



volume commands

ONTAP 9.15.1 commands

NetApp
December 18, 2024

Table of Contents

- volume commands 1
 - volume autosize 1
 - volume create 2
 - volume delete 16
 - volume expand 17
 - volume make-vsroot 19
 - volume modify 19
 - volume mount 32
 - volume offline 33
 - volume online 34
 - volume rehost 35
 - volume rename 36
 - volume restrict 36
 - volume show-footprint 37
 - volume show-space 46
 - volume show 52
 - volume size 78
 - volume transition-prepare-to-downgrade 79
 - volume unmount 80
- volume activity-tracking commands 80
- volume analytics commands 83
- volume clone commands 88
- volume conversion commands 110
- volume efficiency commands 112
- volume encryption commands 146
- volume file commands 153
- volume flexcache commands 188
- volume flexgroup commands 208
- volume inode-upgrade commands 208
- volume move commands 211
- volume object-store commands 228
- volume qtree commands 231
- volume quota commands 242
- volume reallocation commands 274
- volume rebalance commands 283
- volume recovery-queue commands 311
- volume schedule-style commands 316
- volume snaplock commands 316
- volume snapshot commands 323

volume commands

volume autosize

Set/Display the autosize settings of the flexible volume.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume autosize` command allows the user to specify the maximum size that a volume automatically grows to when it is out of space or the minimum size that it shrinks to when the amount of used space is below a certain threshold. If only the volume/Vserver name is specified, then the current settings are displayed. This command is not supported on infinite volumes.

Parameters

-vserver <vserver name> - Vserver Name

Use this parameter to specify the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This parameter specifies the volume for which the user wants to set or display the autosize configuration.

{ [-maximum-size {<integer>[KB|MB|GB|TB|PB] }] - Maximum Autosize

Use this parameter to specify the maximum size to which a flexible volume can grow. The default for flexible volumes is 120% of the volume size. If the value of this parameter is invalidated by manually growing the volume to a size more than the value of the maximum size, the maximum size is reset to 120% of the volume size. The value for `-maximum-size` cannot be set larger than the platform-dependent maximum flexible volume size. If you specify a larger value, the value of `-maximum-size` is automatically reset to the supported maximum without returning an error. The value for `-maximum-size` supports relative sizing.

[-minimum-size {<integer>[KB|MB|GB|TB|PB] }] - Minimum Autosize

This parameter specifies the minimum size to which the volume can automatically shrink. If the volume was created with the `grow_shrink` autosize mode enabled, then the default minimum size is equal to the initial volume size. If the value of the `-minimum-size` parameter is invalidated by a manual volume resize or is invalid when autosize is enabled, the minimum size is reset to the volume size. The value for `-minimum-size` supports relative sizing. This parameter is not supported on infinite volumes.

[-grow-threshold-percent <percent>] - Grow Threshold Used Space Percentage

This parameter specifies the used space threshold for the automatic growth of the volume. When the volume's used space becomes greater than this threshold, the volume will automatically grow unless it has reached the maximum autosize. This parameter is not supported on infinite volumes.

[-shrink-threshold-percent <percent>] - Shrink Threshold Used Space Percentage

This parameter specifies the used space threshold for the automatic shrinking of the volume. When the amount of used space in the volume drops below this threshold, the volume will shrink unless it has reached the specified minimum size. This parameter is not supported on infinite volumes.

[`-mode {off|grow|grow_shrink}`] - Autosize Mode

This parameter specifies the autosize mode for the volume. The supported autosize modes are:

- *off* - The volume does not grow or shrink in size in response to the amount of used space.
- *grow* - The volume will automatically grow when used space in the volume is above the grow threshold.
- *grow_shrink* - The volume will grow or shrink in size in response to the amount of used space.

By default, `-mode` is *off* for new flexible volumes, except for data protection (DP) mirrors, for which the default value is *grow_shrink*. The *grow* and *grow_shrink* modes work together with Snapshot autodelete to automatically reclaim space when a volume is about to become full. The volume parameter `-space-mgmt-try-first` controls the order in which these two space reclamation policies are attempted.

[`-reset <true>`] - Autosize Reset }

Use this option to reset the values of autosize, max-autosize, min-autosize, autosize-grow-threshold-percent, autosize-shrink-threshold-percent and autosize-mode to their default values based on the current size of the volume. For example, the max-autosize value is set to 120% of the current size of the volume.

Examples

The following example sets the autosize settings on a volume named vol1. The maximum size to grow is 1TB and autogrow is enabled.

```
cluster1::> vol autosize vol1 -maximum-size 1t -mode grow
(volume autosize)
vol autosize: Flexible volume 'vs1:vol1' autosize settings UPDATED.
```

The following example shows the autosize settings on a volume named vol1. The maximum size to grow is 1TB and autogrow is enabled.

```
cluster1::> vol autosize vol1
(volume autosize)
Volume autosize is currently ON for volume 'vs1:vol1'.
The volume is set to grow to a maximum of 1t.
```

volume create

Create a new volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume create` command creates a volume on a specified Vserver and storage aggregates. You can optionally specify the following attributes for the new volume:

- Size
- State (online, offline, or restricted)
- Type (read-write or data-protection)
- Export policy
- User ID
- Group ID
- Security style (All volume types: UNIX mode bits, CIFS ACLs, or mixed NFS and CIFS permissions.)
- Default UNIX permissions for files on the volume
- Language
- Junction path
- Whether the junction path is active (advanced privilege level or higher only)
- Whether the volume is the root volume for its Vserver (advanced privilege level or higher only)
- Comment
- Whether autosizing is enabled for FlexVols
- Maximum size for autosizing FlexVols
- Minimum size for autosize
- Grow used space threshold percentage for autosize
- Shrink used space threshold percentage for autosize
- Whether autosizing is enabled for FlexVols
- Current mode of operation of volume autosize
- Maximum directory size (advanced privilege level or higher only)
- Space guarantee style (none or volume)
- Space SLO type (none, thick or semi-thick)
- Snapshot policy
- Snapshot reserve percentage
- Use logical space reporting
- Use logical space enforcement
- Whether the volume create operation runs as a foreground or background process
- Caching policy
- Encrypt
- File system analytics state
- Tiering object tags
- Key Manager Attribute
- SnapLock type
- Activity tracking state
- Cache retention priority
- Efficiency policy

- Tiering minimum cooling days
- Cloud retrieval policy
- Whether the volume's total number of files will be set to the highest possible value
- Application IO Size
- Preserve unlink

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver on which the volume is located. If only one data Vserver exists, you do not need to specify this parameter.

-volume <volume name> - Volume Name

This specifies the name of the volume that is to be created. A volume's name must start with an alphabetic character (a to z or A to Z) and be 197 or fewer characters in length for FlexGroups, and 203 or fewer characters in length for all other volume types. Volume names must be unique within a Vserver.

{ [-aggregate <aggregate name>] - Aggregate Name

This specifies the storage aggregate on which the volume is to be created. This parameter only applies to FlexVol volumes.

| [-aggr-list <aggregate name>, ...] - List of Aggregates for FlexGroup Constituents

Specifies an array of names of aggregates to be used for FlexGroup constituents. Each entry in the list will create a constituent on the specified aggregate. The root constituent will always be placed on the first aggregate in the list, unless `optimize-aggr-list` is specified as `true`. An aggregate may be specified multiple times to have multiple constituents created on it. This parameter only applies to FlexGroups.

[-aggr-list-multiplier <integer>] - Aggregate List Repeat Count

Specifies the number of times to iterate over the aggregates listed with the `-aggr-list` parameter when creating a FlexGroup. The aggregate list will be repeated the specified number of times. Example:

```
-aggr-list aggr1,aggr2 -aggr-list-multiplier 2
```

will cause four constituents to be created in the order `aggr1`, `aggr2`, `aggr1`, `aggr2`.

+

The default value is 4.

+

This parameter only applies to FlexGroups.

[-optimize-aggr-list {true|false}] - Have the System Optimize the Order of the Aggregate List

Specifies whether to create the constituents of the FlexGroup volume on the aggregates specified in the `aggr-list` in the order they are specified, or whether the system should optimize the ordering of the aggregates. If this value is `true`, the system will optimize the ordering of the aggregates specified in the `aggr-list`. If this value is `false` the order of the `aggr-list` will be unchanged. The default value is `false`. This parameter only applies to FlexGroups.

[`-auto-provision-as <FlexGroup>`] - Automatically Provision as Volume of Type

Use this parameter to automatically select existing aggregates for provisioning FlexGroup volumes. Note that the fastest aggregate type with at least one aggregate on each node of the cluster will be selected. When auto provisioning a FlexGroup volume, the size of the FlexGroup volume should be a minimum of 800 GB per node.

When the UDO License is installed, this will be the default for volume creation if neither the `aggregate` nor `aggr-list` parameters are specified.

This parameter only applies to FlexGroups.

[`-support-tiering {true|false}`] - Automatically Provision FlexGroup on FabricPools

This parameter specifies whether or not FabricPools are selected when provisioning a FlexGroup during the protection workflows using the `auto-provision-as` parameter. Only FabricPool aggregates are used if this parameter is set to `true` and only non FabricPool aggregates are used if this parameter is set to `false`. Tiering support for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is `false`

This parameter only applies to FlexGroups created using the `-auto-provision-as` parameter.

[`-use-mirrored-aggregates {true|false}`] - Automatically Provision FlexGroup on Mirrored Aggregates

Use this parameter to specify whether mirrored aggregates are selected when creating a FlexGroup using the `auto-provision-as` parameter. Only mirrored aggregates are used if this parameter is set to `true` and only unmirrored aggregates are used if this parameter is set to `false`. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is `true` for a MetroCluster configuration and is `false` for a non-MetroCluster configuration

This parameter only applies to FlexGroups created using the `-auto-provision-as` parameter.

[`-encryption-type {none|volume|aggregate|mixed}`] - Encryption Type

Use this parameter to specify the encryption-type while creating a FlexGroup using the `auto-provision-as` parameter. If the value is `none`, the FlexGroup created will be unencrypted. If the value is `volume`, the FlexGroup created will be of type NVE (NetApp Volume Encryption) and if the value is `aggregate`, the FlexGroup created will be of type NAE (NetApp Aggregate Encryption).

This parameter only applies to FlexGroups created using the `-auto-provision-as` parameter.

[`-nodes {<nodename>|local}`] - List of Nodes Hosting the Volume }

Specifies an array of node names to be used for provisioning the FlexGroup. If an array of node names is specified, only aggregates from the specified nodes will be considered for provisioning. If no value is specified, all nodes in the cluster will be used.

[`-size {<integer>[KB|MB|GB|TB|PB]}`] - Volume Size

This optionally specifies the size of the volume. The size is specified as a number followed by a unit designation: k (kilobytes), m (megabytes), g (gigabytes), or t (terabytes). If the unit designation is not specified, bytes are used as the unit, and the specified number is rounded up to the nearest 4 KB. The minimum size for a FlexVol volume is 20 MB. The minimum size for a volume guaranteed FlexGroup is 20 MB per constituent. The minimum size for a none guaranteed FlexGroup is 200 MB per constituent. However, the recommended size for all FlexGroups is a minimum of 100 GB per constituent. For all volumes, the default size is set to the minimum size. The volume's maximum size is limited by the platform maximum. If the volume's guarantee is set to `volume`, the volume's maximum size can also be limited by

the available space in the hosting aggregates. Volumes can be increased and decreased in size with the [volume modify](#) command. The maximum number of files a volume is configured for is listed under "Total Files" when running the command `volume show -instance`.

[`-state {online|restricted|offline|force-online|force-offline|mixed}`] - Volume State

This optionally specifies the volume's state. A restricted volume does not provide client access to data but is available for administrative operations.



The *mixed* state applies to FlexGroups only and cannot be specified as a target state.

[`-policy <text>`] - Export Policy

This optionally specifies the ID number of the export policy associated with the volume. For information on export policies, see the documentation for the [vserver export-policy create](#) command. FlexGroups do not support policies that allow NFSv4 protocol access.

[`-user <user name>`] - User ID

This optionally specifies the name or ID of the user that is set as the owner of the volume's root.

[`-group <group name>`] - Group ID

This optionally specifies the name or ID of the group that is set as the owner of the volume's root.

[`-security-style <security style>`] - Security Style

This optionally specifies the security style for the volume. Possible values include *unix* (for UNIX mode bits), *ntfs* (for CIFS ACLs), *mixed* (for mixed NFS and CIFS permissions) and *unified* (for mixed NFS and CIFS permissions with unified ACLs). Regardless of the security style, both NFS and CIFS clients can read from and write to the volume.

[`-unix-permissions <unix perm>`] - UNIX Permissions

This optionally specifies the default UNIX permissions for files on the volume. Specify UNIX permissions either as a four-digit octal value (for example, 0700) or in the style of the UNIX `ls` command (for example, `-rwxr-x---`). For information on UNIX permissions, see the UNIX or Linux documentation. The default setting is 0755 or `---rwxr-xr-x`.

[`-junction-path <junction path>`] - Junction Path

This optionally specifies the volume's junction path. The junction path name is case insensitive and must be unique within a Vserver's namespace.

[`-junction-active {true|false}`] - Junction Active (privilege: advanced)

This optionally specifies whether the volume's junction path is active. The default setting is `true`. If the junction path is inactive, the volume does not appear in the Vserver's namespace. This parameter is available only at the advanced privilege level and higher.

[`-vsroot {true|false}`] - Vserver Root Volume (privilege: advanced)

This optionally specifies whether the volume is the root volume of its Vserver. The default setting is `false`. If this parameter is set to `true`, the default size of the newly created volume is 1GB. This parameter is not supported on FlexGroups.

[`-comment <text>`] - Comment

This optionally specifies a comment for the volume.

[-max-autosize {<integer>[KB|MB|GB|TB|PB] }] - Maximum Autosize

This parameter allows the user to specify the maximum size to which a volume can grow. The default for volumes is 120% of the volume size. If the value of this parameter is invalidated by manually resizing the volume, the maximum size is reset to 120% of the volume size. The value for `-max-autosize` cannot be set larger than the platform-dependent maximum volume size. If you specify a larger value, the value of `-max-autosize` is automatically reset to the supported maximum without returning an error.

[-min-autosize {<integer>[KB|MB|GB|TB|PB] }] - Minimum Autosize

This parameter specifies the minimum size to which the volume can automatically shrink. If the volume was created with the `grow_shrink` autosize mode enabled, then the default minimum size is equal to the initial volume size. If the value of the `-min-autosize` parameter is invalidated by a manual volume resize, the minimum size is reset to the volume size.

[-autosize-grow-threshold-percent <percent>] - Autosize Grow Threshold Percentage

This parameter specifies the used space threshold for the automatic growth of the volume. When the volume's used space becomes greater than this threshold, the volume will automatically grow unless it has reached the maximum autosize.

[-autosize-shrink-threshold-percent <percent>] - Autosize Shrink Threshold Percentage

This parameter specifies the used space threshold for the automatic shrinking of the volume. When the amount of used space in the volume drops below this threshold, the volume will shrink unless it has reached the specified minimum size.

[-autosize-mode {off|grow|grow_shrink}] - Autosize Mode

This parameter specifies the autosize mode for the volume. The supported autosize modes are:

- *off* - The volume will not grow or shrink in size in response to the amount of used space.
- *grow* - The volume will automatically grow when used space in the volume is above the grow threshold.
- *grow_shrink* - The volume will grow or shrink in size in response to the amount of used space.

By default, `-autosize-mode` is *off* for new volumes, except for data protection mirrors, for which the default value is *grow_shrink*. The *grow* and *grow_shrink* modes work together with Snapshot autodelete to automatically reclaim space when a volume is about to become full. The volume parameter `-space-mgmt-try-first` controls the order in which these two space reclamation policies are attempted.

[-files-set-maximum {true|false}] - Set Total Files (for user-visible data) to the Highest Value that the Volume can Hold (privilege: advanced)

This optionally specifies whether the volume's total number of files will be set to the highest possible value. By default, it is *false*. If *true*, the volume's total number of files is set to the highest value that the volume can hold.

[-maxdir-size {<integer>[KB|MB|GB|TB|PB] }] - Maximum Directory Size (privilege: advanced)

This optionally specifies the maximum directory size. The default maximum directory size is model-dependent and optimized for the size of system memory.

{ [-space-slo {none|thick|semi-thick}] - Space SLO

This optionally specifies the Service Level Objective for space management (the space SLO setting) for the volume. The space SLO value is used to enforce volume settings so that sufficient space is set aside to meet the space SLO. The default setting is *none*. There are three supported values: *none*, *thick* and

semi-thick.

