ONTAP 9.5. The primary SVM of an SVM DR relationship can have the origin volume; however, if you are running an ONTAP release earlier than ONTAP 9.18.1, when the SVM DR relationship is broken, the FlexCache relationship must be re-created with a new origin volume.</p> <p>Beginning with ONTAP 9.18.1, when an origin SVM fails over, caches automatically switch to the origin at the DR site. Manual recovery steps are eliminated.</p> <p>Learn about creating FlexCache volumes.</p>	<p>No</p> <p>You can have FlexCache volumes in primary SVMs, but not in secondary SVMs. Any FlexCache volume in the primary SVM is not replicated as part of the SVM DR relationship.</p>

Storage-level Access Guard (SLAG)	No	No
Thin provisioning	Yes	Yes Supported beginning with ONTAP 9.7
Volume cloning	Yes Cloning of an origin volume and the files in the origin volume is supported beginning with ONTAP 9.6.	No
Volume move	Yes	Yes (only for volume constituents) Moving volume constituents of a FlexCache volume is supported with ONTAP 9.6 and later.
Volume rehost	No	No
vStorage API for Array Integration (VAAI)	Yes	No

 In ONTAP 9 releases earlier than 9.5, origin FlexVol volumes can only serve data to FlexCache volumes created on systems running Data ONTAP 8.2.x operating in 7-Mode. Beginning with ONTAP 9.5, origin FlexVol volumes can also serve data to FlexCache volumes on ONTAP 9 systems. For information about migrating from 7-Mode FlexCache to ONTAP 9 FlexCache see [NetApp Technical Report 4743: FlexCache in ONTAP](#).

Guidelines for sizing ONTAP FlexCache volumes

You must be aware of the limits for FlexCache volumes before you start provisioning the volumes.

The size limit of a FlexVol volume is applicable to an origin volume. The size of a FlexCache volume can be less than or equal to the origin volume. The best practice for the size of a FlexCache volume is to be at least 10 percent of the size of the origin volume.

You must also be aware of the following additional limits on FlexCache volumes:

Limit	ONTAP 9.8 and later	ONTAP 9.7	ONTAP 9.6 - 9.5
Maximum number of FlexCache volumes that you can create from an origin volume	100	10	10

Recommended maximum number of origin volumes per node	100	100	10
Recommended maximum number of FlexCache volumes per node	100	100	10
Recommended maximum number of FlexGroup constituents in a FlexCache volume per node	800	800	40
Maximum number of constituents per FlexCache volume per node	32	32	32

Related information

- [NetApp Interoperability](#)

Create ONTAP FlexCache volumes

You can create a FlexCache volume in the same ONTAP cluster for improving performance when accessing a hot object. If you have data centers in different locations, you can create FlexCache volumes on remote ONTAP clusters for accelerating data access.

About this task

- Beginning with ONTAP 9.18.1, you can enable NAS S3 bucket access on a FlexCache volume by setting the `-is-s3-enabled` option to `true` when you create the volume. This option is disabled by default.
- Beginning with ONTAP 9.18.1, FlexCache supports creating cache volumes for origin volumes with SVMs that belong to an SVM-DR relationship.

If you are running ONTAP 9.18.1 or later, a storage administrator must peer the cache SVMs with both the primary and secondary origin SVMs that are part of an SVM-DR relationship before creating cache volumes of origin volumes that are part of SVM-DR relationship.

- Beginning with ONTAP 9.14.0, you can create an unencrypted FlexCache volume from an encrypted source.
- Beginning with ONTAP 9.7, both FlexVol volume and FlexGroup volumes are supported as origin volumes.
- Beginning with ONTAP 9.5, FlexCache supports FlexVol volumes as origin volumes and FlexGroup volumes as FlexCache volumes.

Before you begin

- You must be running ONTAP 9.5 or later.
- If you are running ONTAP 9.6 or earlier, you must [add a FlexCache license](#).

A FlexCache license is not required for ONTAP 9.7 or later. Beginning with ONTAP 9.7, FlexCache functionality is included with ONTAP and no longer requires a license or activation.

 If an HA pair is using [encrypting SAS or NVMe drives \(SED, NSE, FIPS\)](#), you must follow the instructions in the topic [Returning a FIPS drive or SED to unprotected mode](#) for all drives within the HA pair prior to initializing the system (boot options 4 or 9). Failure to do this may result in future data loss if the drives are repurposed.

Example 2. Steps

System Manager

1. If the FlexCache volume is on a different ONTAP cluster than the origin volume, create a cluster peer relationship:
 - a. In the local cluster, click **Protection > Overview**.
 - b. Expand **Intercluster Settings**, click **Add Network Interfaces** and add intercluster network interfaces for the cluster.

Repeat this step on the remote cluster.

 - c. In the remote cluster, click **Protection > Overview**. Click **⋮** in the Cluster Peers section and click **Generate Passphrase**.
 - d. Copy the generated passphrase and paste it in the local cluster.
 - e. In the local cluster, under Cluster Peers, click **Peer Clusters** and peer the local and remote clusters.
2. Create an SVM peer relationship:

Under Storage VM Peers, click **⋮** and then **Peer Storage VMs** to peer the storage VMs.

3. Select **Storage > Volumes**.
4. Select **Add**.
5. Select **More Options** and then select **Add as cache for a remote volume**.



If you are running ONTAP 9.8 or later and you want to disable QoS or choose a custom QoS policy, click **More Options**, and then under **Storage and Optimization**, select **Performance Service Level**.

CLI

1. If the FlexCache volume to be created is in a different cluster, create a cluster peer relationship:
 - a. On the destination cluster, create a peer relationship with the data protection source cluster:

```
cluster peer create -generate-passphrase -offer-expiration
MM/DD/YYYY HH:MM:SS|1...7days|1...168hours -peer-addrs
<peer_LIF_IPs> -initial-allowed-vserver-peers <svm_name>,...|*
-ipspace <ipspace_name>
```

Beginning with ONTAP 9.6, TLS encryption is enabled by default when creating a cluster peer relationship. TLS encryption is supported for the intercluster communication between the origin and FlexCache volumes. You can also disable TLS encryption for the cluster peer relationship, if required.

```
cluster02::> cluster peer create -generate-passphrase -offer
-expiration 2days -initial-allowed-vserver-peers *

          Passphrase: UCa+6lRVICXeL/gq1WrK7ShR
          Expiration Time: 6/7/2017 08:16:10 EST
          Initial Allowed Vserver Peers: *
          Intercluster LIF IP: 192.140.112.101
          Peer Cluster Name: Clus_7ShR (temporary generated)
```

Warning: make a note of the passphrase - it cannot be displayed again.

b. On the source cluster, authenticate the source cluster to the destination cluster:

```
cluster peer create -peer-addrs <peer_LIF_IPs> -ipspace <ipspace>
```

```
cluster01::> cluster peer create -peer-addrs
192.140.112.101,192.140.112.102
```

Notice: Use a generated passphrase or choose a passphrase of 8 or more characters.

To ensure the authenticity of the peering relationship, use a phrase or sequence of characters that would be hard to guess.

Enter the passphrase:

Confirm the passphrase:

Clusters cluster02 and cluster01 are peered.

2. If the FlexCache volume is in a different SVM than that of the origin volume, create an SVM peer relationship with `flexcache` as the application:

a. If the SVM is in a different cluster, create an SVM permission for the peering SVMs:

```
vserver peer permission create -peer-cluster <cluster_name>
-vserver <svm-name> -applications flexcache
```

The following example illustrates how to create an SVM peer permission that applies for all of the local SVMs:

```
cluster1::> vserver peer permission create -peer-cluster cluster2
-vserver "*" -applications flexcache

Warning: This Vserver peer permission applies to all local
Vservers. After that no explicit
"vserver peer accept" command required for Vserver peer
relationship creation request
from peer cluster "cluster2" with any of the local Vservers. Do
you want to continue? {y|n}: y
```

b. Create the SVM peer relationship:

```
vserver peer create -vserver <local_SVM> -peer-vserver
<remote_SVM> -peer-cluster <cluster_name> -applications flexcache
```

3. Create a FlexCache volume:

```
volume flexcache create -vserver <cache_svm> -volume
<cache_vol_name> -auto-provision-as flexgroup -size <vol_size>
-origin-vserver <origin_svm> -origin-volume <origin_vol_name> -is-s3
-enabled true|false
```

The following example creates a FlexCache volume and automatically selects existing aggregates for provisioning:

```
cluster1::> volume flexcache create -vserver vs_1 -volume fc1 -auto
-provision-as flexgroup -origin-volume vol_1 -size 160MB -origin
-vserver vs_1
[Job 443] Job succeeded: Successful
```

The following example creates a FlexCache volume and sets the junction path:

```
cluster1::> volume flexcache create -vserver vs34 -volume fc4 -aggr
-list aggr34,aggr43 -origin-volume origin1 -size 400m -junction-path
/fc4
[Job 903] Job succeeded: Successful
```

The following example enables S3 access on a FlexCache volume:

```
cluster1::> volume flexcache create -vserver vs3 -volume
cache_vs3_vol33 -origin-volume vol33 -origin-vserver vs3 -junction
-path /cache_vs3_vol33 -is-s3-enabled true
```

4. Verify the FlexCache relationship from the FlexCache volume and the origin volume.

- View the FlexCache relationship in the cluster:

```
volume flexcache show
```

```
cluster1::> volume flexcache show
Vserver Volume      Size      Origin-Vserver Origin-Volume
Origin-Cluster
-----
-----
vs_1   fc1          160MB    vs_1          vol_1
cluster1
```

- View all of the FlexCache relationships in the origin cluster:

```
volume flexcache origin show-caches
```

```
cluster::> volume flexcache origin show-caches
Origin-Vserver Origin-Volume  Cache-Vserver  Cache-Volume
Cache-Cluster
-----
-----
vs0          ovoll        vs1          cfg1
clusA
vs0          ovoll        vs2          cfg2
clusB
vs_1          vol_1       vs_1          fc1
cluster1
```

Result

The FlexCache volume is successfully created. Clients can mount the volume by using the junction path of the FlexCache volume.

