



Manage FlexGroup volumes

ONTAP 9

NetApp
February 12, 2026

This PDF was generated from <https://docs.netapp.com/us-en/ontap/flexgroup/monitor-space-usage-task.html> on February 12, 2026. Always check docs.netapp.com for the latest.

Table of Contents

Manage FlexGroup volumes	1
Monitor the space usage of ONTAP FlexGroup volumes	1
Increase the size of ONTAP FlexGroup volumes	3
Reduce the size of ONTAP FlexGroup volumes	5
Configure ONTAP FlexGroup volumes to automatically grow and shrink their size	6
Delete directories asynchronously from ONTAP FlexGroup volumes	7
Delete directories asynchronously	7
Cancel a directory delete job	9
Manage client rights to delete ONTAP directories asynchronously with FlexGroups	9
Enable client asynchronous directory delete	10
Disable client asynchronous directory delete	11
Create qtrees with ONTAP FlexGroup volumes	11
Use quotas for ONTAP FlexGroup volumes	12
Quota targets and types	12
Behavior of FlexGroup volumes when quota limits are exceeded	13
Examples of quota enforcement for FlexGroup volumes	14
Apply rules and limits on the FlexGroup volume	19
Enable storage efficiency on ONTAP FlexGroup volumes	21
Protect ONTAP FlexGroup volumes using snapshots	23
Move constituents from ONTAP FlexGroup volumes	24
Use aggregates in FabricPool for existing ONTAP FlexGroup volumes	26
Balance ONTAP FlexGroup volumes by redistributing file data	28
How files are distributed	28
Supported protocols	28
Enable advanced capacity balancing	29
Rebalance ONTAP FlexGroup volumes by moving files	31
FlexGroup rebalancing considerations	32
Enable FlexGroup rebalancing	33
Modify FlexGroup rebalance configurations	34
Stop FlexGroup rebalance	35
View FlexGroup rebalance status	36

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 quotas 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 Mapping	Disk Limit	Disk Limit	Files Limit	Files 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
----- ----- ----- ----- ----- ----- ----- -----
-----
```

Volume	Tree	Type	ID	Used	Limit	Used	Limit
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
Type Target Qtree Mapping Limit Limit Limit Limit
Threshold
----- ----- ----- ----- ----- ----- -----
-----
```

Type	Target	Qtree	Mapping	Soft Limit	Soft Limit	Files Limit	Files Limit
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%
-----  -----
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
-----  -----
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:

- 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
```

- 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
```

- 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:

- 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
```

- 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

```

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.