- *none* : The value of *none* does not provide any guarantee for overwrites or enforce any restrictions. It should be used if the admin plans to manually manage space consumption in the volume and aggregate, and out of space errors.
- *thick* : The value of *thick* guarantees that the hole fills and overwrites to space-reserved files in this volume will always succeed by reserving space. To meet this space SLO, the following volume-level settings are automatically set and cannot be modified:
- Space Guarantee: *volume* - The entire size of the volume is preallocated in the aggregate. Changing the volume's *space-guarantee* type is not supported.
- Fractional Reserve: *100* - 100% of the space required for overwrites is reserved. Changing the volume's *fractional-reserve* setting is not supported.
- *semi-thick* : The value of *semi-thick* is a best-effort attempt to ensure that overwrites succeed by restricting the use of features that share blocks and auto-deleting backups and Snapshot copies in the volume. To meet this space SLO, the following volume-level settings are automatically set and cannot be modified:
- Space Guarantee: *volume* - The entire size of the volume is preallocated in the aggregate. Changing the volume's *space-guarantee* type is not supported.
- Fractional Reserve: *0* - No space will be reserved for overwrites by default. However, changing the volume's *fractional-reserve* setting is supported. Changing the setting to 100 means that 100% of the space required for overwrites is reserved.
- Snapshot Autodelete: *enabled* - Automatic deletion of Snapshot copies is enabled to reclaim space. To ensure that the overwrites can be accommodated when the volume reaches threshold capacity, the following volume Snapshot autodelete parameters are set automatically to the specified values and cannot be modified:
- *enabled* : *true*
- *commitment* : *destroy*
- *trigger* : *volume*
- *defer-delete* : *none*
- *destroy-list* : *vol_clone, lun_clone, file_clone, cifs_share*

In addition, with a value of ```_semi-thick```, the following technologies are not supported for the volume:

- File Clones with autodelete disabled: Only full file clones of files, LUNs or NVMe namespaces that can be autodeleted can be created in the volume. The use of *autodelete* for file clone create is required.
- Partial File Clones: Only full file clones of files or LUNs that can be autodeleted can be created in the volume. The use of *range* for file clone create is not supported.
- Volume Efficiency: Enabling volume efficiency is not supported to allow autodeletion of Snapshot copies.

| [-s, -space-guarantee {none|volume}] - Space Guarantee Style

This optionally specifies the space guarantee style for the volume. A value of *volume* reserves space on

the aggregates for the entire volume. A value of *none* reserves no space on the aggregates, meaning that writes can fail if an aggregate runs out of space. The default setting for the volumes on All Flash FAS systems is *none*, otherwise the default setting is *volume*. The *file* setting is no longer supported.

[`-type {RW|DP}`] - Volume Type }

This optionally specifies the volume's type, either read-write (RW) or data-protection (DP). If you do not specify a value for this parameter, a RW volume is created by default.

[`-snapdir-access {true|false}`] - Snapshot Directory Access Enabled

This optionally specifies whether clients have access to `.snapshot` directories. The default setting is `true`.

[`-percent-snapshot-space <percent>`] - Space Reserved for Snapshot Copies

This optionally specifies the amount of space that is reserved in the volume for Snapshot copies. The default setting is 5 percent, except for data protection mirrors for which the default is 0 percent.

[`-snapshot-policy <snapshot policy>`] - Snapshot Policy

This optionally specifies the Snapshot policy for the volume. The default is the Snapshot policy for all volumes on the Vserver, as specified by the `-snapshot-policy` parameter of the `vserver create` and `vserver modify` commands. The schedules associated with the `snapshot-policy` for a FlexGroup cannot have an interval shorter than 30 minutes.

[`-language <Language code>`] - Language

This optionally specifies the language encoding setting for the volume. By default, the volume inherits the Vserver language encoding setting.



You cannot modify the language encoding setting of a volume.

[`-foreground {true|false}`] - Foreground Process

This specifies whether the operation runs in the foreground. The default setting is `true` (the operation runs in the foreground). When set to `true`, the command will not return until the operation completes. This parameter applies only to FlexGroups. For FlexVol volumes, the command always runs in the foreground.

[`-nvfail {on|off}`] - NVFAIL Option

Setting this optional parameter to `true` causes the volume to set the `in-nvfailed-state` flag to `true`, if committed writes to the volume are lost due to a failure. The `in-nvfailed-state` flag fences the volume from further data access and prevents possible corruption of the application data. Without specifying a value, this parameter is automatically set to `false`.

[`-constituent-role <Constituent Roles>`] - Constituent Volume Role

This parameter is no longer supported.

{ [`-qos-policy-group <text>`] - QoS Policy Group Name

This optional parameter specifies which QoS policy group to apply to the volume. This policy group defines measurable service level objectives (SLOs) that do not adjust based on the volume allocated space or used space. If you do not assign a policy group to a volume, the system will not monitor and control the traffic to it.

| [`-qos-adaptive-policy-group <text>`] - QoS Adaptive Policy Group Name }

This optional parameter specifies which QoS adaptive policy group to apply to the volume. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs) that adjust

based on the volume allocated space or used space.

[-caching-policy <text>] - Caching Policy Name

This optionally specifies the caching policy to apply to the volume. A caching policy defines how the system caches this volume's data in a Flash Pool aggregate or Flash Cache modules. If a caching policy is not assigned to this volume, the system uses *auto* as the default caching policy.

Both metadata and user data are eligible for caching. Metadata consists of directories, indirect blocks and system metabytes. They are eligible for read caching only. When a random write pattern is detected on user data, the first such write is eligible for read caching while all subsequent overwrites are eligible for write caching. The available caching policies are:

- none - Does not cache any user data or metadata blocks.
- auto - Read caches all metadata and randomly read user data blocks, and write caches all randomly overwritten user data blocks.
- meta - Read caches only metadata blocks.
- random_read - Read caches all metadata and randomly read user data blocks.
- random_read_write - Read caches all metadata, randomly read and randomly written user data blocks.
- all_read - Read caches all metadata, randomly read and sequentially read user data blocks.
- all_read_random_write - Read caches all metadata, randomly read, sequentially read and randomly written user data.
- all - Read caches all data blocks read and written. It does not do any write caching.
- noread-random_write - Write caches all randomly overwritten user data blocks. It does not do any read caching.
- meta-random_write - Read caches all metadata and write caches randomly overwritten user data blocks.
- random_read_write-random_write - Read caches all metadata, randomly read and randomly written user data blocks. It also write caches randomly overwritten user data blocks.
- all_read-random_write - Read caches all metadata, randomly read and sequentially read user data blocks. It also write caches randomly overwritten user data blocks.
- all_read_random_write-random_write - Read caches all metadata, randomly read, sequentially read and randomly written user data. It also write caches randomly overwritten user data blocks.
- all-random_write - Read caches all data blocks read and written. It also write caches randomly overwritten user data blocks.

Note that in a caching-policy name, a hyphen (-) separates read and write policies. Default caching-policy is *auto*.

[-cache-retention-priority {normal|low|high}] - Cache Retention Priority (privilege: advanced)

This optionally specifies the cache retention priority to apply to the volume. A cache retention priority defines how long the blocks of a volume will be cached in flash pool once they become cold. If a cache retention priority is not assigned to this volume, the system uses the default policy. This parameter is available only at the advanced privilege level and higher.

The available cache retention priority are:

- low - Cache the cold blocks for the lowest time.

- normal - Cache the cold blocks for the default time.
- high - Cache the cold blocks for the highest time.

[`-is-autobalance-eligible {true|false}`] - Is Eligible for Auto Balance Aggregate (privilege: advanced)

If the Auto Balance feature is enabled, this parameter specifies whether the volume might be considered for system workload balancing. When set to `true`, the Auto Balance Aggregate feature might recommend moving this volume to another aggregate. The default value is `true`.

[`-max-constituent-size {<integer>[KB|MB|GB|TB|PB]}`] - Maximum size of a FlexGroup Constituent (privilege: advanced)

This optionally specifies the maximum size of a FlexGroup constituent. The default value is determined by checking the maximum FlexVol size setting on all nodes used by the FlexGroup. The smallest value found is selected as the default for the `-max-constituent-size` for the FlexGroup. This parameter applies to FlexGroups only.

[`-efficiency-policy <efficiency policy>`] - Storage Efficiency Policy (privilege: advanced)

This optionally specifies which storage efficiency policy to apply to the volume. This parameter is applicable only for All-Flash FAS. This parameter is not supported on data protection volumes on any platform. To disable all efficiency features on the volume in All-Flash FAS, use the value `none`. The default value is `auto`. The policies `inline-only` and `none` are not supported on Capacity optimized Flash with QAT supported platforms.

[`-snaplock-type {non-snaplock|compliance|enterprise}`] - SnapLock Type

This optionally specifies the SnapLock type of the volume. The default value is `non-snaplock`.

[`-vserver-dr-protection {protected|unprotected}`] - Vserver DR Protection

This optionally specifies whether the volume should be protected by Vserver level SnapMirror. This parameter is applicable only if the Vserver is the source of a Vserver level SnapMirror relationship. The default value for a volume of type "RW" is `protected`.

[`-encrypt {true|false}`] - Enable or Disable Encryption

This parameter allows the user to create an encrypted volume. When it is set to `true`, a new key is generated, and the volume is encrypted using the generated key. When it is set to `false`, the volume created is unencrypted.

[`-is-space-reporting-logical {true|false}`] - Logical Space Reporting

This optionally specifies whether to report space logically on the volume. When space is reported logically, ONTAP reports the volume space such that all the physical space saved by the storage efficiency features are also reported as used. The default setting is `false`.

[`-is-space-enforcement-logical {true|false}`] - Logical Space Enforcement

This optionally specifies whether to perform logical space accounting on the volume. When space is enforced logically, ONTAP enforces volume settings such that all the physical space saved by the storage efficiency features will be calculated as used. The default setting is `false`.

[`-tiering-policy <Tiering Policy>`] - Volume Tiering Policy

This optional parameter specifies the tiering policy to apply to the volume. This policy determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud tier when they become cold. FabricPool combines Flash (performance tier) with an object store (cloud tier) into a single aggregate. The

default tiering policy is *snapshot-only* for a FlexVol and *none* for a FlexGroup. Temperature of a volume block increases if it is accessed frequently and decreases when it is not.

The available tiering policies are:

- *snapshot-only* - This policy allows tiering of only the volume Snapshot copies not associated with the active file system. The default minimum cooling period is 2 days. The `-tiering-minimum-cooling-days` parameter can be used to override the default.
- *auto* - This policy allows tiering of both Snapshot copy data and active file system user data to the cloud tier. The default cooling period is 31 days. The `-tiering-minimum-cooling-days` parameter can be used to override the default.
- *none* - Volume blocks will not be tiered to the cloud tier.
- *all* - This policy allows tiering of both Snapshot copy data and active file system user data to the cloud tier as soon as possible without waiting for a cooling period. On DP volumes, this policy allows all transferred user data blocks to start in the cloud tier.

[`-cloud-retrieval-policy {default|on-read|never|promote}`] - Volume Cloud Retrieval Policy (privilege: advanced)

This optional parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier.

The available cloud retrieval policies are:

- *default* - This policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot-only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved.
- *on-read* - This policy retrieves tiered data for all client driven data reads.
- *never* - This policy never retrieves tiered data.
- *promote* - This policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot-only'. If the tiering policy is 'snapshot-only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud.

[`-tiering-minimum-cooling-days <integer>`] - Volume Tiering Minimum Cooling Days (privilege: advanced)

This optional parameter specifies the minimum number of days that user data blocks of the volume must be cooled before they can be considered cold and tiered out to the cloud tier. For volumes hosted on FabricPools, this parameter is used for tiering purposes and does not affect the reporting of inactive data. For volumes hosted on non-FabricPools, this parameter affects the cooling window used for the reporting of inactive data. The value specified should be greater than the frequency with which applications in the volume shift between different sets of data. Valid values are between 2 and 183. This parameter cannot be set when volume tiering policy is either "none" or "all". The default value of this parameter depends on the volume's tiering policy. See the tiering policy section of this documentation for corresponding default values. If the tiering policy on the volume gets changed, then this parameter will be reset to the default value corresponding to the new tiering policy.

[`-tiering-object-tags <text>, ...`] - Tags to be Associated with Objects Stored on a FabricPool

This optional parameter specifies tiering object tags to be associated with objects stored on a FabricPool.

Object tags should follow these rules:

- Each object tag should be a key-value pair separated by '='.
- Multiple tags should be separated by ','. Overall tags should be in format key1=value1[,key2=value2,...].
- All tags of a volume must have a unique key.
- Each tag key should start with either a letter or an underscore. Keys should contain only alphanumeric characters and underscores. Maximum allowed limit is of 127 characters.
- Each tag value should be of maximum 127 characters consisting of only alphanumeric characters and underscores.
- Maximum 4 object tags are allowed per volume.

{ [-application-io-size <Application I/O Size>] - Block Size to Use by Compression (privilege: advanced)

This optionally specifies the application IO size to apply to the volume. This parameter specifies the compression group size for enabling compression on the volume. This parameter is applicable only for All Flash FAS. This parameter is not supported on data protection volumes on any platform.

The available application I/O Sizes are:

- 8K
- auto - ONTAP will select compression group size automatically, based on temperature of the block.

[-storage-efficiency-mode {default|efficient}] - Storage Efficiency Mode

This optionally specifies efficiency mode of the volume.

The available efficiency modes are:

- default
- efficient

The `default` mode sets 8k adaptive compression on the volume. The `efficient` mode sets auto adaptive compression on the volume. This parameter is supported only on AFF platform. Auto adaptive compression will be set on QAT supported platforms irrespective of the modes.

[-analytics-state <Analytics State>] - File System Analytics State

This optionally specifies whether the volume will keep track of file system usage for the purpose of analytics. The default value is `off`.

[-activity-tracking-state <Activity Tracking State>] - Activity Tracking State

This optionally specifies whether the volume will keep track of volume activity, in real time. Tracked activities include read and write traffic for files, directories, users and clients. The default value is `off`.

[-key-manager-attribute <text>] - Key Manager Attribute

This optional parameter specifies an additional key manager attribute that should be an identifier-value pair separated by '=', ex. CRN=my-unique-value

The following identifiers are currently supported:

- CRN

[`-anti-ransomware-state` {`disabled`|`enabled`|`dry-run`|`paused`|`dry-run-paused`|`enable-paused`|`disable-in-progress`}] - Anti-ransomware State

Use this parameter to specify the anti-ransomware-state while creating a volume. If the value is *disabled*, the volume created will have anti-ransomware feature disabled. If the value is *enabled*, the volume created will have anti-ransomware feature enabled and if the value is *dry-run*, the volume created will be in the anti-ransomware evaluation mode.

[`-granular-data` {`disabled`|`basic`}] - Granular data

This parameter specifies the type of granular data storage supported on the volume. Once this property is changed to any value other than *disabled* on a volume, it cannot be disabled on that volume using this parameter. It can only be disabled on that volume by restoring a Snapshot copy. This property is only supported on FlexGroup volumes. The default value is *basic* for AWS S3 FlexGroup volumes and *disabled* otherwise. Files modified by the granular data algorithms use granular data, which might impact read/write I/O performance.

The available options are:

- `disabled` - No volume capacity rebalancing operations are performed.
- `basic` - Volume capacity rebalancing operations are enabled.

[`-atime-update-period` <integer>] - Access Time Update Period (Seconds) (privilege: advanced)

This parameter specifies the time period that must be elapsed in between consecutive atime update events. Note: The value of this parameter is only used if `-atime-update` is enabled on the volume.

[`-snapshot-locking-enabled` {`true`|`false`}] - Enable Snapshot Copy Locking

This parameter specifies whether the Snapshot copy locking feature is enabled, based on an independent compliance-clock time. Once this property is enabled, it cannot be disabled and the volume cannot be deleted until all locked Snapshot copies are past their expiry time. The default value is `false`.

[`-is-large-size-enabled` {`true`|`false`}] - Are Large Size Volumes and Files Enabled

This parameter specifies whether support for large Flexible volumes and large files is enabled. If this value is `true`, the maximum volume size for a Flexible volume will be 300TB and the maximum size of a single file will be 128TB. If this value is `false`, the maximum volume size for a Flexible volume will be 100TB and the maximum size of a single file will be 16TB. The default value is `false`.

[`-is-preserve-unlink-enabled` {`true`|`false`}] - Is Preserve Unlink Enabled (privilege: advanced)

This parameter specifies whether preserve unlink is enabled for NFSV4.1 shares on the volume. Once this feature is enabled, a file removal op on a file with existing share locks will result in the file being moved to the trash directory. Clients will continue to be able to access the file using its file handle. The last close for that file will result in deletion of the file. The default value is `false`.

[`-is-cloud-write-enabled` {`true`|`false`}] - Is Cloud Write Enabled (privilege: advanced)

This parameter specifies whether cloud write is enabled on the volume. This feature is only available on volumes in FabricPools. The default value is `false`.

[`-aggressive-readahead-mode` {`none`|`file_prefetch`}] - Aggressive readahead mode (privilege: advanced)

This parameter specifies the aggressive readahead mode of the volume. Possible values are *none* and *file_prefetch*. When set to *file_prefetch*, on a file read, the system aggressively issues readaheads for all of the blocks in the file and retains those blocks in a cache for a finite period of time. This

feature is only available on FabricPool volumes on FSx for ONTAP and Cloud Volumes ONTAP. The default value is *none*.

[`-in-consistency-group {true|false}`] - If this Volume is part of a Consistency Group

This parameter specified if this volume is associated with a consistency group.