Related information

[Cluster and SVM peering](#)

FlexCache write-back

Learn about ONTAP FlexCache write-back

Introduced in ONTAP 9.15.1, FlexCache write-back is an alternate mode of operation for writing at a cache. Write-back allows the write to be committed to stable storage at the cache and acknowledged to the client without waiting for the data to make it to the origin. The data is asynchronously flushed back to the origin. The result is a globally distributed file system that enables writes to perform at near-local speeds for specific workloads and environments, offering significant performance benefits.



ONTAP 9.12.1 introduced a write-back feature as a public preview. This is referred to as write-back version 1 (wbv1) and shouldn't be thought of as the same as write-back in ONTAP 9.15.1, which is referred to as write-back version 2 (wbv2).

Write-back vs write-around

Since FlexCache was introduced in ONTAP 9.5, it has been a read-writable cache; however, it operated in write-around mode. Writes at the cache were shipped to the origin to be committed to stable storage. After the origin successfully committed the write to stable storage, it acknowledged the write to the cache. The cache would then acknowledge the write to the client. This made every write incur the penalty of traversing the network between the cache and origin. FlexCache write-back changes this.



After upgrading to ONTAP 9.15.1, you can convert a traditional write-around cache to a write-back cache, and, if necessary, back to write-around. This can, however, make reading diagnostic logs harder should a problem arise.

	Write-around	Write-back
ONTAP Version	9.6+	9.15.1+
Use case	Read-heavy workload	Write-heavy workload
Data committed at	Origin	Cache
Client experience	WAN-like	LAN-like
Limits	100 per origin	10 per origin
CAP Theorem	Available and tolerant to partition	Available and consistent

FlexCache write-back terminology

Understand key concepts and terms working with FlexCache write-back.

Term	Definition
Dirty data	Data that has been committed to stable storage at the cache, but has not been flushed to the origin.
Exclusive Lock Delegation (XLD)	A protocol-level lock authority granted on a per-file basis to a cache. This authority allows the cache to hand out exclusive write locks to clients without contacting the origin.

Term	Definition
Shared Lock Delegation (SLD)	A protocol-level lock authority granted on a per-file basis to a cache. This authority allows the cache to hand out shared read locks to clients without contacting the origin.
Write-back	A mode of FlexCache operation where writes to a cache are committed to stable storage at that cache and immediately acknowledged to the client. Data is asynchronously written back to the origin.
Write-around	A mode of FlexCache operation where writes to a cache are forwarded to the origin to be committed to stable storage. Once committed, the origin will acknowledge the write to the cache, and the cache will acknowledge the write to the client.
Dirty Data Record System (DDRS)	A proprietary mechanism that keeps track of the dirty data in a write-back-enabled cache on a per-file basis.
Origin	A FlexGroup or FlexVol that contains the source data for all FlexCache cache volumes. It is the single source of truth, orchestrates locking, and ensures 100% data consistency, currency, and coherency.
Cache	A FlexGroup that is a sparse cache volume of the FlexCache origin.

Consistent, current, and coherent

FlexCache is NetApp's solution to having the right data, everywhere, every time. FlexCache is 100% consistent, current, and coherent 100% of the time:

- **Consistent:** The data is the same wherever it is accessed.
- **Current:** The data is always up-to-date.
- **Coherent:** The data is correct/uncorrupted.

ONTAP FlexCache write-back guidelines

FlexCache write-back involves many complex interactions between the origin and caches. For optimal performance, you should ensure your environment follows these guidelines. These guidelines are based on the latest major ONTAP version (ONTAP 9.17.1.) available at the time of content creation.

As a best practice, test your production workload in a non-production environment. This is even more important if you are implementing FlexCache write-back outside of these guidelines.

The following guidelines are well-tested internally at NetApp. It is **strongly** recommended you stay within them. If you do not, unexpected behavior could occur.

- Significant enhancements for FlexCache write-back were introduced in ONTAP 9.17.1P1. It is **strongly** advised you run the current recommended release after 9.17.1P1 at both the origin and cache clusters. If you are unable to run 9.17.1 codeline, the latest P release of 9.16.1 is the next suggested release. ONTAP 9.15.1 does not have all the necessary fixes and improvements for FlexCache write-back, and is not recommended for production workloads.
- In its current iteration, FlexCache write-back caches should be configured with a single constituent for the

entire FlexCache volume. Multi-constituent FlexCaches can result in unwanted evictions of data from the cache.

- Testing has been executed for files smaller than 100GB and WAN round-trip times between the cache and origin not exceeding 200ms. Any workloads outside of these limits might result in unexpected performance characteristics.
- Writing to SMB alternate data streams causes the main file to be evicted from the cache. All dirty data for the main file needs to be flushed to the origin before any other operations can take place on that file. The alternate data stream is also forwarded to the origin.
- Renaming a file causes the file to be evicted from the cache. All dirty data for the file needs to be flushed to the origin before any other operations can take place on that file.
- At this time, the only attributes that can be changed or set on a file on the write-back-enabled FlexCache volume are:
 - Timestamps
 - Mode bits
 - NT ACLs
 - Owner
 - Group
 - Size

Any other attributes that are changed or set are forwarded to origin which might result in evicting the file from the cache. If you require other attributes to be changed or set at the cache, ask your account team to open a PVR.

- Snapshots taken at the origin cause recalling all outstanding dirty data from every write-back-enabled cache associated with that origin volume. This might require multiple retries of the operation if there is significant write-back activity in progress, as evicts of those dirty files might take some time.
- SMB Opportunistic Locks (Oplocks) for writes are not supported on write-back-enabled FlexCache volumes.
- The origin must remain under 80% full. Cache volumes are not granted exclusive lock delegations if there isn't at least 20% space remaining in the origin volume. Calls to a write-back-enabled cache are forwarded to the origin in this situation. This helps prevent running out of space at the origin, which would result in leaving dirty data orphaned at a write-back-enabled cache.
- Low bandwidth and/or lossy intercluster networks can have a significant negative effect on FlexCache write-back performance. While there isn't a specific bandwidth requirement, as it is highly dependent on your workload, it is **strongly** recommended you ensure the health of the intercluster link between the cache(s) and origin.

ONTAP FlexCache write-back architecture

FlexCache was designed with strong consistency in mind, including both modes of write operation: write-back and write-around. Both the traditional write-around mode of operation and the new write-back mode of operation introduced in ONTAP 9.15.1 guarantee that the data accessed will always be 100% consistent, current, and coherent.

The following concepts detail how FlexCache write-back operates.

Delegations

Lock delegations and data delegations help FlexCache keep both write-back and write-around caches data consistent, coherent, and current. The origin orchestrates both delegations.

Lock delegations

A lock delegation is a protocol-level lock authority the origin grants on a per-file basis to a cache to issue protocol locks to clients as needed. These include [exclusive lock delegations \(XLD\)](#) and [shared lock delegations \(SLD\)](#).

XLD and write-back

To ensure ONTAP never has to reconcile a conflicting write, an XLD is granted to a cache where a client requests to write to a file. Importantly, only one XLD can exist for any file at any time, meaning there will never be more than one writer to a file at a time.

When the request to write to a file comes into a write-back enabled cache, the following steps take place:

1. The cache checks if it already has an XLD for the requested file. If so, it will grant the write lock to the client as long as another client isn't writing to the file at the cache. If the cache doesn't have an XLD for the requested file, it will request one from the origin. This is a proprietary call that traverses the intercluster network.
2. Upon receiving the XLD request from the cache, the origin will check if there is an outstanding XLD for the file at another cache. If so, it will recall that file's XLD, which triggers a flush of any [dirty data](#) from that cache back to the origin.
3. Once the dirty data from that cache is flushed back and committed to stable storage at the origin, the origin will grant the XLD for the file to the requesting cache.
4. Once the file's XLD is received, the cache grants the lock to the client, and the write commences.

A high-level sequence diagram covering some of these steps is covered in the [Write-back sequence diagram](#).

From a client perspective, all locking will work as if it were writing to a standard FlexVol or FlexGroup with a potential small delay when the write lock is requested.

In its current iteration, if a write-back enabled cache holds the XLD for a file, ONTAP will block **any** access to that file at other caches, including `READ` operations.



There is a limit of 170 XLDs per origin constituent.

Data delegations

A data delegation is a per-file guarantee given to a cache by the origin that the data cached for that file is up-to-date. As long as the cache has a data delegation for a file, it can serve the cached data for that file to the client without having to contact the origin. If the cache doesn't have a data delegation for the file, it must contact the origin to receive the data requested by the client.

In write-back mode, a file's data delegation is revoked if an XLD is taken for that file at another cache or the origin. This effectively fences off the file from clients at all other caches and the origin, even for reads. This is a trade off that must be made to ensure old data is never accessed.

Reads at a write-back-enabled cache generally operate like reads at a write-around cache. In both write-around and write-back-enabled caches, there could be an initial `READ` performance hit when the requested file has an exclusive write lock at a write-back-enabled cache other than where the read is issued. The XLD has to

be revoked, and the dirty data must be committed to the origin before the read at the other cache can be serviced.

Tracking dirty data

Write-back from cache to origin happens asynchronously. This means that dirty data isn't immediately written back to the origin. ONTAP employs a dirty data record system to keep track of dirty data per file. Each dirty data record (DDR) represents approximately 20MB of dirty data for a particular file. When a file is actively being written, ONTAP will start flushing dirty data back after two DDRs have been filled and the third DDR is being written. This results in approximately 40MB of dirty data remaining in a cache during writes. For stateful protocols (NFSv4.x, SMB), the remaining 40MB of data will be flushed back to the origin when the file is closed. For stateless protocols (NFSv3), the 40MB of data will be flushed back when either access to the file is requested at a different cache or after the file is idle for two or more minutes, up to a maximum of five minutes. For more information on timer-triggered or space-triggered dirty data flushing, see [Cache scrubbers](#).

In addition to the DDRs and scrubbers, some front-end NAS operations also trigger the flushing of all dirty data for a file:

- SETATTR
 - `SETATTR`'s that modify only mtime, atime, and/or ctime can be processed at the cache, avoiding the penalty of the WAN.
- CLOSE
- OPEN at another cache
- READ at another cache
- REaddir at another cache
- REaddirplus at another cache
- WRITE at another cache

Disconnected mode

When an XLD for a file is held at a write-around cache and that cache gets disconnected from the origin, reads for that file are still allowed at the other caches and origin. This behavior differs when an XLD is held by a write-back-enabled cache. In this case, if the cache is disconnected, reads to the file will hang everywhere. This helps ensure 100% consistency, currency, and coherence are maintained. The reads are allowed in write-around mode because the origin is guaranteed to have all of the data available that has been write-acknowledged to the client. In write-back mode during a disconnect, the origin can not guarantee that all of the data written to and acknowledged by the write-back-enabled cache made it to the origin before the disconnect occurred.

In the event a cache with an XLD for a file is disconnected for an extended period of time, a system administrator can manually revoke the XLD at the origin. This will allow IO to the file to resume at the surviving caches and the origin.



Manually revoking the XLD will result in the loss of any dirty data for the file at the disconnected cache. Manually revoking an XLD should only be done in the event of a catastrophic disruption between the cache and origin.

Cache scrubbers

There are scrubbers in ONTAP that run in response to specific events, such as a timer expiring or space

thresholds being breached. The scrubbers take an exclusive lock on the file being scrubbed, effectively freezing IO to that file until the scrub completes.


Scrubbers include:

- **mtime-based scrubber on the cache:** This scrubber starts every five minutes and scrubs any file sitting unmodified for two minutes. If any dirty data for the file is still in the cache, IO to that file is quiesced and write-back is triggered. IO will resume after the write-back is complete.
- **mtime-based scrubber on origin:** Much like the mtime-based scrubber at the cache, this also runs every five minutes. However, it scrubs any file sitting unmodified for 15 minutes, recalling the inode's delegation. This scrubber doesn't initiate any write-back.
- **RW limit-based scrubber on origin:** ONTAP monitors how many RW lock delegations are handed out per origin constituent. If this number surpasses 170, ONTAP starts scrubbing write lock delegations on a least-recently-used (LRU) basis.
- **Space-based scrubber on the cache:** If a FlexCache volume reaches 90% full, the cache is scrubbed, evicting on an LRU basis.
- **Space-based scrubber on the origin:** If a FlexCache origin volume reaches 90% full, the cache is scrubbed, evicting on an LRU basis.


Sequence diagrams

These sequence diagrams depict the difference in write acknowledgements between write-around and write-back mode.

Write-around



Write-back



ONTAP FlexCache write-back use cases

These are write profiles best suited for a write-back-enabled FlexCache. You should test your workload to see if write-back or write-around provides the best performance.



Write-back is not a replacement for write-around. Although write-back is designed with write-heavy workloads, write-around is still the better choice for many workloads.

Target workloads

File size

File size is less important than the number of writes issued between the `OPEN` and `CLOSE` calls for a file. Small files inherently have fewer `WRITE` calls, making them less ideal for write-back. Large files might have more writes between `OPEN` and `CLOSE` calls, but this isn't guaranteed.

Refer to the [FlexCache write-back guidelines](#) page for the most current recommendations regarding max file size.

Write size

When writing from a client, other modifying NAS calls are involved other than write calls. These include, but are not limited to:

- CREATE
- OPEN
- CLOSE
- SETATTR
- SET_INFO

SETATTR and SET_INFO calls that set `mtime`, `atime`, `ctime`, `owner`, `group`, or `size` are processed at the cache. The rest of these calls must be processed at the origin and trigger a write-back of any dirty data accumulated at the write-back-enabled cache for the file being operated on. IO to the file will be quiesced until the write-back is complete.

Knowing that these calls must traverse the WAN helps you to identify workloads suited for write-back. Generally, the more writes that can be done between OPEN and CLOSE calls without one of the other calls listed above being issued, the better the performance gain write-back provides.

Read-after-write

Read-after-write workloads have historically performed poorly at FlexCache. This is due to the write-around mode of operation before 9.15.1. The WRITE call to the file has to be committed at the origin, and the subsequent READ call would have to pull the data back to the cache. This results in both operations incurring the penalty of the WAN. Therefore, read-after-write workloads are discouraged for FlexCache in write-around mode. With the introduction of write-back in 9.15.1, data is now committed at the cache, and can immediately be read from the cache, eliminating the WAN penalty. If your workload includes read-after-write at FlexCache volumes, you should configure the cache to operate in write-back mode.



If read-after-write is a critical part of your workload, you should configure your cache to operate in write-back mode.


Write-after-write

When a file accumulates dirty data in a cache, the cache asynchronously writes the data back to the origin. This naturally leads to times when the client closes the file with dirty data still waiting to be flushed back to origin. If another open or write comes in for the file that was just closed and still has dirty data, the write will be suspended until all the dirty data has been flushed to origin.

Latency considerations

When FlexCache operates in write-back mode, it becomes more beneficial to NAS clients as latency increases. There is a point, however, at which the overhead of write-back outweighs the advantages gained in low-latency environments. In some NetApp tests, write-back benefits started around a minimum latency between cache and origin of 8ms. This latency varies with workload, so be sure to test to know your workload's point-of-return.

The following graph shows the point-of-return for write-back in NetApp lab tests. The x axis is the file-size, and the y axis is the elapsed time. The test used NFSv3, mounting with an `rsize` and `wsize` of 256KB, and 64ms of WAN latency. This test was performed using a small ONTAP Select instance for both the cache and origin, and a single threaded-write operation. Your results might vary.



Write-back should not be used for intracluster caching. Intracluster caching occurs when the origin and cache are in the same cluster.

ONTAP FlexCache write-back prerequisites

Before you deploy FlexCache in write-back mode, ensure you have met these performance, software, licensing, and system configuration requirements.

CPU and Memory

It is **strongly recommended** that each origin cluster node have at least 128GB of RAM and 20 CPUs to absorb the write-back messages initiated by write-back enabled caches. This is the equivalent of an A400 or greater. If the origin cluster serves as the origin to multiple write-back enabled FlexCaches, it will require more CPU and RAM.



Using an undersized origin for your workload can have profound impacts on performance at the write-back-enabled cache or the origin.

ONTAP version

- The origin **must** be running ONTAP 9.15.1 or later.
- Any caching cluster that needs to operate in write-back mode **must** be running ONTAP 9.15.1 or later.
- Any caching cluster that does not need to operate in write-back mode can run any generally supported ONTAP version.

Licensing

FlexCache, including the write-back mode of operation, is included with your ONTAP purchase. No extra license is required.

Peering

- The origin and cache clusters must be [cluster peered](#)
- The server virtual machines (SVMs) on the origin and cache cluster must be [vserver peered](#) with the FlexCache option.



You do not need to peer a cache cluster to another cache cluster. There is also no need to peer a cache SVM to another cache SVM.

ONTAP FlexCache write-back interoperability

Understand these interoperability considerations when deploying FlexCache in write-back mode.

ONTAP version

To use the write-back mode of operation, both the cache and origin **must** be running ONTAP 9.15.1 or later.



Clusters where a write-back-enabled cache is unnecessary can run earlier versions of ONTAP, but that cluster can only operate in write-around mode.

You can have a mix of ONTAP versions in your environment.

Table 1. Mixed cluster versions example 1

Cluster	ONTAP version	Write-back supported?
Origin	ONTAP 9.15.1	N/A †
Cluster 1	ONTAP 9.15.1	Yes
Cluster 2	ONTAP 9.14.1	No

Table 2. Mixed cluster versions example 2

Cluster	ONTAP version	Write-back supported?
Origin	ONTAP 9.14.1	N/A †
Cluster 1	ONTAP 9.15.1	No
Cluster 2	ONTAP 9.15.1	No

† *Origins aren't a cache, so neither write-back nor write-around support is applicable.*



In [Mixed cluster versions example 2](#), neither cluster can enable write-back mode because the origin is not running ONTAP 9.15.1 or later, which is a strict requirement.

Client interoperability

Any client generally supported by ONTAP can access a FlexCache volume regardless of whether it is operating in write-around or write-back mode. For an up-to-date list of supported clients, refer to NetApp's [interoperability matrix](#).

Although the client version doesn't matter specifically, the client must be new enough to support NFSv3, NFSv4.0, NFSv4.1, SMB2.x, or SMB3.x. SMB1 and NFSv2 are deprecated protocols and are not supported.

Write-back and write-around

As seen in [Mixed cluster versions example 1](#), FlexCache operating in write-back mode can co-exist with caches operating in write-around mode. It is advised to compare write-around against write-back with your specific workload.



If the performance for a workload is the same between write-back and write-around, use write-around.

ONTAP feature interoperability

For the most up-to-date list of FlexCache feature interoperability, refer to [the supported and unsupported features for FlexCache volumes](#).

Enable and manage ONTAP FlexCache write-back

Beginning with ONTAP 9.15.1, you can enable FlexCache write-back mode on FlexCache volumes to provide better performance for edge computing environments and caches with write-heavy workloads. You can also determine whether write-back is enabled on a FlexCache volume or disable write-back on the volume when necessary.

When write-back is enabled on the cache volume, write requests are sent to the local cache rather than to the origin volume.

Before you begin

You must be in advanced privilege mode.

Create a new FlexCache volume with write-back enabled

Steps

You can create a new FlexCache volume with write-back enabled by using ONTAP System Manager or the ONTAP CLI.

System Manager

1. If the FlexCache volume is on a different cluster than the origin volume, create a cluster peer relationship:
 - a. On the local cluster, click **Protection > Overview**.
 - b. Expand **Intercluster Settings**, click **Add Network Interfaces**, and add intercluster interfaces to the cluster.

Repeat this on the remote cluster.

 - c. On the remote cluster, click **Protection > Overview**. Click **⋮** in the Cluster Peers section and click **Generate Passphrase**.
 - d. Copy the generated passphrase and paste it in the local cluster.
 - e. On the local cluster, under Cluster Peers, click **Peer Clusters** and peer the local and remote clusters.
2. If the FlexCache volume is on a different cluster than the origin volume, create an SVM peer relationship:

Under **Storage VM Peers**, click **⋮** and then **Peer Storage VMs** to peer the storage VMs.