Examples

The following example creates a new volume named `user_jdoe` on a Vserver named `vs0` and a storage aggregate named `aggr1`. Upon its creation, the volume is placed in the online state. It uses the export policy named `default_expolicy`. The owner of the volume's root is a user named `jdoe` whose primary group is named `dev`. The volume's junction path is `/user/jdoe`. The volume is 250 GB in size, space for the entire volume is reserved on the aggregate, and the create operation runs in the background.

```
cluster1::> volume create -vserver vs0 -volume user_jdoe -aggregate aggr1
               -state online -policy default_expolicy -user jdoe -group dev
               -junction-path /user/jdoe -size 250g -space-guarantee volume
               -percent-snapshot-space 20 -foreground false
```

The following example creates a new volume named `vol_cached` on a Vserver named `vs0` and a Flash Pool storage aggregate named `aggr1`. The newly created volume is placed online and uses `auto` as the caching policy.

```
cluster1::> volume create -vserver vs0 -volume vol_cached -aggregate aggr1
               -state online -caching-policy auto
```

The following example creates a new FlexGroup named `media_vol` on a Vserver named `vs0` with four constituents on aggregates `aggr1` and `aggr2`. Upon its creation, the volume is placed in the online state. The volume's junction path is `/media`. The volume is 200 TB in size, no space for the volume is reserved on the aggregates, and the create operation runs in the background.

```
cluster1::> volume create -vserver vs0 -volume media_vol
               -aggr-list aggr1,aggr1,aggr2,aggr2 -junction-path /media -size
200TB
               -space-guarantee none -foreground false
```

The following example creates a new FlexGroup volume named `fg` on a Vserver named `vs0` on aggregates selected by Data ONTAP.

```
cluster1::> volume create -vserver vs0 -volume fg -auto-provision-as
flexgroup
```

Related Links

- [volume modify](#)

- [volume show](#)
- [vserver export-policy create](#)
- [vserver create](#)
- [vserver modify](#)

volume delete

Delete an existing volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume delete` command deletes the specified volumes. Before deleting a volume, the user is prompted to confirm the operation unless the `-force` flag is specified. If this volume was associated with a policy group the underlying qos workload is deleted.

NOTE:

- If there is a qtree or quota policy associated with a volume, it is deleted when you delete the volume.
- A volume must be offline (see [volume offline](#)) to be deleted.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the name of the Vserver from which the volume is to be deleted. If only one data Vserver exists, you do not need to specify this parameter.

-volume <volume name> - Volume Name

This specifies the name of the volume that is to be deleted.

[-force <true>] - Force Delete (privilege: advanced)

If this parameter is specified, the user is not prompted to confirm each deletion operation. In addition, the operation is run only on the local node, and several potential errors are ignored. By default, this setting is `false`. This parameter is available only at the advanced privilege level and higher.

[-foreground {true|false}] - Foreground Process

This specifies whether the operation runs in the foreground. The default setting is `true` (the operation runs in the foreground). When set to `true`, the command will not return until the operation completes.

Examples

The following example deletes a volume named `vol1_old` from a Vserver named `vs0`:

```
cluster1::> volume delete -vserver vs0 -volume vol1_old
```

Related Links

- [volume offline](#)

volume expand

Expand the size of a volume by adding constituents

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume expand` command allows the user to increase the size of a FlexGroup by adding constituents. The size of the new constituents is determined by the size of the smallest existing constituent. This command only applies to FlexGroups.

Parameters

-vserver <vserver name> - Vserver Name

This parameter can be used to specify the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This parameter specifies the volume for which the user wants to expand.

-aggr-list <aggregate name>,... - List of Aggregates for FlexGroup Constituents

Specifies an array of names of aggregates to be used for new FlexGroup constituents. Each entry in the list will create a constituent on the specified aggregate. An aggregate may be specified multiple times to have multiple constituents created on it.

[-aggr-list-multiplier <integer>] - Aggregate List Repeat Count

Specifies the number of times to iterate over the aggregates listed with the `-aggr-list` parameter when expanding a FlexGroup. The aggregate list will be repeated the specified number of times. Example:

```
-aggr-list aggr1,aggr2 -aggr-list-multiplier 2
```

will cause four constituents to be created in the order `aggr1` , `aggr2` , `aggr1` , `aggr2` . The default value is 1.

[-foreground {true|false}] - Foreground Process

If `false` is specified for this parameter, the command runs as a job in the background. If `true` is specified, the command will not return until the operation is complete. The default value is `true` .

Examples

The following example increases the size of a FlexGroup by adding 3 constituents:

```

cluster1::> volume show -vserver vs1 -volume flexgroup -fields size
vserver volume      size
-----
vs1      flexgroup 180TB

cluster1::> volume expand -vserver vs1 -volume flexgroup -aggr-list
aggr1,aggr2,aggr3
Warning: The following number of constituents of size 20TB will be added
to
FlexGroup "flexgroup": 3. 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 52] Job succeeded: Successful

cluster1::> volume show -vserver vs1 -volume flexgroup -fields size
vserver volume      size
-----
vs1      flexgroup 240TB

```

The following example increase the size of a FlexGroup by adding 6 constituents using the `-aggr-list-multiplier`:

```

cluster1::> volume show -vserver vs1 -volume flexgroup -fields size
vserver volume      size
-----
vs1      flexgroup 240TB

cluster1::> volume expand -vserver vs1 -volume flexgroup -aggr-list
aggr1,aggr2 -aggr-list-multiplier 3
Warning: The following number of constituents of size 20TB will be added
to
FlexGroup "flexgroup": 6. 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 53] Job succeeded: Successful

cluster1::> volume show -vserver vs1 -volume flexgroup -fields size
vserver volume      size
-----
vs1      flexgroup 360TB

```

volume make-vsroot

Designate a non-root volume as a root volume of the Vserver

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `volume make-vsroot` command promotes a non-root volume of the Vserver to be the Vserver's root volume. The Vserver's root volume must be a FlexVol volume with a size of at least 1 GB.

For instance, if you run this command on a volume named `user` that is located on a Vserver named `vs0`, the volume `user` is made the root volume of the Vserver `vs0`.

This command is available only at the advanced privilege level and higher.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

This specifies the Vserver on which a non-root volume is to be made the root volume.

-volume <volume name> - Volume Name (privilege: advanced)

This specifies the non-root volume that is to be made the root volume of its Vserver. This must be an existing FlexVol volume. Using a SnapLock volume as the root volume for a Vserver is not supported.

Examples

The following example makes a volume named `root_vs0_backup` the root volume of its Vserver with FlexVol volumes, which is named `vs0`.

```
node::> volume make-vsroot -vserver vs0 -volume root_vs0_backup
```

The following example makes a volume named `root_vs1` the root volume of the Vserver with Infinite Volume `vs1`.

```
node::> volume make-vsroot -vserver vs1 -volume root_vs1 -aggregate aggr1
```

volume modify

Modify volume attributes

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume modify` command can be used to modify the following attributes of a volume:

- Size
- State (online, offline, restricted, force-online or force-offline)
- Export policy
- User ID
- Group ID
- Security style (All volume types: UNIX mode bits, CIFS ACLs, or mixed NFS and CIFS permissions.)
- Default UNIX permissions for files on the volume
- Whether the junction path is active
- Comment
- Volume nearly full threshold percent
- Volume full threshold percent
- Maximum size for autosizing
- Minimum size for autosize
- Grow used space threshold percentage for autosize
- Shrink used space threshold percentage for autosize
- Whether autosizing is enabled
- Current mode of operation of volume autosize
- Reset the autosize values to their defaults
- Total number of files for user-visible data permitted on the volume
- Space guarantee style (none or volume)
- Space SLO type (none, thick or semi-thick)
- Snapshot policy
- Use logical space reporting
- Use logical space enforcement
- Tiering object tags
- Convert ucode
- Whether the volume's total number of files will be set to the highest possible value
- Caching policy
- Cache retention priority
- Tiering minimum cooling days
- Cloud retrieval policy
- Preserve unlink

You can use the `volume move` command to change a volume's aggregate or node. You can use the [volume rename](#) command to change a volume's name. You can use the [volume make-vsroot](#) command to make a volume the root volume of its Vserver.

You can change additional volume attributes by using this command at the advanced privilege level and higher.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver on which the volume is located. If only one data Vserver exists, you do not need to specify this parameter. Although node Vservers are not displayed when using <Tab> completion, this parameter supports node Vservers for modifying the root volume of the specified node Vserver.

-volume <volume name> - Volume Name

This specifies the volume that is to be modified.

[-size {<integer>[KB|MB|GB|TB|PB]}] - Volume Size

This optionally specifies the new size of the volume. The size is specified as a number followed by a unit designation: k (kilobytes), m (megabytes), g (gigabytes), or t (terabytes). If the unit designation is not specified, bytes are used as the unit, and the specified number is rounded up to the nearest 4 KB. A relative rather than absolute size change can be specified by adding + or - before the given size: for example, specifying +30m adds 30 megabytes to the volume's current size. The minimum size for a volume is 20 MB (the default setting). The volume's maximum size is limited by the platform maximum. If the volume's guarantee is set to `volume`, the volume's maximum size can also be limited by the available space in the hosting aggregate. If the volume's guarantee is currently disabled, its size cannot be increased.

[-state {online|restricted|offline|force-online|force-offline|mixed}] - Volume State

This optionally specifies the volume's state. A restricted volume does not provide client access to data but is available for administrative operations.



The *mixed* state applies to FlexGroups only and cannot be specified as a target state.

[-policy <text>] - Export Policy

This optionally specifies the ID number of the export policy associated with the volume. For information on export policy, see the documentation for the [vserver export-policy create](#) command. FlexGroups do not support export policies that allow NFSv4 protocol access.

[-user <user name>] - User ID

This optionally specifies the name or ID of the user that is set as the owner of the volume's root.

[-group <group name>] - Group ID

This optionally specifies the name or ID of the group that is set as the owner of the volume's root.

[-security-style <security style>] - Security Style

This optionally specifies the security style for the volume. Possible values include *unix* (for UNIX mode bits), *ntfs* (for CIFS ACLs), *mixed* (for mixed NFS and CIFS permissions) and *unified* (for mixed NFS and CIFS permissions with unified ACLs). Regardless of the security style, both NFS and CIFS clients can read from and write to the volume.

[-unix-permissions <unix perm>] - UNIX Permissions

This optionally specifies the default UNIX permissions for files on the volume. Specify UNIX permissions either as a four-digit octal value (for example, 0700) or in the style of the UNIX `ls` command (for example, `-rwxr-x--`). For information on UNIX permissions, see the UNIX or Linux documentation. The default setting is 0755 or `-rwxr-xr-x`.

[`-junction-active {true|false}`] - Junction Active (privilege: advanced)

This optionally specifies whether the volume's junction path is active. The default setting is `true`. If the junction is inactive, the volume does not appear in the Vserver's namespace.

[`-comment <text>`] - Comment

This optionally specifies a comment for the volume.

[`-space-nearly-full-threshold-percent <percent>`] - Volume Nearly Full Threshold Percent

This optionally specifies the percentage at which the volume is considered nearly full, and above which an EMS warning will be generated. The default value is 95%. The maximum value for this option is 99%. Setting this threshold to 0 disables the volume nearly full space alerts.

[`-space-full-threshold-percent <percent>`] - Volume Full Threshold Percent

This optionally specifies the percentage at which the volume is considered full, and above which a critical EMS error will be generated. The default value is 98%. The maximum value for this option is 100%. Setting this threshold to 0 disables the volume full space alerts.

{ [`-max-autosize {<integer>[KB|MB|GB|TB|PB]}`] } - Maximum Autosize

This parameter allows the user to specify the maximum size to which a volume can grow. The default for volumes is 120% of the volume size. If the value of this parameter is invalidated by manually resizing the volume, the maximum size is reset to 120% of the volume size. The value for `-max-autosize` cannot be set larger than the platform-dependent maximum volume size. If you specify a larger value, the value of `-max-autosize` is automatically reset to the supported maximum without returning an error. The value for `-max-autosize` supports relative sizing.

[`-min-autosize {<integer>[KB|MB|GB|TB|PB]}`] - Minimum Autosize

This parameter specifies the minimum size to which the volume can automatically shrink. If the volume was created with the `grow_shrink` autosize mode enabled, then the default minimum size is equal to the initial volume size. If the value of the `-min-autosize` parameter is invalidated by a manual volume resize, the minimum size is reset to the volume size. The value for `-min-autosize` supports relative sizing.

[`-autosize-grow-threshold-percent <percent>`] - Autosize Grow Threshold Percentage

This parameter specifies the used space threshold for the automatic growth of the volume. When the volume's used space becomes greater than this threshold, the volume will automatically grow unless it has reached the maximum autosize.

[`-autosize-shrink-threshold-percent <percent>`] - Autosize Shrink Threshold Percentage

This parameter specifies the used space threshold for the automatic shrinking of the volume. When the amount of used space in the volume drops below this threshold, the volume will shrink unless it has reached the specified minimum size.

[`-autosize-mode {off|grow|grow_shrink}`] - Autosize Mode

This parameter specifies the autosize mode for the volume. The supported autosize modes are:

- `off` - The volume will not grow or shrink in size in response to the amount of used space.
- `grow` - The volume will automatically grow when used space in the volume is above the grow threshold.
- `grow_shrink` - The volume will grow or shrink in size in response to the amount of used space.

By default, `-autosize-mode` is `off` for new volumes, except for DP mirrors, for which the default value is `grow_shrink`. The `grow` and `grow_shrink` modes work together with Snapshot autodelete to

automatically reclaim space when a volume is about to become full. The volume parameter `-space-mgmt -try-first` controls the order in which these two space reclamation policies are attempted.

[`-autosize-reset <true>`] - Autosize Reset }

This allows the user to reset the values of `autosize`, `max-autosize`, `min-autosize`, `autosize-grow-threshold-percent`, `autosize-shrink-threshold-percent` and `autosize-mode` to their default values. For example, the `max-autosize` value will be set to 120% of the current size of the volume.

[`-files <integer>`] - Total User Visible Files

This optionally specifies the total number of files for user-visible data permitted on the volume. This value can be raised or lowered. Raising the total number of files does not immediately cause additional disk space to be used to track files. Instead, as more files are created on the volume, the system dynamically increases the number of disk blocks that are used to track files. The space assigned to track files is never freed, and the `files` value cannot be decreased below the current number of files that can be tracked within the assigned space for the volume.

[`-files-set-maximum {true|false}`] - Set Total Files (for user-visible data) to the Highest Value that the Volume can Hold (privilege: advanced)

This optionally specifies whether the volume's total number of files will be set to the highest possible value. If `true`, the volume's total number of files is set to the highest value that the volume can hold. Only `true` is a valid input. `false` is not permitted. To modify the total number of files to a specific value, use option `files`.

[`-maxdir-size {<integer>[KB|MB|GB|TB|PB]}`] - Maximum Directory Size (privilege: advanced)

This optionally specifies the maximum directory size. The default maximum directory size is model-dependent, and optimized for the size of system memory. You can increase it for a specific volume by using this option, but doing so could impact system performance. If you need to increase the maximum directory size, work with customer support.

{ [`-space-slo {none|thick|semi-thick}`] - Space SLO

This optionally specifies the Service Level Objective for space management (the space SLO setting) for the volume. The space SLO value is used to enforce volume settings so that sufficient space is set aside to meet the space SLO. The default setting is `none`. There are three supported values: `none`, `thick` and `semi-thick`.

- `none`: The value of `none` does not provide any guarantee for overwrites or enforce any restrictions. It should be used if the admin plans to manually manage space consumption in the volume and aggregate, and out of space errors.
- `thick`: The value of `thick` guarantees that the hole fills and overwrites to space-reserved files in this volume will always succeed by reserving space. To meet this space SLO, the following volume-level settings are automatically set and cannot be modified:
 - Space Guarantee: `volume` - The entire size of the volume is preallocated in the aggregate. Changing the volume's `space-guarantee` type is not supported.
 - Fractional Reserve: `100` - 100% of the space required for overwrites is reserved. Changing the volume's `fractional-reserve` setting is not supported.
- `semi-thick`: The value of `semi-thick` is a best-effort attempt to ensure that overwrites succeed by restricting the use of features that share blocks and auto-deleting backups and Snapshot copies in the volume. To meet this space SLO, the following volume-level settings are automatically set and cannot be modified:

- **Space Guarantee:** *volume* - The entire size of the volume is preallocated in the aggregate. Changing the volume's *space-guarantee* type is not supported.
- **Fractional Reserve:** *0* - No space will be reserved for overwrites by default. However, changing the volume's *fractional-reserve* setting is supported. Changing the setting to 100 means that 100% of the space required for overwrites is reserved.
- **Snapshot Autodelete:** *enabled* - Automatic deletion of Snapshot copies is enabled to reclaim space. To ensure that the overwrites can be accommodated when the volume reaches threshold capacity, the following volume Snapshot autodelete parameters are set automatically to the specified values and cannot be modified:
 - *enabled: true*
 - *commitment: destroy*
 - *trigger: volume*
 - *defer-delete: none*
 - *destroy-list: vol_clone, lun_clone, file_clone, cifs_share*

In addition, with a value of ```_semi-thick```, the following technologies are not supported for the volume:

- **File Clones with autodelete disabled:** Only full file clones of files, LUNs or NVMe namespaces that can be autodeleted can be created in the volume. The use of *autodelete* for file clone create is required.
- **Partial File Clones:** Only full file clones of files or LUNs that can be autodeleted can be created in the volume. The use of *range* for file clone create is not supported.
- **Volume Efficiency:** Enabling volume efficiency is not supported to allow autodeletion of Snapshot copies.

[*-s*, *-space-guarantee* {*none*|*volume*}] - Space Guarantee Style

This option controls whether the volume is guaranteed some amount of space in the aggregate. The default setting for the volumes on All Flash FAS systems is *none*, otherwise the default setting is *volume*. The *file* setting is no longer supported. Volume guaranteed means that the entire size of the volume is preallocated. The *none* value means that no space is preallocated, even if the volume contains space-reserved files or LUNs; if the aggregate is full, space is not available even for space-reserved files and LUNs within the volume. Setting this parameter to *none* enables you to provision more storage than is physically present in the aggregate (thin provisioning). When you use thin provisioning for a volume, it can run out of space even if it has not yet consumed its nominal size and you should carefully monitor space utilization to avoid unexpected errors due to the volume running out of space. For flexible root volumes, to ensure that system files, log files, and cores can be saved, the *space-guarantee* must be *volume*. This is to ensure support of the appliance by customer support, if a problem occurs. Disk space is preallocated when the volume is brought online and, if not used, returned to the aggregate when the volume is brought offline. It is possible to bring a volume online even when the aggregate has insufficient free space to preallocate to the volume. In this case, no space is preallocated, just as if the *none* option had been selected. In this situation, the *vol* options and *vol status* command display the actual value of the *space-guarantee* option, but indicate that it is disabled.

[*-fractional-reserve* <percent>] - Fractional Reserve }

This option changes the amount of space reserved for overwrites of reserved objects (LUNs, files) in a volume. The option is set to 100 by default with *guarantee* set to *volume*. A setting of 100 means that

100% of the required reserved space is actually reserved so the objects are fully protected for overwrites. The value is set to 0 by default with `guarantee` set to `none`. The value can be either 0 or 100 when `guarantee` is set to `volume` or `none`. Using a value of 0 indicates that no space will be reserved for overwrites. This returns the extra space to the available space for the volume, decreasing the total amount of space used. However, this does leave the protected objects in the volume vulnerable to out of space errors. If the percentage is set to 0%, the administrator must monitor the space usage on the volume and take corrective action.

[`-min-readahead {true|false}`] - Minimum Read Ahead (privilege: advanced)

This optionally specifies whether minimum readahead is used on the volume. The default setting is `false`.

[`-atime-update {true|false}`] - Access Time Update Enabled (privilege: advanced)

This optionally specifies whether the access time on inodes is updated when a file is read. The default setting is `true`.

[`-snapdir-access {true|false}`] - Snapshot Directory Access Enabled

This optionally specifies whether clients have access to `.snapshot` directories. The default setting is `true`.

[`-percent-snapshot-space <percent>`] - Space Reserved for Snapshot Copies

This optionally specifies the amount of space that is reserved on the volume for Snapshot copies. The default setting is 5 percent.

[`-snapshot-policy <snapshot policy>`] - Snapshot Policy

This optionally specifies the Snapshot policy for the volume. The default is the Snapshot policy for all volumes on the SVM, as specified by the `-snapshot-policy` parameter of the `vserver create` and `vserver modify` commands. When replacing a Snapshot policy on a volume, any existing Snapshot copies on the volume that do not match any of the prefixes of the new Snapshot policy will not be deleted. This is because the Snapshot scheduler will not clean up older Snapshot copies if the prefixes do not match. After the new Snapshot policy takes effect, depending on the new retention count, any existing Snapshot copies that continue to use the same prefixes might be deleted. For example, your existing Snapshot policy is set up to retain 150 weekly Snapshot copies and you create a new Snapshot policy that uses the same prefixes but changes the retention count to 50 Snapshot copies. After the new Snapshot policy takes effect, it will start deleting older Snapshot copies until there are only 50 remaining.

[`-language <Language code>`] - Language

Use this parameter to change the volume language from `*.UTF-8` to `utf8mb4`. To change the language of a volume, contact technical support.

[`-foreground {true|false}`] - Foreground Process

This specifies whether the operation runs in the foreground. The default setting is `true` (the operation runs in the foreground). When set to `true`, the command will not return until the operation completes. This parameter applies only to FlexGroups. For FlexVol volumes, the command always runs in the foreground.

[`-nvfail {on|off}`] - NVFAIL Option

Setting this optional parameter to `true` causes the volume to set the `in-nvfailed-state` flag to `true`, if committed writes to the volume are lost due to a failure. The `in-nvfailed-state` flag fences the volume from further data access and prevents possible corruption of the application data. Without specifying a value, this parameter is automatically set to `false`.

[`-in-nvfailed-state {true|false}`] - Volume's NVFAIL State (privilege: advanced)

This field is automatically set to true on a volume when committed writes to the volume are possibly lost due to a failure, and the volume has the nvfail option enabled. With this field set, the client access to the volume is fenced to protect against possible corruptions that result from accessing stale data. The administrator needs to take appropriate recovery actions to recover the volume from the possible data loss. After the recovery is completed, the administrator can clear this field and restore the client access to the volume. This field can be cleared using the CLI but it cannot be set.

[`-dr-force-nvfail {on|off}`] - Force NVFAIL on MetroCluster Switchover

Setting this optional parameter to true on a volume causes the MetroCluster switchover operation to set the in-nvfailed-state flag to true on that volume. The in-nvfailed-state flag prevents further data access to the volume. The default value is false. This parameter has no effect on a negotiated or an automatic switchover.

[`-filesystem-size-fixed {true|false}`] - Is File System Size Fixed

This parameter is only applicable for a DP relationship. This option causes the file system to remain the same size and not grow or shrink when a SnapMirrored volume relationship is broken, or when a volume add is performed on it. It is automatically set to true when a volume becomes a SnapMirrored volume. It stays set to true after the snapmirror break command is issued for the volume. This allows a volume to be SnapMirrored back to the source without needing to add disks to the source volume. If the volume is a flexible volume and the volume size is larger than the file system size, setting this option to false forces the volume size to equal the file system size. The default setting is false.

[`-extent-enabled {off|on|space-optimized}`] - (DEPRECATED)-Extent Option



This parameter has been deprecated and may be removed in a future release of Data ONTAP.

Setting this option to ``on`` or ``space-optimized`` enables extents in the volume. This causes application writes to be written in the volume as a write of a larger group of related data blocks called an extent. Using extents may help workloads that perform many small random writes followed by large sequential reads. However, using extents may increase the amount of disk operations performed on the controller, so this option should only be used where this trade-off is desired. If the option is set to ``space-optimized`` then the reallocation update will not duplicate blocks from Snapshot copies into the active file system, and will result in conservative space utilization. Using ``space-optimized`` may be useful when the volume has Snapshot copies or is a SnapMirror source, when it can reduce the storage used in the volume and the amount of data that SnapMirror needs to move on the next update. The ``space-optimized`` value can result in degraded read performance of Snapshot copies. The default value is ``off`` ; extents are not used.

[`-space-mgmt-try-first {volume_grow|snap_delete}`] - Primary Space Management Strategy

A flexible volume can be configured to automatically reclaim space in case the volume is about to run out of space, by either increasing the size of the volume using autogrow or deleting Snapshot copies in the volume using Snapshot autodelete. If this option is set to `volume_grow` the system will try to first increase the size of volume before deleting Snapshot copies to reclaim space. If the option is set to `snap_delete` the system will first automatically delete Snapshot copies and in case of failure to reclaim space will try to grow

the volume.

[`-read-realloc {off|on|space-optimized}`] - Read Reallocation Option

Setting this option to `on` or `space-optimized` enables read reallocation in the volume. This results in the optimization of file layout by writing some blocks to a new location on disk. The layout is updated only after the blocks have been read because of a user read operation, and only when updating their layout will provide better read performance in the future. Using read reallocation may help workloads that perform a mixture of random writes and large sequential reads. If the option is set to `space-optimized` then the reallocation update will not duplicate blocks from Snapshot copies into the active file system, and will result in conservative space utilization. Using `space-optimized` may be useful when the volume has Snapshot copies or is a SnapMirror source, when it can reduce the storage used in the volume and the amount of data that snapmirror needs to move on the next update. The `space-optimized` value can result in degraded read performance of Snapshot copies. The default value is `off`.

[`-sched-snap-name {create-time|ordinal}`] - Naming Scheme for Automatic Snapshot Copies

This option specifies the naming convention for automatic Snapshot copies. If set to `create-time`, automatic Snapshot copies are named using the format `<schedule_name>.yyyy-mm-dd_hhmm`. Example: "hourly.2010-04-01_0831". If set to `ordinal`, only the latest automatic Snapshot copy is named using the format `<schedule_name>.<n>`. Example: "hourly.0". Older automatic Snapshot copies are named using the format `<schedule_name>.yyyy-mm-dd_hhmm`. Example: "hourly.2010-04-01_0831".

{ [`-qos-policy-group <text>`] - QoS Policy Group Name

This optional parameter specifies which QoS policy group to apply to the volume. This policy group defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated. If you do not assign a policy group to a volume, the system will not monitor and control the traffic to it. To remove this volume from a policy group, enter the reserved keyword "none".

| [`-qos-adaptive-policy-group <text>`] - QoS Adaptive Policy Group Name }

This optional parameter specifies which QoS adaptive policy group to apply to the volume. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs) that adjust based on the volume allocated space or used space. To remove this volume from an adaptive policy group, enter the reserved keyword "none".

[`-caching-policy <text>`] - Caching Policy Name

This parameter specifies the caching policy to apply to the volume. A caching policy defines how the system caches this volume's data in a Flash Pool aggregate or Flash Cache modules.

Both metadata and user data are eligible for caching. Metadata consists of directories, indirect blocks and system metabytes. They are eligible for read caching only. When a random write pattern is detected on user data, the first such write is eligible for read caching while all subsequent overwrites are eligible for write caching. The available caching policies are:

- `none` - Does not cache any user data or metadata blocks.
- `auto` - Read caches all metadata and randomly read user data blocks, and write caches all randomly overwritten user data blocks.
- `meta` - Read caches only metadata blocks.
- `random_read` - Read caches all metadata and randomly read user data blocks.
- `random_read_write` - Read caches all metadata, randomly read and randomly written user data blocks.
- `all_read` - Read caches all metadata, randomly read and sequentially read user data blocks.

- `all_read_random_write` - Read caches all metadata, randomly read, sequentially read and randomly written user data.
- `all` - Read caches all data blocks read and written. It does not do any write caching.
- `noread-random_write` - Write caches all randomly overwritten user data blocks. It does not do any read caching.
- `meta-random_write` - Read caches all metadata and write caches randomly overwritten user data blocks.
- `random_read_write-random_write` - Read caches all metadata, randomly read and randomly written user data blocks. It also write caches randomly overwritten user data blocks.
- `all_read-random_write` - Read caches all metadata, randomly read and sequentially read user data blocks. It also write caches randomly overwritten user data blocks.
- `all_read_random_write-random_write` - Read caches all metadata, randomly read, sequentially read and randomly written user data. It also write caches randomly overwritten user data blocks.
- `all-random_write` - Read caches all data blocks read and written. It also write caches randomly overwritten user data blocks.

Note that in a caching-policy name, a hyphen (-) separates read and write policies. Default caching-policy is `auto`.

`[-is-autobalance-eligible {true|false}] - Is Eligible for Auto Balance Aggregate (privilege: advanced)`

If the Auto Balance feature is enabled, this parameter specifies whether the volume might be considered for system workload balancing. When set to `true`, the Auto Balance Aggregate feature might recommend moving this volume to another aggregate. The default value is `true`.

`[-max-constituent-size {<integer>[KB|MB|GB|TB|PB]}] - Maximum size of a FlexGroup Constituent (privilege: advanced)`

This optionally specifies the maximum size of a FlexGroup constituent. The default value is determined by checking the maximum FlexVol size setting on all nodes used by the FlexGroup. The smallest value found is selected as the default for the `-max-constituent-size` for the FlexGroup. This parameter applies to FlexGroups only.

`[-vserver-dr-protection {protected|unprotected}] - Vserver DR Protection`

This optionally specifies whether the volume should be protected by Vserver level SnapMirror. This parameter is applicable only if the Vserver is the source of a Vserver level SnapMirror relationship.

`[-is-space-reporting-logical {true|false}] - Logical Space Reporting`

This optionally specifies whether to report space logically on the volume. When space is reported logically, ONTAP reports the volume space such that all the physical space saved by the storage efficiency features are also as reported as used. The default setting is `false`.

`[-is-space-enforcement-logical {true|false}] - Logical Space Enforcement`

This optionally specifies whether to perform logical space accounting on the volume. When space is enforced logically, ONTAP enforces volume settings such that all the physical space saved by the storage efficiency features will be calculated as used. The default setting is `false`.

`[-tiering-policy <Tiering Policy>] - Volume Tiering Policy`

This optional parameter specifies the tiering policy to apply to the volume. This policy determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud tier when they become cold.

FabricPool combines Flash (performance tier) with an object store (cloud tier) into a single aggregate. The temperature of a volume block increases if it is accessed frequently, and it decreases when it is not.

The available tiering policies are:

- `snapshot-only` - This policy allows tiering of only the volume Snapshot copies not associated with the active file system. The default cooling period is 2 days. The `-tiering-minimum-cooling-days` parameter can be used to override the default.
- `auto` - This policy allows tiering of both Snapshot copy data and active file system user data to the cloud tier. The default cooling period is 31 days. The `-tiering-minimum-cooling-days` parameter can be used to override the default.
- `none` - Volume blocks will not be tiered to the cloud tier.
- `all` - This policy allows tiering of both Snapshot copy data and active file system user data to the cloud tier as soon as possible without waiting for a cooling period. On DP volumes, this policy allows all transferred user data blocks to start in the cloud tier.

`[-cloud-retrieval-policy {default|on-read|never|promote}]` - Volume Cloud Retrieval Policy (privilege: advanced)

This optional parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier.

The available cloud retrieval policies are:

- `default` - This policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot-only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved.
- `on-read` - This policy retrieves tiered data for all client driven data reads.
- `never` - This policy never retrieves tiered data.
- `promote` - This policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot-only'. If the tiering policy is 'snapshot-only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud.

`[-tiering-minimum-cooling-days <integer>]` - Volume Tiering Minimum Cooling Days (privilege: advanced)

This parameter specifies the minimum number of days that user data blocks of the volume must be cooled before they can be considered cold and tiered out to the cloud tier. For volumes hosted on FabricPools, this parameter is used for tiering purposes and does not affect the reporting of inactive data. For volumes hosted on non-FabricPools, this parameter affects the cooling window used for reporting inactive data. The value specified should be greater than the frequency with which applications in the volume shift between different sets of data. Valid values are between 2 and 183. This parameter cannot be set when volume tiering policy is either "none" or "all".

`[-tiering-object-tags <text>, ...]` - Tags to be Associated with Objects Stored on a FabricPool

This optional parameter specifies tiering object tags to be associated with objects stored on a FabricPool.

Object tags should follow these rules:

- Each object tag should be a key-value pair separated by '='.
- Multiple tags should be separated by ','. Overall tags should be in format `key1=value1[,key2=value2,...]`.

- All tags of a volume must have a unique key.
- Each tag key should start with either a letter or an underscore. Keys should contain only alphanumeric characters and underscores. Maximum allowed limit is of 127 characters.
- Each tag value should be of maximum 127 characters consisting of only alphanumeric characters and underscores.
- Maximum 4 object tags are allowed per volume.
- To remove existing tags of a volume, specify empty list as a parameter value.

`[-anti-ransomware-state {disabled|enabled|dry-run|paused|dry-run-paused|enable-paused|disable-in-progress}] - Anti-ransomware State`

Use this parameter to modify the anti-ransomware-state of the volume. The states `dry-run-paused`, `enable-paused` and `disable-in-progress` are non-modifiable. The following states are supported for modify:

- *enabled* - the feature is enabled.
- *disabled* - the feature is disabled.
- *dry-run* - the feature is in evaluation mode.
- *paused* - the feature is in pause mode.

`[-granular-data {disabled|basic}] - Granular data`

This parameter specifies the type of granular data storage supported on the volume. Once this property is changed to any value other than *disabled* on a volume, it cannot be disabled on that volume using this parameter. It can only be disabled on that volume by restoring a Snapshot copy. This property is only supported on FlexGroup volumes. Files modified by the granular data algorithms use granular data, which might impact read/write I/O performance.

The available options are:

- `disabled` - No volume capacity rebalancing operations are performed.
- `basic` - Volume capacity rebalancing operations are enabled.

`[-atime-update-period <integer>] - Access Time Update Period (Seconds) (privilege: advanced)`

This parameter specifies the time period that must be elapsed in between consecutive atime update events. Note: The value of this parameter is only used if `-atime-update` is enabled on the volume.

`[-snapshot-locking-enabled {true|false}] - Enable Snapshot Copy Locking`

This parameter specifies whether Snapshot copy locking is enabled, based on an independent compliance-clock time. Once this property is enabled, it cannot be disabled and the volume cannot be deleted until all locked Snapshot copies are past their expiry time.

`[-is-large-size-enabled {true|false}] - Are Large Size Volumes and Files Enabled`

This parameter specifies whether support for large Flexible volumes and large files is enabled. If this value is true, the maximum volume size for a Flexible volume will be 300TB and the maximum size of a single file will be 128TB. If this value is false, the maximum volume size for a Flexible volume will be 100TB and the maximum size of a single file will be 16TB. The default value is false.

`[-is-preserve-unlink-enabled {true|false}] - Is Preserve Unlink Enabled (privilege: advanced)`

This parameter specifies whether preserve unlink is enabled for NFSV4.1 shares on the volume. Once this

feature is enabled, a file removal op on a file with existing share locks will result in the file being moved to the trash directory. Clients will continue to be able to access the file using its file handle. The last close for that file will result in deletion of the file. The default value is false.

[`-is-cloud-write-enabled {true|false}`] - Is Cloud Write Enabled (privilege: advanced)

This parameter specifies whether cloud write is enabled on the volume. This feature is only available on volumes in FabricPools. The default value is false.

[`-aggressive-readahead-mode {none|file_prefetch}`] - Aggressive readahead mode (privilege: advanced)

This parameter specifies the aggressive readahead mode of the volume. When set to `file_prefetch`, on a file read, the system aggressively issues readaheads for all of the blocks in the file and retains those blocks in a cache for a finite period of time. This feature is only available on FabricPool volumes on FSx for ONTAP and Cloud Volumes ONTAP. The default value is `none`.