If the FlexCache volume is on the same cluster, you cannot create an SVM peer relationship using System Manager.

3. Select **Storage > Volumes**.
4. Select **Add**.
5. Select **More Options** and then select **Add as cache for a remote volume**.
6. Select **Enable FlexCache write-back**.

CLI

1. If the FlexCache volume to be created is in a different cluster, create a cluster peer relationship:
 - a. On the destination cluster, create a peer relationship with the data protection source cluster:

```
cluster peer create -generate-passphrase -offer-expiration
MM/DD/YYYY HH:MM:SS|1...7days|1...168hours -peer-addrs
<peer_LIF_IPs> -initial-allowed-vserver-peers <svm_name>,...|*
-ipspace <ipspace_name>
```

Beginning with ONTAP 9.6, TLS encryption is enabled by default when creating a cluster peer relationship. TLS encryption is supported for the intercluster communication between the origin and FlexCache volumes. You can also disable TLS encryption for the cluster peer relationship, if required.

```
cluster02::> cluster peer create -generate-passphrase -offer
-expiration 2days -initial-allowed-vserver-peers *

          Passphrase: UCa+6lRVICXeL/gq1WrK7ShR
          Expiration Time: 6/7/2017 08:16:10 EST
          Initial Allowed Vserver Peers: *
          Intercluster LIF IP: 192.140.112.101
          Peer Cluster Name: Clus_7ShR (temporary generated)
```

Warning: make a note of the passphrase - it cannot be displayed again.

b. On the source cluster, authenticate the source cluster to the destination cluster:

```
cluster peer create -peer-addrs <peer_LIF_IPs> -ipspace <ipspace>
```

```
cluster01::> cluster peer create -peer-addrs
192.140.112.101,192.140.112.102
```

Notice: Use a generated passphrase or choose a passphrase of 8 or more characters.

To ensure the authenticity of the peering relationship, use a phrase or sequence of characters that would be hard to guess.

Enter the passphrase:

Confirm the passphrase:

Clusters cluster02 and cluster01 are peered.

2. If the FlexCache volume is in a different SVM than that of the origin volume, create an SVM peer relationship with `flexcache` as the application:

a. If the SVM is in a different cluster, create an SVM permission for the peering SVMs:

```
vserver peer permission create -peer-cluster <cluster_name>
-vserver <svm-name> -applications flexcache
```

The following example illustrates how to create an SVM peer permission that applies for all of the local SVMs:

```
cluster1::> vserver peer permission create -peer-cluster cluster2
-vserver "*" -applications flexcache

Warning: This Vserver peer permission applies to all local
Vservers. After that no explicit
"vserver peer accept" command required for Vserver peer
relationship creation request
from peer cluster "cluster2" with any of the local Vservers. Do
you want to continue? {y|n}: y
```

b. Create the SVM peer relationship:

```
vserver peer create -vserver <local_SVM> -peer-vserver
<remote_SVM> -peer-cluster <cluster_name> -applications flexcache
```

3. Create a FlexCache volume with write-back enabled:

```
volume flexcache create -vserver <cache_vserver_name> -volume
<cache_flexgroup_name> -aggr-list <list_of_aggregates> -origin
-volume <origin_flexgroup> -origin-vserver <origin_vserver_name>
-junction-path <junction_path> -is-writeback-enabled true
```

Enable FlexCache write-back on an existing FlexCache volume

You can enable FlexCache write-back on an existing FlexCache volume using ONTAP System Manager or the ONTAP CLI.

System Manager

1. Select **Storage > Volumes** and select an existing FlexCache volume.
2. On the volume's Overview page, click **Edit** in the upper right corner.
3. In the **Edit Volume** window, select **Enable FlexCache write-back**.

CLI

1. Enable write-back on an existing FlexCache volume:

```
volume flexcache config modify -volume <cache_flexgroup_name> -is
-writeback-enabled true
```

Check if FlexCache write-back is enabled

Steps

You can use System Manager or the ONTAP CLI to determine whether FlexCache write-back is enabled.

System Manager

1. Select **Storage > Volumes** and select a volume.
2. In the volume **Overview**, locate **FlexCache details** and check if FlexCache write-back is set to **Enabled** on the FlexCache volume.

CLI

1. Check if FlexCache write-back is enabled:

```
volume flexcache config show -volume <cache_flexgroup_name> -fields  
is-writeback-enabled
```

Disable write-back on a FlexCache volume

Before you can delete a FlexCache volume you need to disable FlexCache write-back.

Steps

You can use System Manager or the ONTAP CLI to disable FlexCache write-back.

System Manager

1. Select **Storage > Volumes** and select an existing FlexCache volume that has FlexCache write-back enabled.
2. On the volume's Overview page, click **Edit** in the upper right corner.
3. In the **Edit Volume** window, deselect **Enable FlexCache write-back**.

CLI

1. Disable write-back:

```
volume flexcache config modify -volume <cache_vol_name> -is  
-writeback-enabled false
```

Frequently asked questions about ONTAP FlexCache write-back

This FAQ can help if you are looking for a quick answer to a question.

I want to use write-back. What version of ONTAP do I need to run?

Both the cache and the origin must be running ONTAP 9.15.1 or later. It is **strongly** recommended that you run the latest P release. Engineering is constantly improving the performance and functionality of write-back-enabled caches.

Can clients accessing the origin have an effect on clients accessing the write-back-enabled cache?

Yes. The origin has equal right to the data as any of the caches. If an operation is executed on a file that requires the file to be evicted from the cache, or a lock/data delegation to be revoked, the client at the cache might see a delay accessing the file.

Can I apply QoS to write-back-enabled FlexCaches?

Yes. Every cache and the origin can have independent QoS policies applied. This will have no direct effect on any write-back initiated intercluster traffic. Indirectly, you can slow down intercluster write-back traffic by QoS limiting the front-end traffic at the write-back-enabled cache.

Is multi-protocol NAS supported at write-back-enabled FlexCaches?

Yes. Multi-protocol is fully supported at write-back-enabled FlexCaches. Currently, NFSv4.2 and S3 are not supported by FlexCache operating in write-around or write-back mode.

Are SMB alternate data streams supported at write-back-enabled FlexCaches?

SMB alternate data streams (ADS) are supported, but not accelerated by write-back. The write to the ADS is forwarded to the origin, incurring the penalty of the WAN latency. The write also evicts the main file the ADS is a part of from the cache.

Can I switch a cache between write-around and write-back mode after it is created?

Yes. All you have to do is toggle the `is-writeback-enabled` flag in the `flexcache modify` [command](#).

Are there bandwidth considerations I should be aware of for the intercluster link between the cache(s) and origin?

Yes. FlexCache write-back is highly dependent on the intercluster link between the cache(s) and origin. Low bandwidth and/or lossy networks can have a significant negative effect on performance. There isn't a specific bandwidth requirement, as it is highly dependent on your workload.

FlexCache duality

FAQ about FlexCache duality

This FAQ answers common questions about FlexCache duality introduced in ONTAP 9.18.1.

Frequently asked questions

What is "duality"?

Duality enables unified access to the same data using both file (NAS) and object (S3) protocols. Introduced in ONTAP 9.12.1 without FlexCache support, duality was extended in ONTAP 9.18.1 to include FlexCache volumes, allowing S3 protocol access to NAS files cached in a FlexCache volume.

What S3 operations are supported on a FlexCache S3 bucket?

S3 operations supported on standard S3 NAS buckets are supported on FlexCache S3 NAS buckets, with the exception of the `COPY` operation. For an up-to-date list of unsupported operations for a standard S3 NAS bucket, visit the [interoperability documentation](#).

Can I use FlexCache in write-back mode with FlexCache duality?

No. If a FlexCache S3 NAS bucket is created on a FlexCache volume, the FlexCache volume **must** be in write-around mode. If you attempt to create a FlexCache S3 NAS bucket on a FlexCache volume in write-back mode, the operation will fail.

I can't upgrade one of my clusters to ONTAP 9.18.1 because of hardware limitations. Will duality still work in my cluster if only the cache cluster is running ONTAP 9.18.1?

No. Both the cache cluster and origin cluster must have a minimum effective cluster version of 9.18.1. If you attempt to create a FlexCache S3 NAS bucket on a cache cluster peered with an origin running an ONTAP version earlier than 9.18.1, the operation will fail.

I have a MetroCluster configuration. Can I use FlexCache duality?

No. FlexCache duality is not supported in MetroCluster configurations.

Can I audit S3 access to files in a FlexCache S3 NAS bucket?

S3 auditing is provided by the NAS auditing functionality FlexCache volumes use. For more information about NAS auditing of FlexCache volumes, see [Learn more about FlexCache auditing](#).

What should I expect if the cache cluster becomes disconnected from the origin cluster?

S3 requests to a FlexCache S3 NAS bucket will fail with a 503 Service Unavailable error if the cache cluster is disconnected from the origin cluster.

Can I use multipart S3 operations with FlexCache duality?

For multipart S3 operations to work, the underlying FlexCache volume must have the granular-data field set to 'advanced'. This field is set to whatever value is set for the origin volume.

Does FlexCache duality support HTTP and HTTPS access?

Yes. By default, HTTPS is required. You can configure the S3 service to allow HTTP access if needed.

Enable S3 access to NAS FlexCache volumes

Beginning in ONTAP 9.18.1, you can enable S3 access to NAS FlexCache volumes, also referred to as "duality." This allows clients to access data stored in a FlexCache volume using the S3 protocol, in addition to traditional NAS protocols like NFS and SMB. You can use the following information to set up FlexCache duality.

Prerequisites

Before you begin, you must ensure you complete the following prerequisites:

- Make sure the S3 protocol and desired NAS protocols (NFS, SMB, or both) are licensed and configured on the SVM.
- Verify that DNS and any other required services are configured.
- Cluster and SVM Peered
- FlexCache Volume create
- Data-lif created



For more thorough documentation on FlexCache duality, see [ONTAP S3 multiprotocol support](#).