Examples

The following example modifies a volume named `vol4` on a Vserver named `vs0`. The volume's export policy is changed to `default_expolicy` and its size is changed to 500 GB.

```
cluster1::> volume modify -vserver vs0 -volume vol4 -policy
default_expolicy -size 500g
```

The following example modifies a volume named `vol2`. It enables `autogrow` and sets the maximum `autosize` to 500g

```
cluster1::> volume modify -volume vol2 -autosize-mode grow -max-autosize
500g
```

The following example modifies a volume named `vol2` to have a space guarantee of `none`.

```
cluster1::> volume modify -space-guarantee none -volume vol2
```

The following example modifies all volumes in Vserver `vs0` to have a fractional reserve of 30%.

```
cluster1::> volume modify -fractional-reserve 30 -vserver vs0 *
```

The following example modifies a volume named `vol2` to grow in size by 5 gigabytes

```
cluster1::> volume modify -volume vol2 -size +5g
```

The following example modifies a volume named `vol2` to have a different caching policy. The volume must be on a Flash Pool aggregate.

```
cluster1::> volume modify -volume vol2 -caching-policy none
```

Related Links

- [volume rename](#)
- [volume make-vsroot](#)
- [vserver export-policy create](#)
- [vserver create](#)
- [vserver modify](#)

volume mount

Mount a volume on another volume with a junction-path

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume mount` command mounts a volume at a specified junction path.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This specifies the volume that is to be mounted.

-junction-path <junction path> - Junction Path Of The Mounting Volume

This specifies the junction path of the mounted volume. The junction path name is case insensitive and must be unique within a Vserver's namespace.

[-active {true|false}] - Activate Junction Path

This optionally specifies whether the mounted volume is accessible. The default setting is `false`. If the mounted path is not accessible, it does not appear in the Vserver's namespace.

[-policy-override {true|false}] - Override The Export Policy

This optionally specifies whether the parent volume's export policy overrides the mounted volume's export policy. The default setting is `false`.

Examples

The following example mounts a volume named `user_tsmith` on a Vserver named `vs0`. The junction path for the mounted volume is `/user/tsmith`. The mounted volume is accessible, and the mounted volume's export policy is not overridden by the parent volume's export policy.

```
node::> volume mount -vserver vs0 -volume user_tsmith
-junction-path /user/tsmith -active true -policy-override false
```

volume offline

Take an existing volume offline

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume offline` command takes the volume offline. If the volume is already in restricted or `iron_restricted` state, then it is already unavailable for data access, and much of the following description does not apply. The current root volume might not be taken offline. Several operations being performed on the volume in question can prevent the volume offline command from succeeding for various lengths of time. If such operations are required, the command might take additional time to complete. If you set the volume junction path, it is unmounted when the volume is taken offline. The junction path is not restored when the volume is brought online again. When the volume offline command fails, the command is aborted. The `-force` flag can be used to forcibly offline a volume.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the name of the Vserver from which the volume is to be taken offline. If only one data Vserver exists, you do not need to specify this parameter.

-volume <volume name> - Volume Name

This specifies the name of the volume that is to be taken offline.

[-f, -force <true>] - Force Offline

This specifies whether the offline operation is forced. Using this option to force a volume offline can potentially disrupt access to other volumes. The default setting is `false`.

[-foreground {true|false}] - Foreground Process

This specifies whether the operation runs in the foreground. The default setting is `true` (the operation runs in the foreground). When set to `true`, the command will not return until the operation completes. This parameter applies only to FlexGroups. For FlexVol volumes, the command always runs in the foreground.

[-disable-block-storage-check <true>] - Disable Check for Existing LUNs and NVMe Namespaces

Taking the volume offline will make all associated LUNs and NVMe over Fabrics (NVMe-oF) namespaces unavailable, which normally requires a user confirmation. If this parameter is specified, the command proceeds without a confirmation. The default setting is `false`.

Examples

The following example takes the volume named `vol1` offline:

```
cluster1::> volume offline voll
Volume 'vs1:voll' is now offline.
```

volume online

Bring an existing volume online

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume online` command brings the volume online. A volume can only be brought online if it is offline or restricted. If the volume is inconsistent but has not lost data, the user will be cautioned and prompted before bringing it online. It is advisable to run `wafl-iron` (or do a `snapmirror initialize` in case of a replica volume) prior to bringing an inconsistent volume online. Bringing an inconsistent volume online increases the risk of further file system corruption. If the containing aggregate cannot honor the space guarantees required by this volume, the volume online operation will fail. It is not advisable to use volumes with their space guarantees disabled. Lack of free space can lead to failure of writes which in turn can appear as data loss to some applications.

Parameters

-vserver <vserver name> - Vserver Name

This parameter specifies the name of the Vserver from which the volume is to be brought online. If only one data Vserver exists, you do not need to specify this parameter.

-volume <volume name> - Volume Name

This parameter specifies the name of the volume that is to be brought online.

[-f, -force <true>] - Force Online

When this parameter is used, the volume will be brought online even if there is not enough space available in the aggregate to honor the volume's space guarantee.

[-foreground {true|false}] - Foreground Process

This parameter specifies whether the operation runs in the foreground. The default setting is `true` (the operation runs in the foreground). When set to `true`, the command will not return until the operation completes. This parameter applies only to FlexGroups. For FlexVol volumes, the command always runs in the foreground.

Examples

The following example brings a volume named `vol1` online:

```
cluster1::> volume online voll
Volume 'vs1:voll' is now online.
```

volume rehost

Rehost a volume from one Vserver into another Vserver

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `volume rehost` command rehosts a volume from source Vserver onto destination Vserver. The volume name must be unique among the other volumes on the destination Vserver. This command is not supported in a MetroCluster configuration.

Parameters

-vserver <vserver name> - Source Vserver name

This specifies the Vserver on which the volume is located.

-volume <volume name> - Target volume name

This specifies the volume that is to be rehosted.

-destination-vserver <vserver name> - Destination Vserver name

This specifies the destination Vserver where the volume must be located post rehost operation.

{ [-force-unmap-luns {true|false}] - Unmap LUNs in volume

This specifies whether the rehost operation should unmap LUNs present on volume. The default setting is *false* (the rehost operation shall not unmap LUNs). When set to true, the command will unmap all mapped LUNs on the volume.

| [-auto-remap-luns {true|false}] - Automatic Remap of LUNs }

This specifies whether the rehost operation should perform LUN mapping operation at the destination Vserver for the LUNs mapped on the volume at the source Vserver. The default setting is *false* (the rehost operation shall not map LUNs at the destination Vserver). When set to true, at the destination Vserver the command will create initiator groups along with the initiators (if present) with same name as that of source Vserver. Then the LUNs on the volume are mapped to initiator groups at the destination Vserver as mapped in source Vserver.

[-key-manager-attribute <text>] - Optional CRN

This optional parameter specifies an additional key manager attribute that should be an identifier-value pair separated by '=', ex. CRN=my-unique-value

The following identifiers are currently supported:

- CRN

Examples

The following example rehosts a volume named vol3 on Vserver named vs1 to a destination Vserver named vs2:

```
cluster::> volume rehost -vserver vs1 -volume vol3 -destination-vserver
vs2
```

volume rename

Rename an existing volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume rename` command renames a volume. The volume name must be unique among the other volumes in the same Vserver.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver on which the volume is located. For a node's root volume, use the name of the node for this parameter.

-volume <volume name> - Volume Name

This specifies the volume that is to be renamed.

-newname <volume name> - Volume New Name

This specifies the volume's new name. A volume's name must start with an alphabetic character (a to z or A to Z) and be 203 or fewer characters in length.

[-foreground {true|false}] - Foreground Process

This specifies whether the operation runs in the foreground. The default setting is `true` (the operation runs in the foreground). When set to `true`, the command will not return until the operation completes. This parameter applies only to FlexGroups. For FlexVol volumes, the command always runs in the foreground.

Examples

The following example renames a volume named `vol3_backup` as `vol3_save` on a Vserver named `vs0`:

```
node::> volume rename -vserver vs0 -volume vol3_backup -newname vol3_save
```

volume restrict

Restrict an existing volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume restrict` command puts the volume in restricted state. If the volume is online, then it will be made unavailable for data access as described under [volume offline](#) .

Parameters

-vserver <vserver name> - Vserver Name

This specifies the name of the Vserver from which the volume is to be restricted. If only one data Vserver exists, you do not need to specify this parameter.

-volume <volume name> - Volume Name

This specifies the name of the volume that is to be restricted.

[-foreground {true|false}] - Foreground Process

This specifies whether the operation runs in the foreground. The default setting is `true` (the operation runs in the foreground). When set to `true`, the command will not return until the operation completes. This parameter applies only to FlexGroups. For FlexVol volumes, the command always runs in the foreground.

Examples

The following example restricts a volume named `vol1`:

```
cluster1::> volume restrict vol1
Volume 'vs1:vol1' is now restricted.
```

Related Links

- [volume offline](#)

volume show-footprint

Display a list of volumes and their data and metadata footprints in their associated aggregate.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume show-footprint` command displays information about the space used in associated aggregates by volumes and features enabled in volumes. The command output depends on the parameter or parameters specified with the command. If no parameters are specified, the command displays the following information about all volumes. + + If the associated aggregates have an object-store attached to them, then the command displays a more detailed split up of the space used in each tier. This additional information is useful to show per-tier space usage which can be used to estimate the space requirements and transfer duration when moving a volume to a different tier with `volume move` .

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver

If this parameter and the `-volume` parameter are specified, the command displays detailed information about the specified volume. If this parameter is specified by itself, the command displays information about volumes on the specified Vserver.

[-volume <volume name>] - Volume Name

If this parameter and the `-vserver` parameter are specified, the command displays detailed information about the specified volume. If this parameter is specified by itself, the command displays information about all volumes matching the specified name.

[-volume-msid <integer>] - Volume MSID

If this parameter is specified, the command displays information only about the volume that has the specified MSID.

[-volume-dsid <integer>] - Volume DSID

If this parameter is specified, the command displays information only about the volume that has the specified DSID.

[-vserver-uuid <UUID>] - Vserver UUID

If this parameter is specified, the command displays information only about the volume on the vserver which has the specified UUID.

[-aggregate <aggregate name>] - Aggregate Name

If this parameter is specified, the command displays information only about the volumes that are associated with the specified aggregate.

[-aggregate-uuid <UUID>] - Aggregate UUID

If this parameter is specified, the command displays information only about the volumes on the aggregate which have the specified UUID.

[-hostname <text>] - Hostname

If this parameter is specified, the command displays information only about the volumes that belong to the specified host.

[-tape-backup-metafiles-footprint {<integer>[KB|MB|GB|TB|PB] }] - Tape Backup Metadata Footprint

If this parameter is specified, the command displays information only about the volumes whose tape backup metafiles use the specified amount of space in the aggregate.

[-tape-backup-metafiles-footprint-percent <percent>] - Tape Backup Metadata Footprint Percent

If this parameter is specified, the command displays information only about the volumes whose tape backup metafiles use the specified percentage of space in the aggregate.

`[-dedupe-metafiles-footprint {<integer>[KB|MB|GB|TB|PB] }] - Deduplication Footprint`

If this parameter is specified, the command displays information only about the volumes whose deduplication metafiles use the specified amount of space in the aggregate.

`[-dedupe-metafiles-footprint-percent <percent>] - Deduplication Footprint Percent`

If this parameter is specified, the command displays information only about the volumes whose deduplication metafiles use the specified percentage of space in the aggregate.

`[-dedupe-metafiles-temporary-footprint {<integer>[KB|MB|GB|TB|PB] }] - Temporary Deduplication Footprint`

If this parameter is specified, the command displays information only about the volumes whose temporary deduplication metafiles use the specified amount of space in the aggregate.

`[-dedupe-metafiles-temporary-footprint-percent <percent>] - Temporary Deduplication Footprint Percent`

If this parameter is specified, the command displays information only about the volumes whose temporary deduplication metafiles use the specified percentage of space in the aggregate.

`[-cross-volume-dedupe-metafiles-footprint {<integer>[KB|MB|GB|TB|PB] }] - Cross Volume Deduplication Footprint`

If this parameter is specified, the command displays information only about the volumes whose cross volume deduplication metafiles use the specified amount of space in the aggregate.

`[-cross-volume-dedupe-metafiles-footprint-percent <percent>] - Cross Volume Deduplication Footprint Percent`

If this parameter is specified, the command displays information only about the volumes whose cross volume deduplication metafiles use the specified percentage of space in the aggregate.

`[-cross-volume-dedupe-metafiles-temporary-footprint {<integer>[KB|MB|GB|TB|PB] }] - Cross Volume Temporary Deduplication Footprint`

If this parameter is specified, the command displays information only about the volumes whose cross volume deduplication temporary metafiles use the specified amount of space in the aggregate.

`[-cross-volume-dedupe-metafiles-temporary-footprint-percent <percent>] - Cross Volume Temporary Deduplication Footprint Percent`

If this parameter is specified, the command displays information only about the volumes whose cross volume deduplication temporary metafiles use the specified percentage of space in the aggregate.

`[-volume-blocks-footprint {<integer>[KB|MB|GB|TB|PB] }] - Volume Data Footprint`

If this parameter is specified, the command displays information only about the volumes whose data blocks use the specified amount of space in the aggregate.

This field is the total amount of data written to the volume. It includes data in the active file system in the volume as well as data that is consumed by volume Snapshot copies. This row only includes data and not reserved space, so when volumes have reserved files, the volume's total used in the [volume show-space](#) command output can exceed the value in this row.

[`-volume-blocks-footprint-percent` <percent_no_limit>] - Volume Data Footprint Percent

If this parameter is specified, the command displays information only about the volumes whose data blocks use the specified percentage of space in the aggregate.

[`-flexvol-metadata-footprint` {<integer>[KB|MB|GB|TB|PB]}] - Flexible Volume Metadata Footprint

If this parameter is specified, the command displays information only about the volumes whose file system metadata uses the specified amount of space in the aggregate.

This field includes the space used or reserved in the aggregate for metadata associated with this volume.

[`-flexvol-metadata-footprint-percent` <percent>] - Flexible Volume Metadata Footprint Percent

If this parameter is specified, the command displays information only about the volumes whose file system metadata uses the specified percentage of space in the aggregate.

[`-delayed-free-footprint` {<integer>[KB|MB|GB|TB|PB]}] - Delayed Free Blocks

If this parameter is specified, the command displays information only about the volumes whose delayed free blocks use the specified amount of space in the aggregate.

When Data ONTAP frees space in a volume, this space is not always immediately shown as free in the aggregate. This is because the operations to free the space in the aggregate are batched for increased performance. Blocks that are declared free in the FlexVol volume but which are not yet free in the aggregate are called "delayed free blocks" until the associated delayed free blocks are processed. For SnapMirror destinations, this row will have a value of 0 and will not be displayed.

[`-delayed-free-footprint-percent` <percent>] - Delayed Free Blocks Percent

If this parameter is specified, the command displays information only about the volumes that have the specified amount of blocks waiting to be freed in the aggregate. This space is called "delayed free blocks".

[`-snapmirror-destination-footprint` {<integer>[KB|MB|GB|TB|PB]}] - SnapMirror Destination Footprint

If this parameter is specified, the command displays information only about the volumes whose SnapMirror transfer uses the specified amount of space in the aggregate.

During a SnapMirror transfer, this row will include incoming SnapMirror data and SnapMirror-triggered delayed free blocks from previous SnapMirror transfers.

[`-snapmirror-destination-footprint-percent` <percent>] - SnapMirror Destination Footprint Percent

If this parameter is specified, the command displays information only about the volumes whose SnapMirror transfer uses the specified percentage of space in the aggregate.

[`-volume-guarantee-footprint` {<integer>[KB|MB|GB|TB|PB]}] - Volume Guarantee

If this parameter is specified, the command displays information only about the volumes whose guarantees use the specified amount of space in the aggregate.

This field includes the amount of space reserved by this volume in the aggregate for future writes. The amount of space reserved depends on the guarantee type (the provisioning mode) of the volume. + For a "volume" guaranteed volume, this is the size of the volume minus the amount in the Volume Data Footprint row. + For a "file" guaranteed volume, this is the sum of all of the space reserved for hole fills and overwrites in all of the space reserved files in the volume.

[-volume-guarantee-footprint-percent <percent>] - Volume Guarantee Percent

If this parameter is specified, the command displays information only about the volumes whose guarantees use the specified percentage of space in the aggregate.

[-file-operation-metadata {<integer>[KB|MB|GB|TB|PB]}] - File Operation Metadata

If this parameter is specified, the command displays information only about the volumes that have file operation metadata using the specified amount of space in the aggregate.

This metadata is used by file move and copy operations. Although it is not returned as free space once the operations are complete, it can be reused by future file operations.

[-file-operation-metadata-percent <percent>] - File Operation Metadata Percent

If this parameter is specified, the command displays information only about the volumes that have file operation metadata using the specified percentage of space in the aggregate.

[-total-footprint {<integer>[KB|MB|GB|TB|PB]}] - Total Footprint

If this parameter is specified, the command displays information only about the volumes which use the specified amount of space in the aggregate. This field is the sum of the other rows in this table.

[-total-footprint-percent <percent_no_limit>] - Total Footprint Percent

If this parameter is specified, the command displays information only about the volumes which use the specified percentage of space in the aggregate.

[-aggregate-size {<integer>[KB|MB|GB|TB|PB]}] - Containing Aggregate Size

If this parameter is specified, the command displays information only about the volumes that are associated with an aggregate of the specified size.

[-bin0-name <text>] - Name for bin0

If this parameter is specified, the command displays information only about volumes whose associated aggregate has an object store attached to it with an active file system tier name matches the specified value.

[-volume-blocks-footprint-bin0 {<integer>[KB|MB|GB|TB|PB]}] - Volume Footprint for bin0

If this parameter is specified, the command displays information only about volumes whose space in use in the performance tier of the aggregate matches the specified value.

[-volume-blocks-footprint-bin0-percent <percent>] - Volume Footprint bin0 Percent

If this parameter is specified, the command displays information only about volumes whose percentage space in use in the performance tier of the aggregate matches the specified value.

[-bin1-name <text>] - Name for bin1

If this parameter is specified, the command displays information only about volumes whose associated aggregate has a cloud tier attached to it with a configuration name that matches the specified value.

[-volume-blocks-footprint-bin1 {<integer>[KB|MB|GB|TB|PB]}] - Volume Footprint for bin1

If this parameter is specified, the command displays information only about volumes whose space in use in the first cloud tier bucket attached to the associated aggregate matches the specified value. This includes the space used by the blocks in the volume that are staged before being moved to the cloud tier.

[`-volume-blocks-footprint-bin1-percent <percent>`] - Volume Footprint bin1 Percent

If this parameter is specified, the command displays information only about volumes whose percentage space in use in the first cloud tier bucket attached to associated aggregate matches the specified value. This includes the space used by the blocks in the volume that are staged before being moved to the cloud tier.

[`-total-dedupe-metafiles-footprint {<integer>[KB|MB|GB|TB|PB]}`] - Total Deduplication Footprint

If this parameter is specified, the command displays information only about the volumes whose total deduplication metafiles use the specified amount of space in the aggregate.

[`-total-dedupe-metafiles-footprint-percent <percent>`] - Total Deduplication Footprint Percent

If this parameter is specified, the command displays information only about the volumes whose total deduplication metafiles use the specified percentage of space in the aggregate.

[`-auto-adaptive-compression-footprint-data-reduction {<integer>[KB|MB|GB|TB|PB]}`] - Footprint Data Reduction by Auto Adaptive Compression

If this parameter is specified, the command displays information only about the volumes whose estimation of footprint data reduction due to auto adaptive compression use the specified amount of space in the aggregate.

[`-auto-adaptive-compression-footprint-data-reduction-percent <percent>`] - Footprint Data Reduction by Auto Adaptive Compression Percent

If this parameter is specified, the command displays information only about the volumes whose percentage estimation of footprint data reduction due to auto adaptive compression use the specified amount of space in the aggregate.

[`-total-footprint-data-reduction {<integer>[KB|MB|GB|TB|PB]}`] - Total Footprint Data Reduction

If this parameter is specified, the command displays information only about the total footprint data reduction by the volume in the aggregate.

[`-total-footprint-data-reduction-percent <percent>`] - Total Footprint Data Reduction Percent

If this parameter is specified, the command displays information only about the percentage of total footprint data reduction by the volume in the aggregate.

[`-capacity-tier-footprint-data-reduction {<integer>[KB|MB|GB|TB|PB]}`] - Footprint Data Reduction by Capacity Tier

If this parameter is specified, the command displays information only about the footprint data reduction in the cloud tier by the volume in a FabricPool.

[`-capacity-tier-footprint-data-reduction-percent <percent>`] - Footprint Data Reduction by Capacity Tier Percent

If this parameter is specified, the command displays information only about the percentage of footprint data reduction in the cloud tier bucket by the volume in a FabricPool.

[`-effective-total-footprint {<integer>[KB|MB|GB|TB|PB]}`] - Effective Total after Footprint Data Reduction

If this parameter is specified, the command displays information only about the total space used by the volume in the aggregate after total footprint data reduction.

[-effective-total-footprint-percent <percent>] - Effective Total after Footprint Data Reduction Percent

If this parameter is specified, the command displays information only about the percentage of total space used by the volume in the aggregate after total footprint data reduction.

[-compaction-footprint-data-reduction {<integer>[KB|MB|GB|TB|PB]}] - Footprint Data Reduction by Compaction

If this parameter is specified, the command displays information only about the volumes whose estimation of footprint data reduction due to compaction use the specified amount of space in the aggregate.

[-compaction-footprint-data-reduction-percent <percent>] - Footprint Data Reduction by Compaction Percent

If this parameter is specified, the command displays information only about the volumes whose percentage estimation of footprint data reduction due to compaction use the specified amount of space in the aggregate.

[-total-metadata-footprint {<integer>[KB|MB|GB|TB|PB]}] - Total Metadata Footprint

If this parameter is specified, the command displays information only about volume that have the specified amount of space used by volume metadata in the aggregate. The value of this field will be higher than the field 'total-metadata' from the [volume show-space](#) command output since this includes additional metadata of the volume in the aggregate.

[-total-metadata-footprint-percent <percent>] - Total Metadata Footprint Percentage

.If this parameter is specified, the command displays information only about the volume that have the specified percentage of space used by volume metadata in the aggregate.

Examples

The following example displays information about all volumes in the system.

```

cluster1::> volume show-footprint
Vserver : nodevs
  Volume : vol0
Feature                                     Used      Used%
-----
Volume Data Footprint                       103.1MB   11%
Volume Guarantee                            743.6MB   83%
Flexible Volume Metadata                     4.84MB    1%
Delayed Frees                               4.82MB    1%
Total Metadata Footprint                    10.2MB    2%
Total Footprint                             856.3MB   95%
Vserver : thevs
  Volume : therootvol
Feature                                     Used      Used%
-----
Volume Data Footprint                       116KB     0%
Volume Guarantee                            19.83MB   1%
Flexible Volume Metadata                     208KB     0%
Delayed Frees                               60KB      0%
Total Metadata Footprint                    320KB     0%
Total Footprint                             20.20MB   1%
Vserver : thevs
  Volume : thevol
Feature                                     Used      Used%
-----
Volume Data Footprint                       128KB     0%
Volume Guarantee                            2.00GB    76%
Flexible Volume Metadata                     11.38MB   0%
Delayed Frees                               428KB     0%
Total Metadata Footprint                    13.52MB   0%
Total Footprint                             2.01GB    76%

3 entries were displayed.

```

The following example displays information about all volumes in a system and highlights a scenario where the aggregates associated with volumes have a cloud tier attached to them.

```

cluster-1::> vol show-footprint
Vserver : vsim1
  Volume : vol0
Feature                                     Used      Used%
-----
Volume Data Footprint                       2.57GB    81%
Volume Guarantee                            266.1MB   8%
Flexible Volume Metadata                     16.23MB   0%

```

```

    Delayed Frees                27.97MB    1%
    Total Metadata Footprint      48.52MB    1%
Total Footprint                2.87GB    90%
Vserver : vs1
  Volume : svm_root
Feature                          Used      Used%
-----
Volume Data Footprint            2.93MB    0%
  Footprint in Performance Tier  2.99MB   100%
  Footprint in my-store          0B        0%
Volume Guarantee                 17.01MB   0%
Flexible Volume Metadata         208KB     0%
Delayed Frees                    68KB     0%
Total Metadata Footprint        320KB     0%
Total Footprint                 20.20MB   0%
Vserver : vs1
  Volume : vol1
Feature                          Used      Used%
-----
Volume Data Footprint            1.61GB   17%
  Footprint in Performance Tier  1.23GB   72%
  Footprint in my-store          479.0MB  28%
Volume Guarantee                 0B        0%
Flexible Volume Metadata         16.06MB   0%
Delayed Frees                    82.98MB   1%
Total Metadata Footprint        103.52MB  1%
Total Footprint                 1.71GB   18%
Vserver : vs1
  Volume : vol2
Feature                          Used      Used%
-----
Volume Data Footprint            1.22GB   13%
  Footprint in Performance Tier  823.3MB  65%
  Footprint in msl              440MB    35%
Volume Guarantee                 0B        0%
Flexible Volume Metadata         16.06MB   0%
Delayed Frees                    12MB     0%
Total Metadata Footprint        33.52MB   0%
Total Footprint                 1.25GB   13%

4 entries were displayed.

```

Related Links

- [volume show-space](#)

volume show-space

Display space usage for volume(s)

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume show-space` command displays information about space usage within the volume. The command output depends on the parameter or parameters specified with the command. If no parameters are specified, the command displays the following information about all volumes.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver

If this parameter and the `-volume` parameter are specified, the command displays detailed information about the specified volume. If this parameter is specified by itself, the command displays information about volumes on the specified Vserver.

[-volume <volume name>] - Volume Name

If this parameter and the `-vserver` parameter are specified, the command displays detailed information about the specified volume. If this parameter is specified by itself, the command displays information about all volumes matching the specified name.

[-volume-msid <integer>] - Volume MSID

If this parameter is specified, the command displays information only about the volume that has the specified MSID.

[-volume-dsid <integer>] - Volume DSID

If this parameter is specified, the command displays information only about the volume that has the specified DSID.

[-vserver-uuid <UUID>] - Vserver UUID

If this parameter is specified, the command displays information only about the volume on the vserver which has the specified UUID.

[-aggregate <aggregate name>] - Aggregate Name

If this parameter is specified, the command displays information only about the volumes that are associated with the specified aggregate.

[-aggregate-uuid <UUID>] - Aggregate UUID

If this parameter is specified, the command displays information only about the volumes on the aggregate which have the specified UUID.

[-hostname <text>] - Hostname

If this parameter is specified, the command displays information only about the volumes that belong to the specified host.

[-user-data {<integer>[KB|MB|GB|TB|PB] }] - User Data

If this parameter is specified, the command displays information only about the volume or volumes that have the specified amount of space in use by user data in the volume.

This is the amount of data written to the volume via CIFS, NFS or SAN protocols plus the space reserved in the volume for these files (hole and overwrite reserves). This is the same information displayed by running the Unix `du` command on the mount point.

[-user-data-percent <percent_no_limit>] - User Data Percent

If this parameter is specified, the command displays information only about the volume or volumes that have the specified percentage of space in use by user data in the volume.

[-dedupe-metafiles {<integer>[KB|MB|GB|TB|PB] }] - Deduplication

If this parameter is specified, the command displays information only about the volume or volumes that have the specified amount of space in use by deduplication metafiles in the volume.

[-dedupe-metafiles-percent <percent>] - Deduplication Percent

If this parameter is specified, the command displays information only about the volume or volumes that have the specified percentage of space in use by deduplication metafiles in the volume.

[-dedupe-metafiles-temporary {<integer>[KB|MB|GB|TB|PB] }] - Temporary Deduplication

If this parameter is specified, the command displays information only about the volume or volumes that have the specified amount of space in use by temporary deduplication metafiles in the volume.

[-dedupe-metafiles-temporary-percent <percent>] - Temporary Deduplication Percent

If this parameter is specified, the command displays information only about the volume or volumes that have the specified percentage of space in use by temporary deduplication metafiles in the volume.

[-filesystem-metadata {<integer>[KB|MB|GB|TB|PB] }] - Filesystem Metadata

If this parameter is specified, the command displays information only about the volume or volumes that have the specified amount of space in use by file system metadata in the volume.

[-filesystem-metadata-percent <percent>] - Filesystem Metadata Percent

If this parameter is specified, the command displays information only about the volume or volumes that have the specified percentage of space in use by file system metadata in the volume.

[-snapmirror-metadata {<integer>[KB|MB|GB|TB|PB] }] - SnapMirror Metadata

If this parameter is specified, the command displays information only about the volume or volumes that have the specified amount of space in use by SnapMirror metafiles in the volume.

Between SnapMirror transfers, some metadata is maintained to support storage-efficient transfers. During transfers, some additional space is used temporarily. This space is used in all SnapMirror destination volumes.

[-snapmirror-metadata-percent <percent>] - SnapMirror Metadata Percent

If this parameter is specified, the command displays information only about the volume or volumes that have the specified percentage of space in use by SnapMirror metafiles inside the volume.

[`--tape-backup-metadata` {<integer>[KB|MB|GB|TB|PB]}] - Tape Backup Metadata

If this parameter is specified, the command displays information only about the volume or volumes that have the specified amount of space in use by tape backup metafiles in the volume.

[`--tape-backup-metadata-percent` <percent>] - Tape Backup Metadata Percent

If this parameter is specified, the command displays information only about the volume or volumes that have the specified percentage of space in use by tape backup metafiles in the volume.

[`--quota-metafiles` {<integer>[KB|MB|GB|TB|PB]}] - Quota Metadata

If this parameter is specified, the command displays information only about the volume or volumes that have the specified amount of space in use by quota metafiles.

[`--quota-metafiles-percent` <percent>] - Quota Metadata Percent

If this parameter is specified, the command displays information only about the volume or volumes that have the specified percentage of space in use by quota metafiles.

[`--inodes` {<integer>[KB|MB|GB|TB|PB]}] - Inodes

If this parameter is specified, the command displays information only about the volume or volumes that have the specified amount of space in use by the inode metafile in the volume.

This is the amount of space required to store inodes in the file system and is proportional to the maximum number of files ever created in the volume. The inode file is not compacted or truncated, so if a large number of files are created and then deleted, the inode file does not shrink.

[`--inodes-percent` <percent>] - Inodes Percent

If this parameter is specified, the command displays information only about the volume or volumes that have the specified percentage of space in use by the inode metafile in the volume.

[`--inodes-upgrade` {<integer>[KB|MB|GB|TB|PB]}] - Inodes Upgrade

If this parameter is specified, the command displays information only about the volume or volumes that have the specified amount of space in use by the inode subsystem for the purpose of upgrading.

This is the amount of space required to store upgrading inodes in the file system and is proportional to the size of the inode metafile. Once the upgrade is complete, the space used by 'inodes' will be replaced with the space used for upgrade.

[`--inodes-upgrade-percent` <percent>] - Inodes Upgrade Percent

If this parameter is specified, the command displays information only about the volume or volumes that have the specified percentage of space in use for upgrading the inode metafile in the volume.

[`--snapshot-reserve` {<integer>[KB|MB|GB|TB|PB]}] - Snapshot Reserve

If this parameter is specified, the command displays information only about the volume or volumes that have the specified amount of space in use by the Snapshot reserve in the volume.

[`--snapshot-reserve-percent` <percent>] - Snapshot Reserve Percent

If this parameter is specified, the command displays information only about the volume or volumes that have the specified percentage of space in use by the Snapshot reserve in the volume.

[`--snapshot-reserve-unusable` {<integer>[KB|MB|GB|TB|PB]}] - Snapshot Reserve Unusable

If this parameter is specified, the command displays information only about the volume or volumes that have the specified amount of space reserved but unusable in the volume.

Snapshot reserve can be diminished under certain conditions to accommodate volume metadata. Creating space in the volume will make this space available.

[-snapshot-reserve-unusable-percent <integer>] - Snapshot Reserve Unusable Percent

If this parameter is specified, the command displays information only about the volume or volumes that have the specified percentage of space reserved but unusable.

[-snapshot-spill {<integer>[KB|MB|GB|TB|PB]}] - Snapshot Spill

If this parameter is specified, the command displays information only about the volume or volumes that have the specified amount of space in use by their Snapshot spill.

If Snapshot used space exceeds the Snapshot reserve it is considered to spill out of the reserve. This space cannot be used by the active file system until Snapshots are deleted.

[-snapshot-spill-percent <percent>] - Snapshot Spill Percent

If this parameter is specified, the command displays information only about the volume or volumes that have the specified percentage of space in use by the Snapshot spill.

[-performance-metadata {<integer>[KB|MB|GB|TB|PB]}] - Performance Metadata

If this parameter is specified, the command displays information only about the volume or volumes that have the specified amount of space in use for performance optimization in the volume.

[-performance-metadata-percent <percent>] - Performance Metadata Percent

If this parameter is specified, the command displays information only about the volume or volumes that have the specified percentage of space in use for performance optimization in the volume.

[-total-used {<integer>[KB|MB|GB|TB|PB]}] - Total Used

If this parameter is specified, the command displays information only about the volume or volumes that have the specified amount of space in use by the volume, including the space used by the Snapshot reserve. *total-used* differs from *physical-used* space by the sum of the space that is reserved for future writes and the space that is saved by aggregate storage efficiency savings being removed from *total-used*.

This is equivalent to the used field in the output of the [volume show](#) command.

[-total-used-percent <percent_no_limit>] - Total Used Percent

If this parameter is specified, the command displays information only about the volume or volumes that have the specified percentage of space in use by the volume, including the space used by the Snapshot reserve.

[-physical-used {<integer>[KB|MB|GB|TB|PB]}] - Total Physical Used Size

If this parameter is specified, the command displays information only about the volume or volumes that have the specified amount of physical space in use by the volume.

physical-used differs from *total-used* space by the sum of the space that is reserved for future writes and the space that is saved saved by aggregate storage efficiency savings being added to *physical-used*. The value includes blocks in use by Snapshot copies.

[-physical-used-percent <percent_no_limit>] - Physical Used Percentage

If this parameter is specified, the command displays information only about the volume or volumes that have the specified percentage of physical space in use in the volume based on volume size including the

space reserved for Snapshot copies.

[`-logical-used` {<integer>[KB|MB|GB|TB|PB]}] - Logical Used Size

If this parameter is specified, the command displays information only about the volume or volumes that have the specified amount of logical space in use by the volume. This includes space saved by all the storage efficiency features along with physical used space. This does not include Snapshot reserve but does consider Snapshot spill.