Step 1: Create and sign certificates

To enable S3 access to a FlexCache volume, you need to install certificates for the SVM that hosts the FlexCache volume. This example uses self-signed certificates, but in a production environment, you should use certificates signed by a trusted Certificate Authority (CA).

1. Create an SVM root CA:

```
security certificate create -vserver <svm> -type root-ca -common-name <arbitrary_name>
```

2. Generate a Certificate Signing Request:

```
security certificate generate-csr -common-name <dns_name_of_data_lif> -dns-name <dns_name_of_data_lif> -ipaddr <data_lif_ip>
```

Example output:

```
-----BEGIN CERTIFICATE REQUEST-----
MIICzjCCAbYCAQAwHzEdMBsGA1UEAxMUY2FjaGUxZy1kYXRhLm5hcy5sYWIwggEi
MA0GCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQCuJk07508Uh329cHI6x+BaRS2
w5wrqvzoYlidXtYmdCH3m1DDprBiAyfIwBC0/iU3Xd5NpB7nc1wK1CI2VEkrXGUg
...
vMIGN351+FgzLQ4X51KfoMXCV70NqIakxzEmkTIUDKv7n9EVZ4b5DTT1rL03X/nK
+Bim2y2y180PaFB3NauZHTnIIzIc8zCp2IEqmFWyMDcdBjP9KS0+jNm4QhuXiM8F
D7gm3g/O70qa5OxbAEa15o4Nb0195U0T0rwqTaSzFG0XQnK2PmA1OIwS5ET35p3Z
dLU=
-----END CERTIFICATE REQUEST-----
```

Private key example:

```
-----BEGIN PRIVATE KEY-----
MIIEvAIBADANBgkqhkiG9w0BAQEFAASCBKYwggSiAgEAAoIBAQCuJk07508Uh32
9cHI6x+BaRS2w5wrqvzoYlidXtYmdCH3m1DDprBiAyfIwBC0/iU3Xd5NpB7nc1wK
1CI2VEkrXGUgwBtx1K4I1rCTB829Q1aLGAQXVyWnzhQc4tS5PW/DsQ8t7o1Z9zEI
...
rXGEDDaqp7jQGNXUGlbxO3zcB1l1/A9Hc6oalNECgYBKwE3PeZamiwhIHLy9ph7w
dJffCshsPalMuAp2OuKIANa916fT9y5kf9tIbskT+t5Dth8bmV9pwe8UZaK5eC4
Svxm19jHT5QqloDaZVUmMXFKyKoqPDdfvcDk2Eb5gMfIIb0a3TPC/jqqpDn9BzuH
TO02fuRvRR/G/HUz2yRd+A==
-----END PRIVATE KEY-----
```



Keep a copy of your certificate request and private key for future reference.

3. Sign the certificate:

The `root-ca` is the one you created in [Create an SVM root CA](#).

```
certificate sign -ca <svm_root_ca> -ca-serial <svm_root_ca_sn> -expire  
-days 364 -format PEM -vserver <svm>
```

4. Paste the Certificate Signing Request (CSR) generated in [Generate a Certificate Signing Request](#).

Example:

```
-----BEGIN CERTIFICATE REQUEST-----  
MIICzjCCAbYCAQAwHzEdMBsGA1UEAxMUY2FjaGUxZy1kYXRhLm5hcy5sYWIwggEi  
MA0GCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQCuJk07508Uh329cHI6x+BaRS2  
w5wrqvzoYlidXtYmdCH3m1DDprBiAyfIwBC0/iU3Xd5NpB7nc1wK1CI2VEkrXGUg  
...  
vMIGN351+FgzLQ4X51KfoMXCV70NqIakxzEmkTIUDKv7n9EVZ4b5DTT1rL03X/nK  
+Bim2y2y180PaFB3NauZHTnIIzIc8zCp2IEqmFWyMDcdBjP9KS0+jNm4QhuXiM8F  
D7gm3g/O70qa5OxbAEa15o4Nb0195U0T0rwqTaSzFG0XQnK2PmA1OIwS5ET35p3Z  
dLU=  
-----END CERTIFICATE REQUEST-----
```

This prints a signed certificate to the console, similar to the following example.

Signed Certificate example:

```
-----BEGIN CERTIFICATE-----  
MIIDdzCCAl+gAwIBAgIIIGHolbgv5DPowDQYJKoZIhvcNAQELBQAwLjEfMB0GA1UE  
AxMWY2FjaGUtMTY0Zy1zdm0tcm9vdC1jYTELMAkGA1UEBhMCVVMwHhcNMjUxMTIx  
MjIxNTU4WhcNMjYxMTIxMjIxNTU4WjAfMR0wGwYDVQQDExRjYWNoZTFnLWRhdGEu  
...  
qS7zhj3ikWE3Gp9s+QijKWxx/0HDd1UuGqy0QZNqNm/M0mqVnokJNk5F4fBFxMiR  
1o63BxL8xGIRdtTCjjb2Gq2Wj7EC1Uw6CykEkxAcVk+XrRtArGkNtcYdtHfUsKVE  
wswvv0rNydrNnWhJLhS18TW5Tex+OMyTXgk9/3K8kB0mAMrtxxYjt8tm+gztkivf  
J0eo1uDJhaNxqwEZRzFyGaa4k1+56oFzRfTc  
-----END CERTIFICATE-----
```

5. Copy the certificate for the next step.
6. Install the server certificate on the SVM:

```
certificate install -type server -vserver <svm> -cert-name flexcache-  
duality
```

7. Paste the signed certificate from [Sign the certificate](#).

Example:

```

Please enter Certificate: Press <Enter> [twice] when done
-----BEGIN CERTIFICATE-----
MIIDdzCCA1+gAwIBAgIIGHolbgv5DPowDQYJKoZIhvcNAQELBQAwljEfMB0GA1UE
AxMWY2FjaGUtMTY0Zy1zdm0tcm9vdC1jYTELMAkGA1UEBhMCVVMwHhcNMjUxMTIx
MjIxNTU4WhcNMjYxMTIxMjIxNTU4WjAfMR0wGwYDVQQDExRjYWN0ZTFnLWRhdGEu
bmFzLmxhYjCCASIwDQYJKoZIhvcNAQEBBQADggEPADCCAQoCggEBAK6wmTTvk7xS
...
qS7zhj3ikWE3Gp9s+QijKWxx/0HDd1UuGqy0QZNqNm/M0mqVnokJNk5F4fBFxMiR
1o63BxL8xGIRdtTCjbjb2Gq2Wj7EC1Uw6CykEkxAcVx+XrRtArGkNtcYdtHfUsKVE
wswwv0rNydrNnWhJLhS18TW5Tex+OMyTXgk9/3K8kB0mAMrtxxYjt8tm+gztkivf
J0eo1uDJhaNxqwEZRzFyGaa4k1+56oFzRfTc
-----END CERTIFICATE-----

```

8. Paste the private key generated in [Generate a Certificate Signing Request](#).

Example:

```

Please enter Private Key: Press <Enter> [twice] when done
-----BEGIN PRIVATE KEY-----
MIIEvAIBADANBgkqhkiG9w0BAQEFAASCBKYwggSiAgEAAoIBAQCuJk07508Uh32
9cHI6x+BaRS2w5wrqvzoYlidXtYmdCH3m1DDprBiAyfIwBC0/iU3Xd5NpB7nc1wK
1CI2VEkrXGUgwBtx1K4I1rCTB829Q1aLGAQXVwNzhQc4tS5PW/DsQ8t7o1Z9zEI
W/gaEIajgpXIwGNWZ+weKQK+yoolxC+gy4IUE7WvnEUiezaIdoqzyPhYq5GC4XWf
0johpQugOPe0/w2nVFRWJoFQp3ZP3NZAXc8H0qkRB6SjaM243XV2jnuEzX2joXvT
whHH+IBAQ2JDs7s1TY0I20e49J2Fx2+HvUxDx4BHao7CCHA1+MnmEl+9E38wTaEk
NLsU724ZAgMBAAECggEABHUY06wxcIk5h03S9Ik1FDZV3JWzsu5gGdLSQOHd5W+
...
rXGEDDaqp7jQGNXUG1bxO3zcB11/A9Hc6oalNECgYBKwe3PeZamiwhIHLy9ph7w
dJffCshsPalMuAp2OuKIAAnNa916ft9y5kf9tIbskT+t5Dth8bmV9pwe8UZaK5eC4
Svxm19jHT5Qql0DaZVUmMXFKyKoqPDdfvcDk2Eb5gMfIIb0a3TPC/jqqpDn9BzuH
TO02fuRvRR/G/HUz2yRd+A==
-----END PRIVATE KEY-----

```

9. Enter certificates of certification authorities (CA) which form the certificate chain of the server certificate.

This starts with the issuing CA certificate of the server certificate and can range up to the root CA certificate.

```
Do you want to continue entering root and/or intermediate certificates
{y|n}: n
```

You should keep a copy of the private key and the CA-signed digital certificate for future reference.

The installed certificate's CA and serial number for reference:

```
CA: cache-164g-svm-root-ca
serial: 187A256E0BF90CFA
```

10. Get the public key for the SVM root CA:

```
security certificate show -vserver <svm> -common-name <root_ca_cn> -ca
<root_ca_cn> -type root-ca -instance

-----BEGIN CERTIFICATE-----
MIIDgTCCAmgAwIBAgIIGHokTnbsHKEwDQYJKoZIhvcNAQELBQAwLjEfMB0GA1UE
AxMWY2FjaGUtMTY0Zy1zdm0tcm9vdC1jYTELMAkGA1UEBhMCVVMwHhcNMjUxMTIx
MjE1NTIzWhcNMjYxMTIxMjE1NTIzWjAuMR8wHQYDVQQDExZjYWN0ZS0xNjRnLXN2
bs1yb290LWNhMQswCQYDVQQGEwJVUzCCASIwDQYJKoZIhvcNAQEBBQADggEPADCC
...
DoOL7vZFFt44xd+rp0DwafhSnLH5HNhdIAfa2JvZW+eJ7rgevH9wmOzyc1vaih13
Ewtb6cz1a/mtESSYRNBMGkIGM/SFCy5v1ROZXCzF96XPbYQN4cW0AYI3AHYBZP0A
H1NzDR8iml4k9IuKf6BHLFA+VwLTJJZKrdf5Jvjgh0trGAbQGI/Hp2Bjuiopkui+
n4aa5Rz0JFQopqQddAYnMuvcq10CyNn7S0vF/XLd3fJaprH8kQ==
-----END CERTIFICATE-----
```



This is needed to configure the client to trust the certificates signed by the SVM root-ca. The public key is printed to the console. Copy and save the public key. The values in this command are the same ones you entered in [Create an SVM root CA](#).

Step 2: Configure the S3 server

1. Enable S3 protocol access:

```
vserver show -vserver <svm> -fields allowed-protocols
```



S3 is allowed at the SVM level by default.

2. Clone an existing policy:

```
network interface service-policy clone -vserver <svm> -policy default-data-files -target-vserver <svm> -target-policy <any_name>
```

3. Add S3 to the cloned policy:

```
network interface service-policy add-service -vserver <svm> -policy <any_name> -service data-s3-server
```

4. Add the new policy to the data lif:

```
network interface modify -vserver <svm> -lif <data_lif> -service-policy duality
```



Modifying the service policy of an existing LIF can be disruptive. It requires the LIF to be taken down and brought back up with a listener for the new service. TCP **should** recover from this quickly, but be aware of potential impact.

5. Create the S3 object store server on the SVM:

```
vserver object-store-server create -vserver <svm> -object-store-server <dns_name_of_data_lif> -certificate-name flexcache-duality
```

6. Enable S3 capability on FlexCache volume:

The `flexcache config` option `-is-s3-enabled` must be set to `true` before you can create a bucket. You must also set the option `-is-writeback-enabled` to `false`.

The following command modifies an existing FlexCache:

```
flexcache config modify -vserver <svm> -volume <fcache_vol> -is-writeback-enabled false -is-s3-enabled true
```

7. Create an S3 bucket:

```
vserver object-store-server bucket create -vserver <svm> -bucket <bucket_name> -type nas -nas-path <flexcache_junction_path>
```

8. Create a bucket policy:

```
vserver object-store-server bucket policy add-statement -vserver <svm>  
-bucket <bucket_name> -effect allow
```

9. Create an S3 user:

```
vserver object-store-server user create -user <user> -comment ""
```

Example output:

```
Vserver: <svm>>  
User: <user>>  
Access Key: WCOT7...Y7D6U  
Secret Key: 6143s...pd__P  
Warning: The secret key won't be displayed again. Save this key for  
future use.
```

10. Regenerate keys for the root user:

```
vserver object-store-server user regenerate-keys -vserver <svm> -user  
root
```

Example output:

```
Vserver: <svm>>  
User: root  
Access Key: US791...2F1RB  
Secret Key: tgYmn...8_3o2  
Warning: The secret key won't be displayed again. Save this key for  
future use.
```

Step 3: Set up the client

There are many S3 clients available. A good place to start is with the AWS CLI. For more information, see [Installing the AWS CLI](#).

Manage FlexCache volumes

Learn about auditing ONTAP FlexCache volumes

Beginning with ONTAP 9.7, you can audit NFS file access events in FlexCache relationships using native ONTAP auditing and file policy management with FPolicy.

Beginning with ONTAP 9.14.1, FPolicy is supported for FlexCache volumes with NFS or SMB. Previously, FPolicy was not supported for FlexCache volumes with SMB.

Native auditing and FPolicy are configured and managed with the same CLI commands used for FlexVol volumes. However, there is some different behavior with FlexCache volumes.

- **Native auditing**

- You can't use a FlexCache volume as the destination for audit logs.
- If you want to audit read and writes on FlexCache volumes, you must configure auditing on both the cache SVM as well as on the origin SVM.

This is because file system operations are audited where they are processed. That is, reads are audited on the cache SVM and writes are audited on the origin SVM.

- To track the origin of write operations, the SVM UUID and MSID are appended in the audit log to identify the FlexCache volume from which the write originated.

- **FPolicy**

- Although writes to a FlexCache volume are committed on the origin volume, FPolicy configurations monitor the writes on the cache volume. This is unlike native auditing, in which the writes are audited on the origin volume.
- While ONTAP does not require the same FPolicy configuration on cache and origin SVMs, it is recommended that you deploy two similar configurations. You can do so by creating a new FPolicy policy for the cache, configured like that of the origin SVM but with the scope of the new policy limited to the cache SVM.
- The size of extensions in an FPolicy configuration is limited to 20KB (20480 bytes). When the size of extensions used in an FPolicy configuration on a FlexCache volume exceeds 20KB, the EMS message `nblade.fpolicy.extn.failed` is triggered.

Synchronize properties of an ONTAP FlexCache volume from an origin volume

Some of the volume properties of the FlexCache volume must always be synchronized with those of the origin volume. If the volume properties of a FlexCache volume fail to synchronize automatically after the properties are modified at the origin volume, you can manually synchronize the properties.

About this task

The following volume properties of a FlexCache volume must always be synchronized with those of the origin volume:

- Security style (`-security-style`)
- Volume name (`-volume-name`)
- Maximum directory size (`-maxdir-size`)
- Minimum read ahead (`-min-readahead`)

Step

1. From the FlexCache volume, synchronize the volume properties:

```
volume flexcache sync-properties -vserver svm_name -volume flexcache_volume
```

```
cluster1::> volume flexcache sync-properties -vserver vs1 -volume fc1
```

Update the configuration of ONTAP FlexCache relationships

After events such as volume move, aggregate relocation, or storage failover, the volume configuration information on the origin volume and FlexCache volume is updated automatically. In case the automatic updates fail, an EMS message is generated and then you must manually update the configuration for the FlexCache relationship.

If the origin volume and the FlexCache volume are in the disconnected mode, you might need to perform some additional operations to update a FlexCache relationship manually.

About this task

If you want to update the configurations of a FlexCache volume, you must run the command from the origin volume. If you want to update the configurations of an origin volume, you must run the command from the FlexCache volume.

Step

1. Update the configuration of the FlexCache relationship:

```
volume flexcache config-refresh -peer-vserver peer_svm -peer-volume
peer_volume_to_update -peer-endpoint-type [origin | cache]
```

Enable file access time updates on the ONTAP FlexCache volume

Beginning with ONTAP 9.11.1, you can enable the `-atime-update` field on the FlexCache volume to permit file access time updates. You can also set an access time update period with the `-atime-update-period` attribute. The `-atime-update-period` attribute controls how often access time updates can take place and when they can propagate to the origin volume.

Overview

ONTAP provides a volume-level field called `-atime-update`, to manage access time updates on files and directories that are read using READ, READLINK, and REaddir. Atime is used for data lifecycle decisions for files and directories that are infrequently accessed. The infrequently accessed files are eventually migrated to archive storage and are often later moved to tape.

The `atime-update` field is disabled by default on existing and newly created FlexCache volumes. If you are using FlexCache volumes with ONTAP releases earlier than 9.11.1, you should leave the `atime-update` field disabled so caches aren't unnecessarily evicted when a read operation is performed on the origin volume. With large FlexCache caches, however, administrators use special tools to manage data and help to ensure that hot data remains in the cache and cold data is purged. This is not possible when `atime-update` is disabled.

However, beginning with ONTAP 9.11.1, you can enable `-atime-update` and `-atime-update-period`, and use the tools required to manage the cached data.

Before you begin

- All FlexCache volumes must be running ONTAP 9.11.1 or later.
- You must use the advanced privilege mode.

About this task

Setting `-atime-update-period` to 86400 seconds allows no more than one access time update per 24-hour period, regardless of the number of read-like operations performed on a file.

Setting the `-atime-update-period` to 0 sends messages to the origin for each read access. The origin then informs each FlexCache volume that the atime is outdated, which impacts performance.

Steps

1. Set the privilege mode to advanced:

```
set -privilege advanced
```

2. Enable file access time updates and set the update frequency:

```
volume modify -volume vol_name -vserver <SVM name> -atime-update true -atime-update-period <seconds>
```

The following example enables `-atime-update` and sets `-atime-update-period` to 86400 seconds, or 24 hours:

```
c1: volume modify -volume origin1 vs1_c1 -atime-update true -atime-update-period 86400
```

3. Verify that `-atime-update` is enabled:

```
volume show -volume vol_name -fields atime-update,atime-update-period
```

```
c1::>*> volume show -volume cache1_origin1 -fields atime-update,atime-update-period
vserver volume          atime-update atime-update-period
-----
vs2_c1    cache1_origin1  true          86400
```

4. After `-atime-update` is enabled, you can specify if the files on a FlexCache volume can be scrubbed automatically and a scrubbing interval:

```
volume flexcache config modify -vserver <SVM name> -volume <volume_name> -is-atime-scrub-enabled <true|false> -atime-scrub-period <integer>
```

Learn more about `-is-atime-scrub-enabled` parameter in the [ONTAP command reference](#).

Enable global file locking on ONTAP FlexCache volumes

Beginning with ONTAP 9.10.1, global file locking can be applied to prevent reads across all related cached files.

With global file locking enabled, modifications to the origin volume are suspended until all FlexCache volumes are online. You should only enable global file locking when you have control over the reliability of the connections between cache and origin due to suspension and possible timeouts of modifications when FlexCache volumes are offline.

Before you begin

- Global file locking requires the clusters containing the origin and all associated caches to be running ONTAP 9.9.1 or later. Global file locking can be enabled on new or existing FlexCache volumes. The command can be run on one volume and applies to all associated FlexCache volumes.
- You must be in the advanced privilege level to enable global file locking.
- If you revert to a version of ONTAP earlier than 9.9.1, global file locking must first be disabled on the origin and associated caches. To disable, from the origin volume, run: `volume flexcache prepare-to-downgrade -disable-feature-set 9.10.0`
- The process to enable global file locking depends on whether the origin has existing caches:
 - [Enable global file locking on new FlexCache volumes](#)
 - [Enable global file locking on existing FlexCache volumes](#)

Enable global file locking on new FlexCache volumes

Steps

1. Create the FlexCache volume with `-is-global-file-locking` set to true:

```
volume flexcache create volume volume_name -is-global-file-locking-enabled true
```



The default value of `-is-global-file-locking` is “false”. When any subsequent `volume flexcache create` commands are run on a volume, they must be passed with `-is-global-file-locking` enabled set to “true”.