[`-logical-used-percent` <percent_no_limit>] - Logical Used Percent

If this parameter is specified, the command displays information only about the volume or volumes that have the specified percentage of logical space used in the volume.

[`-logical-available` {<integer>[KB|MB|GB|TB|PB]}] - Logical Available

If this parameter is specified, the command displays information only about the volume or volumes that have the specified amount of logical available space in the volume.

[`-total-metadata` {<integer>[KB|MB|GB|TB|PB]}] - Total Metadata

If this parameter is specified, the command displays information only about the volume or volumes that have the specified amount of space used by metadata in the volume. This field will not include the metadata of the volume which is residing inside the aggregate.

[`-total-metadata-percent` <percent>] - Total Metadata Percentage

If this parameter is specified, the command displays information only about the volume or volumes that have the specified percentage of space used by metadata in the volume.

Examples

The following example shows how to display details for all volumes.

```

cluster1::> volume show-space
Vserver : nodevs
  Volume : vol0
Feature                                     Used      Used%
-----
User Data                                   163.4MB   3%
Filesystem Metadata                         172KB    0%
Inodes                                      2.93MB   0%
Snapshot Reserve                           292.9MB  5%
Total Metadata      185KB      0%
Total Used          459.4MB   8%
Total Physical Used 166.4MB   3%
Vserver : thevs
  Volume : rootvol
Feature                                     Used      Used%
-----
User Data                                   100KB    0%
Filesystem Metadata                         76KB    0%
Inodes                                      24KB    0%
Snapshot Reserve                           1MB     5%
Total Metadata      89KB     0%
Total Used          1.20MB   6%
Total Physical Used 200KB    1%
Vserver : vs1
  Volume : vol1
Feature                                     Used      Used%
-----
User Data                                   180.8MB  74%
Filesystem Metadata                         280KB   0%
Inodes                                      12KB   0%
Snapshot Reserve                           12.20MB  5%
Total Metadata      295KB   0%
Total Used          193.3MB  79%
Total Physical Used 192.9MB  79%
3 entries were displayed.

```

The following example shows all Volumes that have a snap reserve greater than 2 MB:

```

cluster1::> volume show-space -snapshot-reserve >2m
Vserver : nodevs
  Volume : vol0
Feature                                     Used      Used%
-----
User Data                                   163.4MB   3%
Filesystem Metadata                        172KB    0%
Inodes                                     2.93MB   0%
Snapshot Reserve                          292.9MB  5%
Total Metadata                             3.21MB   0%
Total Used                                 459.4MB  8%
Total Physical Used                       166.4MB  3%
Vserver : vs1
  Volume : vol1
Feature                                     Used      Used%
-----
User Data                                   180.8MB  74%
Filesystem Metadata                        280KB    0%
Inodes                                     12KB     0%
Snapshot Reserve                          12.20MB  5%
Total Metadata                             312KB    0%
Total Used                                 193.3MB  79%
Total Physical Used                       192.9MB  79%
2 entries were displayed.

```

Related Links

- [volume show](#)

volume show

Display a list of volumes

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume show` command displays information about volumes. The command output depends on the parameter or parameters specified with the command. If no parameters are specified, the command displays the following information about all volumes:

- Vserver name
- Volume name
- Aggregate name
- State (online, offline, restricted, or mixed)

- Type (RW for read-write or DP for data-protection)
- Size
- Available size
- Percentage of space used

To display detailed information about a single volume, run the command with the `-vserver` and `-volume` parameters. The detailed view provides all of the information in the previous list and the following additional information:

- Name ordinal
- Volume data set ID
- Volume master data set ID
- Volume style (trad or flex)
- FlexCache Endpoint Type (none or cache or origin)
- Whether the volume is a cluster volume or node volume
- Export policy name
- User ID
- Group ID
- Security style (unix, ntfs, mixed or unified)
- UNIX permissions
- Junction path
- Junction path source
- Whether the junction path is active
- Parent volume name
- Vserver root volume
- Comment
- Filesystem size
- Total user-visible size
- Used size
- Used percentage
- Volume nearly full threshold percent
- Volume full threshold percent
- Autosize enabled
- Maximum autosize
- Minimum autosize
- Autosize grow threshold percent
- Autosize shrink threshold percent
- Autosize mode
- Total files

- Files used
- Expected available size
- Over provisioned size
- Snapshot reserve available size
- Logical used size
- Logical used percent
- Logical available size
- Active filesystem logical used size
- Snapshot copy logical used size
- Use logical space reporting
- Use logical space enforcement
- Maximum directory size
- Space guarantee style
- Whether a space guarantee is in effect
- Space SLO type (none, thick or semi-thick)
- Whether space SLO is in effect
- Whether minimum readahead is enabled
- Whether access time update is enabled
- Whether Snapshot copy directory access is enabled
- Percentage of space reserved for Snapshot copies
- Percentage of Snapshot copy space used
- Snapshot policy name
- Creation time
- If the filesystem size is fixed
- Overwrite reserve
- Fractional reserve
- Which space management strategy to try first
- Language
- Whether there's one data volume per member aggregate
- Concurrency level
- Optimization policy
- Whether the volume is a clone
- Volume UUID
- Whether failover is enabled
- Failover state
- (DEPRECATED)-Extent option
- Read reallocation option

- Consistency state
- Whether volume is quiesced on disk
- Whether volume is quiesced in memory
- Whether volume contains shared or compressed data
- Space saved by storage efficiency
- Percentage of space saved by storage efficiency
- Space saved by deduplication
- Percentage of space saved by deduplication
- Space shared by deduplication
- Space saved by compression
- Percentage of space saved by compression
- Volume size used by Snapshot copies
- Caching policy
- FlexGroup volume master data set ID
- FlexGroup volume index
- FlexGroup volume UUID
- Maximum size of the FlexGroup volume constituent
- Whether the volume has FlexGroup volume enabled
- Whether a FlexGroup volume is Qtree enabled
- Whether the volume is the destination of a move that is currently in cutover
- List of the aggregates used by the FlexGroup volume
- List of the nodes used by the FlexGroup volume
- SnapLock Type
- Is in pre-commit phase of copy-free Transition
- Application that the volume belongs to
- File system analytics state of the volume
- File system analytics scan progress percentage
- File system analytics number of files scanned
- File system analytics total number of files
- Activity tracking state of the volume
- Application UUID
- Whether the volume's total number of files is set to the highest possible value

To display detailed information about all volumes, run the command with the `-instance` parameter.

You can specify additional parameters to display information that matches only those parameters. For example, to display information only about data-protection volumes, run the command with the `-type DP` parameter.

Parameters

{ [-fields <fieldname>,...]

This specifies the fields that need to be displayed. The fields Vserver and policy are the default fields (see example).

| [-encryption]

If this parameter is specified, the command displays the following information:

- Vserver name
- Volume name
- Aggregate name
- Volume state
- Encryption state

| [-junction]

If this parameter is specified, the command displays the following information:

- Vserver name
- Volume name
- Whether the volume's junction is active
- Junction path
- Junction path source (if the volume is a mirror)

| [-settings] (privilege: advanced)

If this parameter is specified, the command displays the following information:

- Vserver name
- Volume name
- Whether minimum readahead is enabled on the volume
- Whether the access time is updated on inodes when a file on the volume is read
- Whether clients have access to .snapshot directories
- Whether automatic Snapshot copies are enabled on the volume

| [-instance] }

If this parameter is specified, the command displays information about all entries.

[-vserver <vserver name>] - Vserver Name

If this parameter and the `-volume` parameter are specified, the command displays detailed information about the specified volume. If this parameter is specified by itself, the command displays information about volumes on the specified Vserver.

[-volume <volume name>] - Volume Name

If this parameter and the `-vserver` parameter are specified, the command displays detailed information about the specified volume. If this parameter is specified by itself, the command displays information about all volumes matching the specified name.

[-aggregate <aggregate name>] - Aggregate Name

If this parameter is specified, the command displays information only about the volume or volumes that are located on the specified storage aggregate. This field is displayed as "-" for FlexGroup volumes.

[-aggr-list <aggregate name>, ...] - List of Aggregates for FlexGroup Constituents

If this parameter is specified, the command displays information only about the FlexGroup volume or FlexGroup volumes that are located on the specified list of storage aggregates. This parameter applies to FlexGroup volumes only.

[-encryption-type {none|volume|aggregate|mixed}] - Encryption Type

If this parameter is specified, the command displays information about the type of encryption key used for encrypting the volume. The possible values are *none*, *volume*, and *aggregate*. The value *none* is used for non-encrypted volumes, the value *volume* is used for volumes encrypted with NVE (NetApp Volume Encryption) and *aggregate* is used for volumes encrypted with NAE (NetApp Aggregate Encryption).

[-nodes {<nodename>|local}] - List of Nodes Hosting the Volume

If this parameter is specified, the command displays information only about the FlexGroup volume or FlexGroup volumes that are located on the specified list of storage systems. This parameter applies to FlexGroup volumes only.

[-size {<integer>[KB|MB|GB|TB|PB]}] - Volume Size

If this parameter is specified, the command displays information only about the volume or volumes that have the specified size. Size is the maximum amount of space a volume can consume from its associated aggregate(s), including user data, metadata, Snapshot copies, and Snapshot reserve. Note that for volumes without a *-space-guarantee* of *volume*, the ability to fill the volume to this maximum size depends on the space available in the associated aggregate or aggregates.

[-name-ordinal <text>] - Name Ordinal (privilege: advanced)

If this parameter is specified, it denotes the ordinal assignment used in relation to this volume's name. Ordinals are used to disambiguate volumes that have the same base name on the same controller. A value of "0" indicates that the base volume name is unique on the controller. A value greater than zero indicates that the volume's base name is used by two or more volumes on the same controller, and that appending "(n)" to this volume's name uniquely identifies it on this controller.

[-dsid <integer>] - Volume Data Set ID

If this parameter is specified, the command displays information only about the volume or volumes that match the specified data set ID. This field is displayed as "-" for FlexGroup volumes.

[-msid <integer>] - Volume Master Data Set ID

If this parameter is specified, the command displays information only about the volume or volumes that match the specified master data set ID.

[-state {online|restricted|offline|force-online|force-offline|mixed}] - Volume State

If this parameter is specified, the command displays information only about the volume or volumes that have the specified state. The *mixed* state only applies to FlexGroup volumes. If the state of a FlexGroup volume is *mixed*, that indicates that not all of the constituents are in the same state. If this is the case use the "volume show -is-constituent true" command to find out which constituents are not in the proper state.

[-volume-style <flex>] - Volume Style

If this parameter is specified, the command displays information only about the volumes that have the specified style. Possible values are *flex* for FlexVol volumes.

[`-volume-style-extended` {`flexvol`|`flexgroup`|`flexgroup-constituent`}] - Extended Volume Style

If this parameter is specified, the command displays information only about the volumes that are configured with the specified extended style. Possible values are *flexvol* for FlexVol volumes, *flexgroup* for FlexGroup volumes and *flexgroup-constituent* for FlexGroup volume constituents.

[`-flexcache-endpoint-type` {`none`|`cache`|`origin`}] - FlexCache Endpoint Type

If this parameter is specified, the command displays information only about the volumes that are of the specified flexcache-endpoint-type. Possible values are *none* for volumes that are not part of a FlexCache relationship, *cache* for FlexCache volumes and *origin* for origin of FlexCache volumes.

[`-is-cluster-volume` {`true`|`false`}] - Is Cluster-Mode Volume

If this parameter is specified, the command displays information only about cluster volumes (true) or node root volumes and other node scoped volumes (false).

[`-is-constituent` {`true`|`false`}] - Is Constituent Volume

If this parameter is specified, the command displays information only about volumes that either are or are not constituents of a FlexGroup volume, depending on the value provided.

[`-constituent-count` <integer>] - Number of Constituent Volumes

The number of constituents in the FlexGroup volume. This parameter applies to FlexGroup volumes only.

[`-policy` <text>] - Export Policy

If this parameter is specified, the command displays information only about the volume or volumes that use the specified export policy.

[`-user` <user name>] - User ID

If this parameter is specified, the command displays information only about the volume or volumes whose root is owned by the specified user.

[`-group` <group name>] - Group ID

If this parameter is specified, the command displays information only about the volume or volumes whose root is owned by the specified group.

[`-security-style` <security style>] - Security Style

If this parameter is specified, the command displays information only about the volume or volumes that have the specified security style (*unix* for UNIX mode bits, *ntfs* for CIFS ACLs, *mixed* for both styles or *unified* for Unified UNIX, NFS and CIFS permissions).

[`-unix-permissions` <unix perm>] - UNIX Permissions

If this parameter is specified, the command displays information only about the volume or volumes whose default UNIX permissions match the specified permissions. Specify UNIX permissions either as a four-digit octal value (for example, 0700) or in the style of the UNIX ls command (for example, -rwxr-x---). For information on UNIX permissions, see the UNIX or Linux documentation.

[`-junction-path` <junction path>] - Junction Path

If this parameter is specified, the command displays information only about the volume or volumes that have the specified junction path.

[-junction-path-source {RW_volume|LS_mirror}] - Junction Path Source

If this parameter is specified, the command displays information only about the volume or volumes that have the specified junction path source.

[-junction-active {true|false}] - Junction Active

If this parameter is specified, the command displays information only about the volume or volumes whose junction paths have the specified status.

[-junction-parent <volume name>] - Junction Parent Volume

If this parameter is specified, the command displays information only about the volume or volumes that have the specified parent volume.

[-vsroot {true|false}] - Vserver Root Volume (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that match the specified setting; that is, whether they are the root volumes for their Vservers.

[-comment <text>] - Comment

If this parameter is specified, the command displays information only about the volume or volumes that match the specified comment text.

[-available {<integer>[KB|MB|GB|TB|PB] }] - Available Size

If this parameter is specified, the command displays information only about the volume or volumes that have the specified available size. Available is the amount of free space currently available to be used by this volume. For a volume with a `-space-guarantee` of type *volume*, available is always `-total` minus `-used`. For volumes that do not have a `-space-guarantee` of type *volume*, available could be reduced if the volume's associated aggregate or aggregates are space constrained.

[-filesystem-size {<integer>[KB|MB|GB|TB|PB] }] - Filesystem Size

If this parameter is specified, the command displays information only about the volume or volumes that have the specified filesystem size. Filesystem size is the same as the volume's `-size` unless the volume is or was a physical replica destination. In this case, the file system size corresponds to the `-size` of the source volume, until `-filesys-size-fixed` is set to *false*.

[-total {<integer>[KB|MB|GB|TB|PB] }] - Total User-Visible Size

If this parameter is specified, the command displays information only about the volume or volumes that have the specified total size. Total is the total space available for user data and file system metadata. It does not include the Snapshot reserve.

[-used {<integer>[KB|MB|GB|TB|PB] }] - Used Size

If this parameter is specified, the command displays information only about the volume or volumes that have the specified used size. Used is the amount of space occupied by user data and file system metadata. *used* differs from *physical-used* space by the sum of the space that is reserved for future writes and the space that is saved by aggregate storage efficiency savings being removed from *used*. It includes Snapshot spill (the amount of space by which Snapshot copies exceed Snapshot reserve). It does not include the Snapshot reserve.

[-percent-used <percent>] - Used Percentage

If this parameter is specified, the command displays information only about the volume or volumes that have the specified percentage of used space. This row is based on a value of used space that includes the space used by Snapshot copies or the Snapshot reserve (whichever is greater) in relation to the current volume size.

[-space-nearly-full-threshold-percent <percent>] - Volume Nearly Full Threshold Percent

If this parameter is specified, the command displays information only about the volume or volumes that have the specified nearly full threshold percent.

[-space-full-threshold-percent <percent>] - Volume Full Threshold Percent

If this parameter is specified, the command displays information only about the volume or volumes that have the specified full threshold percent.

[-max-autosize {<integer>[KB|MB|GB|TB|PB]}] - Maximum Autosize

If this parameter is specified, the command displays information only about the volume or volumes that have the specified maximum automatic size.

[-min-autosize {<integer>[KB|MB|GB|TB|PB]}] - Minimum Autosize

If this parameter is specified, the command displays information only about the volume or volumes that have the specified minimum automatic size. This field is displayed as "-" for FlexGroup volumes.

[-autosize-grow-threshold-percent <percent>] - Autosize Grow Threshold Percentage

If this parameter is specified, the command displays information only about the volume or volumes that have the specified automatic grow used space threshold percentage. This field is displayed as "-" for FlexGroup volumes.

[-autosize-shrink-threshold-percent <percent>] - Autosize Shrink Threshold Percentage

If this parameter is specified, the command displays information only about the volume or volumes that have the specified automatic shrink used space threshold percentage. This field is displayed as "-" for FlexGroup volumes.

[-autosize-mode {off|grow|grow_shrink}] - Autosize Mode

If this parameter is specified, the command displays information only about the volume or volumes that have the specified automatic sizing mode setting. This field is displayed as "-" for FlexGroup volumes.

[-files <integer>] - Total User Visible Files

If this parameter is specified, the command displays information only about the volume or volumes that have the specified number of files.

[-files-used <integer>] - User Visible Files Used

If this parameter is specified, the command displays information only about the volume or volumes that have the specified number of files used. This includes all user-visible file system entities such as directories, regular files, streams, ACLS, etc.

[-inodefile-public-capacity <integer>] - User Visible Inode File Capacity (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that have the specified public inode file capacity. This includes all user-invisible or metadata files such as private metabytes and zombies.