Enable global file locking on existing FlexCache volumes

Steps

1. Global file locking must be set from the origin volume.
2. The origin cannot have any other existing relationships (for example, SnapMirror). Any existing relationships must be dissociated. All caches and volumes must be connected at the time of running the command. To check the connection status, run:

```
volume flexcache connection-status show
```

The status for all the listed volumes should display as connected. For more information, see [View the status of a FlexCache relationship](#) or [Synchronize properties of a FlexCache volume from an origin](#).

3. Enable global file locking on the caches:

```
volume flexcache origin config show/modify -volume volume_name -is-global-file  
-locking-enabled true
```

Related information

- [ONTAP command reference](#)

Prepopulate ONTAP FlexCache volumes

You can prepopulate a FlexCache volume to reduce the time it takes to access cached data.

Before you begin

- You must be a cluster administrator at the advanced privilege level
- The paths you pass for prepopulation must exist or the prepopulate operation fails.

About this task

- Prepopulate reads files only and crawls through directories
- The `-isRecursion` flag applies to the entire list of directories passed to prepopulate

Steps

1. Prepopulate a FlexCache volume:

```
volume flexcache prepopulate -cache-vserver vserver_name -cache-volume -path  
-list path_list -isRecursion true|false
```

- The `-path-list` parameter indicates the relative directory path you want to prepopulate starting from the origin root directory. For example, if the origin root directory is named `/origin` and it contains directories `/origin/dir1` and `/origin/dir2`, you can specify the path list as follows: `-path-list dir1, dir2` or `-path-list /dir1, /dir2`.
- The default value of the `-isRecursion` parameter is True.

This example prepopulates a single directory path:

```
cluster1::>*> flexcache prepopulate start -cache-vserver vs2 -cache  
-volume fg_cachevol_1 -path-list /dir1  
  (volume flexcache prepopulate start)  
[JobId 207]: FlexCache prepopulate job queued.
```

This example prepopulates files from several directories:

```
cluster1::>*> flexcache prepopulate start -cache-vserver vs2 -cache  
-volume fg_cachevol_1 -path-list /dir1,/dir2,/dir3,/dir4  
  (volume flexcache prepopulate start)  
[JobId 208]: FlexCache prepopulate job queued.
```

This example prepopulates a single file:

```
cluster1::*> flexcache prepopulate start -cache-vserver vs2 -cache
-volume fg_cachevol_1 -path-list /dir1/file1.txt
  (volume flexcache prepopulate start)
[JobId 209]: FlexCache prepopulate job queued.
```

This example prepopulates all files from the origin:

```
cluster1::*> flexcache prepopulate start -cache-vserver vs2 -cache
-volume fg_cachevol_1 -path-list / -isRecursion true
  (volume flexcache prepopulate start)
[JobId 210]: FlexCache prepopulate job queued.
```

This example includes an invalid path for prepopulation:

```
cluster1::*> flexcache prepopulate start -cache-volume
vol_cache2_vs3_c2_vol_origin1_vs1_c1 -cache-vserver vs3_c2 -path-list
kdir1, dir5, dir6
  (volume flexcache prepopulate start)

Error: command failed: Path(s) "dir5, dir6" does not exist in origin
volume
  "vol_origin1_vs1_c1" in Vserver "vs1_c1".
```

2. Display the number of files read:

```
job show -id job_ID -ins
```

Related information

- [job show](#)

Delete ONTAP FlexCache relationships

You can delete a FlexCache relationship and the FlexCache volume if you no longer require the FlexCache volume.

Steps

1. From the cluster that has the FlexCache volume, take the FlexCache volume offline:

```
volume offline -vserver svm_name -volume volume_name
```

2. Delete the FlexCache volume:

```
volume flexcache delete -vserver svm_name -volume volume_name
```

The FlexCache relationship details are removed from the origin volume and the FlexCache volume.

FlexCache for hotspot remediation

Remediating hotspotting in high-performance compute workloads with ONTAP FlexCache volumes

A common problem with many high-performance compute workloads, such as animation rendering or EDA, is hotspotting. Hotspotting is a situation that occurs when a specific part of the cluster or network experiences a significantly higher load compared to other areas, leading to performance bottlenecks and reduced overall efficiency due to excessive data traffic concentrated in that location. For example, a file, or multiple files, is in high demand for the job running which results in a bottleneck at the CPU used to service requests (via a volume affinity) to that file. FlexCache can help alleviate this bottleneck, but it must be set up properly.

This documentation explains how to set up FlexCache to remediate hotspotting.



Beginning July 2024, content from technical reports previously published as PDFs has been integrated with ONTAP product documentation. This ONTAP hotspot remediation technical report content is net new as of the date of its publication and no earlier format was ever produced.

Key concepts

When planning hotspot remediation, it's important to understand these essential concepts.

- **High-density FlexCache (HDF):** A FlexCache that is condensed to span as few nodes as the cache capacity requirements allow
- **HDF Array (HDFA):** A group of HDFs that are caches of the same origin, distributed across the cluster
- **Inter-SVM HDFA:** One HDF from the HDFA per server virtual machine (SVM)
- **Intra-SVM HDFA:** All HDFs in the HDFA in one SVM
- **East-west traffic:** Cluster backend traffic generated from indirect data access

What's next

- Understand how to architect with high-density FlexCache to help remediate hotspotting
- Decide on FlexCache array density
- Determine the density of your HDFs and decide whether you will be accessing the HDFs using NFS with inter-SVM HDFA and intra-SVM HDFA
- Configure HDFA and the data LIFs to realize the benefits of using intracluster caching with ONTAP configuration
- Learn how to configure clients to distribute ONTAP NAS connections with client configuration

Architecting an ONTAP FlexCache hotspot remediation solution


To remediate hotspotting, explore the underlying causes of bottlenecks, why auto-provisioned FlexCache isn't sufficient, and the technical details necessary to effectively architect a FlexCache solution. By understanding and implementing high-density FlexCache arrays (HDFA), you can optimize performance and eliminate bottlenecks in your high-demand workloads.

Understanding the bottleneck

The following [image](#) shows a typical single-file hotspotting scenario. The volume is a FlexGroup with a single constituent per node, and the file resides on node 1.

If you distribute all of the NAS clients' network connections across different nodes in the cluster, you still bottleneck on the CPU that services the volume affinity where the hot file resides. You also introduce cluster network traffic (east-west traffic) to the calls coming from clients connected to nodes other than where the file resides. The east-west traffic overhead is typically small, but for high-performance compute workloads every little bit counts.

Figure 1: FlexGroup single-file hotspot scenario




Why an auto-provisioned FlexCache isn't the answer

To remedy hotspotting, eliminate the CPU bottleneck and preferably the east-west traffic too. FlexCache can

help if set up properly.

In the following example, FlexCache is auto-provisioned with System Manager, NetApp Console, or default CLI arguments. [Figure 1](#) and [figure 2](#) at first appear the same: both are four-node, single-constituent NAS containers. The only difference is that figure 1's NAS container is a FlexGroup, and figure 2's NAS container is a FlexCache. Each figure profiles the same bottleneck: node 1's CPU for volume affinity servicing access to the hot file, and east-west traffic contributing to latency. An auto-provisioned FlexCache hasn't eliminated the bottleneck.

Figure 2: Auto-provisioned FlexCache scenario



Anatomy of a FlexCache

To effectively architect a FlexCache for hotspot remediation, you need to understand some technical details about FlexCache.

FlexCache is always a sparse FlexGroup. A FlexGroup is made up of multiple FlexVols. These FlexVols are called FlexGroup constituents. In a default FlexGroup layout, there are one or more constituents per node in the cluster. The constituents are "sewn together" under an abstraction layer and presented to the client as a single large NAS container. When a file is written to a FlexGroup, ingest heuristics determine which constituent the file will be stored on. It might be a constituent containing the client's NAS connection or it might be a different node. The location is irrelevant because everything operates under the abstraction layer and is invisible to the client.

Let's apply this understanding of FlexGroup to FlexCache. Because FlexCache is built on a FlexGroup, by default you have a single FlexCache that has constituents on all the nodes in the cluster, as depicted in [figure 1](#). In most cases, this is a great thing. You are utilizing all the resources in your cluster.

For remediating hot files, however, this isn't ideal because of the two bottlenecks: CPU for a single file and east-west traffic. If you create a FlexCache with constituents on every node for a hot file, that file will still reside on only one of the constituents. This means there's one CPU to service all access to the hot file. You also want to limit the amount of east-west traffic required to reach the hot file.

The solution is an array of high-density FlexCaches.

Anatomy of a high-density FlexCache

A high-density FlexCache (HDF) will have constituents on as few nodes as the capacity requirements for the cached data allow. The goal is to get your cache to live on a single node. If capacity requirements make that impossible, you can have constituents on only a few nodes instead.

For example, a 24-node cluster could have three high-density FlexCaches:

- One that spans nodes 1 through 8
- A second that spans nodes 9 through 16
- A third that spans nodes 17 through 24

These three HDFs would make up one high-density FlexCache array (HDFA). If the files are evenly distributed within each HDF, you will have a one-in-eight chance that the file requested by the client resides local to the front-end NAS connection. If you were to have 12 HDFs that span only two nodes each, you have a 50% chance of the file being local. If you can collapse the HDF down to a single node, and create 24 of them, you are guaranteed that the file is local.

This configuration will eliminate all east-west traffic and, most importantly, will provide 24 CPUs/volume affinities for accessing the hot file.

What's next?

[Decide on FlexCache array density](#)

Related information

[Documentation on FlexGroup and TRs](#)

Determine ONTAP FlexCache density

Your first hotspot remediation design decision is to figure out FlexCache density. The following examples are four-node clusters. Assume that the file count is evenly distributed among all the constituents in each HDF. Assume also an even distribution of frontend NAS connections across all nodes.

Although these examples aren't the only configurations you can use, you should understand the guiding design principle to make as many HDFs as your space requirements and available resources allow.




HDFAs are represented using the following syntax: HDFs per HDFA x nodes per HDF x constituents per node per HDF

2x2x2 HDFA configuration

Figure 1 is an example of a 2x2x2 HDFA configuration: two HDFs, each spanning two nodes, and each node containing two constituent volumes. In this example, each client has a 50% chance of having direct access to the hot file. Two of the four clients have east-west traffic. Importantly, there are now two HDFs, which means two distinct caches of the hot file. There are now two CPUs/volume affinities servicing access to the hot file.


Figure 1: 2x2x2 HDFA configuration



4x1x4 HDFA configuration

Figure 2 represents an optimal configuration. It is an example of a 4x1x4 HDFA configuration: four HDFs, each contained to a single node, and each node containing four constituents. In this example, each client is guaranteed to have direct access to a cache of the hot file. Since there are four cached files on four different nodes, four different CPUs/volume affinities help service access to the hot file. Additionally, there is zero east-west traffic generated.

Figure 2: 4x1x4 HDFA configuration



What's next

After you decide how dense you want to make your HDFs, you must make another design decision if you will be accessing the HDFs with NFS with [inter-SVM HDFAs](#) and [intra-SVM HDFAs](#).

Determine an ONTAP inter-SVM or intra-SVM HDFA option

After you determine the density of your HDFs, decide whether you will be accessing the HDFs using NFS and learn about inter-SVM HDFA and intra-SVM HDFA options.




If only SMB clients will be accessing the HDFs, you should create all HDFs in a single SVM. Refer to Windows client configuration to see how to use DFS targets for load balancing.

Inter-SVM HDFA deployment

An inter-SVM HDFA requires an SVM be created for each HDF in the HDFA. This allows all HDFs within the HDFA to have the same junction-path, allowing for easier configuration on the client side.

In the [figure 1](#) example, each HDF is in its own SVM. This is an inter-SVM HDFA deployment. Each HDF has a junction-path of /hotspots. Also, every IP has a DNS A record of hostname cache. This configuration leverages DNS round-robin to load balance mounts across the different HDFs.

Figure 1: 4x1x4 inter-SVM HDFA configuration




Intra-SVM HDFA deployment

An intra-SVM requires each HDF to have a unique junction-path, but all HDFs are in one SVM. This setup is easier in ONTAP because it requires only one SVM, but it needs more advanced configuration on the Linux side with autofs and data LIF placement in ONTAP.

In the [figure 2](#) example, every HDF is in the same SVM. This is an intra-SVM HDFA deployment and requires the junction-paths to be unique. To make load balancing work appropriately, you'll need to create a unique DNS name for each IP and place the data LIFs the hostname resolves to only on the nodes where the HDF resides. You'll also need to configure autofs with multiple entries as covered in [Linux client configuration](#).

Figure 2: 4x1x4 intra-SVM HDFA configuration



What's next

Now that you have an idea of how you want to deploy your HDFA, [deploy the HDFA and configure the clients to access them in a distributed fashion](#).

Configure HDFA and ONTAP data LIFs

You'll need to configure the HDFA and the data LIFs appropriately to realize the benefits of this hotspot remediation solution. This solution uses intracluster caching with the origin and HDFA in the same cluster.

The following are two HDFA sample configurations:

- 2x2x2 inter-SVM HDFA
- 4x1x4 intra-SVM HDFA

About this task


Perform this advanced configuration using the ONTAP CLI. There are two configurations you must use in the `flexcache create` command, and one configuration you must make sure isn't configured:

- `-aggr-list`: Provide an aggregate, or list of aggregates, that reside on the node or subset of nodes you want to restrict the HDFA to.
- `-aggr-list-multiplier`: Determine how many constituents will be created per aggregate listed in the `aggr-list` option. If you have two aggregates listed, and set this value to 2, you will end up with four constituents. NetApp recommends up to 8 constituents per aggregate, but 16 is also sufficient.
- `-auto-provision-as`: If you tab out, the CLI will try to autofill and set the value to `flexgroup`. Make sure this isn't configured. If it appears, delete it.

Create a 2x2x2 inter-SVM HDFA configuration

1. To assist in configuring a 2x2x2 inter-SVM HDFA as shown in Figure 1, complete a prep sheet.

Figure 1: 2x2x2 Inter-SVM HDFA layout



SVM	Nodes per HDF	Aggregates	Constituents per node	Junction path	Data LIF IPs
svm1	node1, node2	aggr1, aggr2	2	/hotspot	192.168.0.11, 192.168.0.12
svm2	node3, node4	aggr3, aggr4	2	/hotspot	192.168.0.13, 192.168.0.14

2. Create the HDFs. Run the following command twice, once for each row in the prep sheet. Make sure you adjust the `vserver` and `aggr-list` values for the second iteration.

```
cache::> flexcache create -vserver svm1 -volume hotspot -aggr-list
aggr1,aggr2 -aggr-list-multiplier 2 -origin-volume <origin_vol> -origin
-vserver <origin_svm> -size <size> -junction-path /hotspot
```

3. Create the data LIFs. Run the command four times, creating two data LIFs per SVM on the nodes listed in the prep sheet. Make sure you adjust the values appropriately for each iteration.

```
cache::> net int create -vserver svm1 -home-port e0a -home-node node1
-address 192.168.0.11 -netmask-length 24
```


What's next

Now you need to configure your clients to utilize the HDFA appropriately. See [client configuration](#).

Create a 4x1x4 intra-SVM HDFA

1. To assist in configuring a 4x1x4 inter-SVM HDFA as shown in figure 2, fill out a prep sheet.

Figure 2: 4x1x4 intra-SVM HDFA layout



SVM	Nodes per HDF	Aggregates	Constituents per node	Junction path	Data LIF IPs
svm1	node1	aggr1	4	/hotspot1	192.168.0.11
svm1	node2	aggr2	4	/hotspot2	192.168.0.12
svm1	node3	aggr3	4	/hotspot3	192.168.0.13
svm1	node4	aggr4	4	/hotspot4	192.168.0.14

2. Create the HDFs. Run the following command four times, once for each row in the prep sheet. Make sure you adjust the `aggr-list` and `junction-path` values for each iteration.

```
cache::> flexcache create -vserver svm1 -volume hotspot1 -aggr-list aggr1 -aggr-list-multiplier 4 -origin-volume <origin_vol> -origin -vserver <origin_svm> -size <size> -junction-path /hotspot1
```

3. Create the data LIFs. Run the command four times, creating a total of four data LIFs in the SVM. There should be one data LIF per node. Make sure you adjust the values appropriately for each iteration.

```
cache::> net int create -vserver svm1 -home-port e0a -home-node node1 -address 192.168.0.11 -netmask-length 24
```

What's next

Now you need to configure your clients to utilize the HDFA appropriately. See [client configuration](#).

Configure clients to distribute ONTAP NAS connections

To remedy hotspotting, configure the client properly to do its part in preventing CPU bottleneck.

Linux client configuration

Whether you chose an intra-SVM or inter-SVM HDFA deployment, you should use `autofs` in Linux to make sure clients are load-balancing across the different HDFs. The `autofs` configuration will differ for inter- and intra-SVM.

Before you begin

You'll need `autofs` and the appropriate dependencies installed. For help with this, refer to Linux documentation.

About this task

The steps described will use an example `/etc/auto_master` file with the following entry:

```
/flexcache auto_hotspot
```

This configures `autofs` to look for a file called `auto_hotspot` in the `/etc` directory any time a process tries to access the `/flexcache` directory. The contents of the `auto_hotspot` file will dictate which NFS server and junction-path to mount inside the `/flexcache` directory. The examples described are different configurations for the `auto_hotspot` file.


Intra-SVM HDFA autofs configuration

In the following example, we'll create an `autofs` map for the diagram in [figure 1](#). Because each cache has the same junction-path, and the hostname `cache` has four DNS A records, we only need one line:

```
hotspot cache:/hotspot
```

This one simple line will cause the NFS client to do a DNS lookup for hostname `cache`. DNS is setup to return the IPs in a round-robin fashion. This will result in an even distribution of front-end NAS connections. After the client receives the IP, it will mount the junction-path `/hotspot` at `/flexcache/hotspot`. It could be connected to SVM1, SVM2, SVM3, or SVM4, but the particular SVM doesn't matter.

Figure 1: 2x2x2 inter-SVM HDFA



Intra-SVM HDFA configuration

In the following example, we'll create an `autofs` map for the diagram in [figure 2](#). We need to make sure the NFS clients mount the IPs that are a part of the HDF junction-path deployment. In other words, we don't want to mount `/hotspot1` with anything other than IP 192.168.0.11. To do this, we can list all four IP/junction-path pairs for one local mount location in the `auto_hotspot` map.




The backslash (\) in the following example continues the entry to the next line, making it easier to read.

```
hotspot      cach1:/hostspot1 \
              cache2:/hostspot2 \
              cache3:/hostspot3 \
              cache4:/hostspot4
```

When the client tries to access `/flexcache/hotspot`, `autofs` is going to do a forward-lookup for all four hostnames. Assuming all four IPs are either in the same subnet as the client or in a different subnet, `autofs` will issue an NFS NULL ping to each IP.

This NULL ping requires the packet to be processed by ONTAP's NFS service, but it doesn't require any disk access. The first ping to return is going to be the IP and junction-path `autofs` chooses to mount.

Figure 2: 4x1x4 intra-SVM HDFA



Windows client configuration

With Windows clients, you should use an intra-SVM HDFA. To load balance across the different HDFs in the SVM, you must add a unique share name to each HDF. After that, follow the steps in [Microsoft documentation](#) to implement multiple DFS targets for the same folder.

Copyright information

Copyright © 2026 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.