[-files-maximum-possible <integer>] - Maximum Possible Total User Visible Files (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that have the specified number of maximum files possible. This includes all user-visible file system entities such as directories, regular files, streams, ACLS, etc.

[-files-set-maximum {true|false}] - Set Total Files (for user-visible data) to the Highest Value that the Volume can Hold (privilege: advanced)

If this parameter is specified, the command displays information about whether the volume's total number of files is set to the highest value that the volume can hold. This includes all user-visible file system entities such as directories, regular files, streams, ACLS, etc.

[-files-private-used <integer>] - Private System Metadata Files Used (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that have the specified number of private system metadata files used. This includes all user-invisible or metadata files such as private metafiles and zombies.

[-inodefile-private-capacity <integer>] - Private System Metadata Inode File Capacity (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that have the specified private inode file capacity. This includes all user-visible file system entities such as directories, regular files, streams, ACLS, etc.

[-maxdir-size {<integer>[KB|MB|GB|TB|PB]}] - Maximum Directory Size (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that have the specified maximum directory size.

[-space-guarantee-enabled {true|false}] - Space Guarantee in Effect

If this parameter is specified, the command displays information only about the volume or volumes that have the specified space-guarantee setting. If the value of `-space-guarantee` is *none*, the value of `-space-guarantee-enabled` is always *true*. In other words, because there is no guarantee, the guarantee is always in effect. If the value of `-space-guarantee` is *volume*, the value of `-space-guarantee-enabled` can be *true* or *false*, depending on whether the guaranteed amount of space was available when the volume was mounted.

[-is-space-slo-enabled {true|false}] - Space SLO in Effect

If this parameter is specified, the command displays information only about the volume or volumes that have their `space-slo` setting in effect or not, depending on the value specified for this parameter. If the value of `space-slo` is *none*, the space SLO is always considered to be in effect. If the value of `space-slo` is *semi-thick* or *thick*, the space SLO may be in effect depending on whether the required amount of space was available when the volume was mounted.

[-space-slo {none|thick|semi-thick}] - Space SLO

If this parameter is specified, the command displays information only about the volume or volumes that have the specified `space-slo` setting. The space SLO setting is the Service Level Objective for space management for the volume.

[-s, -space-guarantee {none|volume}] - Space Guarantee Style

If this parameter is specified, the command displays information only about the volume or volumes that have the specified space guarantee style. If the value of `-space-guarantee` is *none*, the value of `-space-guarantee-enabled` is always *true*. In other words, because there is no guarantee, the guarantee is always in effect. If the value of `-space-guarantee` is *volume*, the value of `-space-guarantee-enabled` can be *true* or *false*, depending on whether the guaranteed amount of space was available when the volume was mounted.

[-fractional-reserve <percent>] - Fractional Reserve

If this parameter is specified, the command displays information only about the volume or volumes that

have the specified `fractional-reserve` setting.

[`-type {RW|DP}`] - Volume Type

If this parameter is specified, the command displays information only about the volume or volumes of the specified volume type (RW for read-write or DP for data-protection).

[`-min-readahead {true|false}`] - Minimum Read Ahead (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that have the specified `minimum-readahead` setting.

[`-atime-update {true|false}`] - Access Time Update Enabled (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that have the specified `access-time update` setting.

[`-snapdir-access {true|false}`] - Snapshot Directory Access Enabled

If this parameter is specified, the command displays information only about the volume or volumes that have the specified `Snapshot-copy access` setting.

[`-percent-snapshot-space <percent>`] - Space Reserved for Snapshot Copies

If this parameter is specified, the command displays information only about the volume or volumes that have the specified percentage of space reserved for Snapshot copies.

[`-snapshot-space-used <percent_no_limit>`] - Snapshot Reserve Used

If this parameter is specified, the command displays information only about the volume or volumes that have the specified used percentage of the reserve for Snapshot copies.

[`-snapshot-policy <snapshot policy>`] - Snapshot Policy

If this parameter is specified, the command displays information only about the volume or volumes that use the specified Snapshot policy.

[`-create-time <Date>`] - Creation Time

If this parameter is specified, the command displays information only about the volume or volumes that have the specified creation time.

[`-language <Language code>`] - Language

If this parameter is specified, the command displays information only about the volume or volumes that store data in the specified language. To determine the available languages, enter `volume show -language` ?`` at the clustershell command prompt.

[`-clone-volume {true|false}`] - Clone Volume

If this parameter is specified, the command displays information only about volumes that are clones (true) or not clones (false).

[`-node {<nodename>|local}`] - Node name

If this parameter is specified, the command displays information only the volume or volumes that are located on the specified storage system. This field is displayed as "-" for FlexGroup volumes.

[`-clone-parent-vserver <vserver name>`] - Clone Parent Vserver Name

If this parameter is specified, the command displays information only about the volumes with a matching FlexClone parent Vserver name.

[-clone-parent-name <volume name>] - FlexClone Parent Volume

If this parameter is specified, the command displays information only about the volumes with a matching FlexClone parent volume name.

[-uuid <UUID>] - UUID of the Volume (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that match the specified UUID.

[-nvfail {on|off}] - NVFAIL Option

If this parameter is specified, the command displays information only about volumes for which failover is enabled (on) or disabled (off).

[-in-nvfailed-state {true|false}] - Volume's NVFAIL State

If this parameter is specified, the command displays information only about volumes which are in the failed over state (true) or not (false). This field is only available when the volume is online.

[-dr-force-nvfail {on|off}] - Force NVFAIL on MetroCluster Switchover

If this parameter is specified, the command displays information only about volumes for which dr-force-nvfail is enabled (on) or disabled (off).

[-filesystem-size-fixed {true|false}] - Is File System Size Fixed

If this parameter is specified, the command displays information only about the volume or volumes that have the specified `filesystem-size-fixed` setting.

[-extent-enabled {off|on|space-optimized}] - (DEPRECATED)-Extent Option



This parameter has been deprecated and might be removed in a future release of ONTAP.

If this parameter is specified, the command displays information only about volumes that have extents enabled (on), not enabled (off) or space optimized (space-optimized).

[-overwrite-reserve {<integer>[KB|MB|GB|TB|PB]}] - Reserved Space for Overwrites

If this parameter is specified, the command displays information only about the volume or volumes that have the specified `overwrite-reserve` setting.

[-space-mgmt-try-first {volume_grow|snap_delete}] - Primary Space Management Strategy

If this parameter is specified, the command displays information only about the volume or volumes that have the specified `space-mgmt-try-first` setting. Possible values are `volume_grow` and `snap_delete`. This field is displayed as "-" for FlexGroup volumes.

[-read-realloc {off|on|space-optimized}] - Read Reallocation Option

If this parameter is specified, the command displays information only about volumes that have read reallocation enabled (on), not enabled (off) or space optimized (space-optimized).

[-sched-snap-name {create-time|ordinal}] - Naming Scheme for Automatic Snapshot Copies

If this parameter is specified, the command displays information only about the volume or volumes that have the specified automatic Snapshot-copy naming convention.

[-is-inconsistent {true|false}] - Inconsistency in the File System

If this parameter is specified, the command displays information only about volumes that are inconsistent (true) or consistent (false) in the file system.

[`-is-quieted-on-disk {true|false}`] - Is Volume Quiesced (On-Disk)

If this parameter is specified, the command displays information only about volumes that are quiesced (true) or not quiesced (false) on disk.

[`-is-quieted-in-memory {true|false}`] - Is Volume Quiesced (In-Memory)

If this parameter is specified, the command displays information only about volumes that are quiesced (true) or not quiesced (false) in memory.

[`-transition-state <state>`] - Transition Operation State (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that match the specified transition state.

[`-transition-behavior {data-move|data-protection|none}`] - Transition Behavior (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that match the specified transition behavior. Possible values are:

- data-move: Volumes that are being moved from a system operating in 7-Mode.
- data-protection: Volumes that are being replicated from a system operating in 7-Mode for disaster recovery.
- none: Volumes that are not part of transition.

[`-is-copied-for-transition {true|false}`] - Copied for Transition (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that match the specified value based on whether the volume is copied for transition or not.

[`-is-transitioned {true|false}`] - Transitioned (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that match the specified value based on whether the volume is transitioned or not.

[`-is-sis-volume {true|false}`] - Volume Contains Shared or Compressed Data

If this parameter is specified, the command displays information only about those volumes that match the specified storage efficiency setting.

[`-sis-space-saved {<integer>[KB|MB|GB|TB|PB]}`] - Space Saved by Storage Efficiency

If this parameter is specified, the command displays information only about those volumes that have the specified amount of space saved by the storage efficiency technology.

[`-sis-space-saved-percent <percent>`] - Percentage Saved by Storage Efficiency

If this parameter is specified, the command displays information only about those volumes that have the specified percentage of space saved by the storage efficiency technology.

[`-dedupe-space-saved {<integer>[KB|MB|GB|TB|PB]}`] - Space Saved by Deduplication Along With VBN ZERO Savings

If this parameter is specified, the command displays information only about those volumes that have the specified amount of space saved due to deduplication.

[`-dedupe-space-saved-percent <percent>`] - Percentage Saved by Deduplication

If this parameter is specified, the command displays information only about those volumes that have the specified percentage of space saved due to deduplication.

[`-dedupe-space-shared` {<integer>[KB|MB|GB|TB|PB]}] - Unique Data Which Got Shared by Deduplication

If this parameter is specified, the command displays information only about those volumes that have the specified amount of shared space due to deduplication.

[`-compression-space-saved` {<integer>[KB|MB|GB|TB|PB]}] - Space Saved by Compression

If this parameter is specified, the command displays information only about those volumes that have the specified amount of space saved due to compression.

[`-compression-space-saved-percent` <percent>] - Percentage Space Saved by Compression

If this parameter is specified, the command displays information only about those volumes that have the specified percentage of space saved due to compression.

[`-size-used-by-snapshots` {<integer>[KB|MB|GB|TB|PB]}] - Volume Size Used by Snapshot Copies

If this parameter is specified, the command displays information about those volumes that have the specified volume size used by Snapshot copies.

[`-block-type` {64-bit|extent|32-bit}] - Block Type

If this parameter is specified, the command displays information about only the volumes with the specified indirect block format. Possible values are *32-bit* to display 32-bit volumes and *64-bit* to display 64-bit volumes.

[`-is-moving` {true|false}] - Is Volume Moving

If this parameter is specified, the command displays information only about volumes that are moving (true) or not moving (false).

[`-hybrid-cache-eligibility` {read|read-write|none}] - Flash Pool Caching Eligibility

If this parameter is specified, the command displays information only about the volume or volumes with the specified Flash Pool caching attributes. Possible caching attributes are:

- 'read' ... Indicates that the volume cannot participate in write caching.
- 'read-write' ... Indicates that the volume can participate in read and write caching.

[`-hybrid-cache-write-caching-ineligibility-reason` <text>] - Flash Pool Write Caching Ineligibility Reason

If this parameter is specified, the command displays information only about the volume or volumes which are ineligible to participate in write caching due to the specified reason.

[`-constituent-role` <Constituent Roles>] - Constituent Volume Role

If this parameter is specified, the command displays information only about the constituent volume or volumes that are of the specified role. This parameter applies to FlexGroup volumes only.

[`-is-cft-precommit` {true|false}] - Is in the precommit phase of Copy-Free Transition (privilege: advanced)

If this parameter is specified with the true value, it displays information only about the volumes that are in the precommit phase of a Copy-Free Transition workflow.

[`-qos-policy-group` <text>] - QoS Policy Group Name

If this parameter is specified, the command displays information only about volumes that match the specified QoS policy group.

[-qos-adaptive-policy-group <text>] - QoS Adaptive Policy Group Name

If this parameter is specified, the command displays information only about volumes that match the specified QoS adaptive policy group.

[-caching-policy <text>] - Caching Policy Name

If this parameter is specified, the command displays the volumes that match the specified caching policy.

A caching policy defines how the system caches a volume's data in a Flash Pool aggregate. Both metadata and user data are eligible for caching. The available caching policies are:

- none - Does not cache any user data or metadata blocks.
- auto - Read caches all metadata and randomly read user data blocks, and write caches all randomly overwritten user data blocks.
- meta - Read caches only metadata blocks.
- random_read - Read caches all metadata and randomly read user data blocks.
- random_read_write - Read caches all metadata, randomly read and randomly written user data blocks.
- all_read - Read caches all metadata, randomly read and sequentially read user data blocks.
- all_read_random_write - Read caches all metadata, randomly read, sequentially read, and randomly written user data.
- all - Read caches all data blocks read and written. It does not do any write caching.
- noread-random_write - Write caches all randomly overwritten user data blocks. It does not do any read caching.
- meta-random_write - Read caches all metadata and write caches randomly overwritten user data blocks.
- random_read_write-random_write - Read caches all metadata, randomly read and randomly written user data blocks. It also write caches randomly overwritten user data blocks.
- all_read-random_write - Read caches all metadata, randomly read and sequentially read user data blocks. It also write caches randomly overwritten user data blocks.
- all_read_random_write-random_write - Read caches all metadata, randomly read, sequentially read, and randomly written user data. It also write caches randomly overwritten user data blocks.
- all-random_write - Read caches all data blocks read and written. It also write caches randomly overwritten user data blocks.

Note that in a caching-policy name, a hyphen (-) separates read and write policies. Default caching-policy is auto.

[-cache-retention-priority {normal|low|high}] - Cache Retention Priority (privilege: advanced)

If this parameter is specified, the command displays the volumes that match the specified cache retention priority policy.

A cache retention priority defines how long the blocks of a volume will be cached in flash pool once they become cold. The available cache retention priority are:

- low - Cache the cold blocks for the lowest time.
- normal - Cache the cold blocks for the default time.
- high - Cache the cold blocks for the highest time.

[-is-volume-in-cutover {true|false}] - Is Volume Move in Cutover Phase

If this parameter is specified, the command displays information only about volumes that are in the cutover phase (true) or not in the cutover phase (false) of a volume move. This field is displayed as "-" for FlexGroup volumes.

[-snapshot-count <integer>] - Number of Snapshot Copies in the Volume

If this parameter is specified, the command displays information only about the volumes that have the specified number of Snapshot copies.

[-vbn-bad-present {true|false}] - VBN_BAD may be present in the active filesystem

If this parameter is specified, the command displays information only about volumes that may have VBN_BAD present in its active filesystem (true) or do not have VBN_BAD present in its active filesystem (false).

[-is-autobalance-eligible {true|false}] - Is Eligible for Auto Balance Aggregate (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that are eligible for consideration by the Auto Balance Aggregate feature.

[-is-vol-on-hybrid-aggr {true|false}] - Is Volume on a hybrid aggregate

If this parameter is specified, the command displays information only about volumes associated with a Flash Pool aggregate (true) or not (false). This field is displayed as "-" for FlexGroup volumes.

[-physical-used {<integer>[KB|MB|GB|TB|PB]}] - Total Physical Used Size

If this parameter is specified, the command displays information only about the volume or volumes that have the specified physical used size. Physical-used represents effective total footprint which is total footprint after deducting auto adaptive compression and compaction savings. Total-footprint includes aggregate metadata used by volume. *physical-used* differs from *total-used* space by the sum of the space that is reserved for future writes and the space that is saved by aggregate storage efficiency savings being added to *physical-used*. The value includes blocks in use by Snapshot copies.

[-physical-used-percent <percent_no_limit>] - Physical Used Percentage

If this parameter is specified, the command displays information only about the volume or volumes that have the specified physical used percent based on volume size including the space reserved for Snapshot copies.

[-flexgroup-msid <integer>] - FlexGroup Master Data Set ID (privilege: advanced)

If this parameter is specified, the command displays information only about the FlexGroup volume or FlexGroup volume constituents that have the specified FlexGroup volume master data-set ID. This parameter applies to FlexGroup volumes and FlexGroup volume constituents only.

[-flexgroup-index <integer>] - FlexGroup Index (privilege: advanced)

If this parameter is specified, the command displays information only about the FlexGroup volume constituents that have the specified FlexGroup volume index. This parameter applies to FlexGroup volume constituents only.

[-flexgroup-uuid <UUID>] - UUID of the FlexGroup (privilege: advanced)

If this parameter is specified, the command displays information only about the FlexGroup volume or FlexGroup volume constituents that have the specified FlexGroup volume UUID. This parameter applies to FlexGroup volumes and FlexGroup volume constituents only.

[-flexgroup-name <volume name>] - FlexGroup Name

If this parameter is specified, the command displays information only about the volumes that have the specified FlexGroup volume name. This parameter only applies to FlexGroup volumes and FlexGroup volume constituents.

[-max-constituent-size {<integer>[KB|MB|GB|TB|PB]}] - Maximum size of a FlexGroup Constituent (privilege: advanced)

If this parameter is specified, the command displays information only about the FlexGroup volume or FlexGroup volumes that have the specified maximum constituent size. This parameter applies to FlexGroup volumes only.

[-inofile-version <integer>] - Inofile Version (privilege: advanced)

If this parameter is specified, the command displays information only about the volumes whose inode files are at the specified version.

[-is-flexgroup {true|false}] - Is Volume a FlexGroup

If this parameter is specified, the command displays information only about the volume or volumes that are either FlexGroup volumes or not, depending on the value provided.

[-is-qtree-caching-enabled {true|false}] - Is Qtree Caching Support Enabled (privilege: advanced)

If this parameter is specified, the command displays information only about the FlexGroup volumes and origins of FlexCache volumes with Qtree caching enabled or disabled, depending on the value provided.

[-is-move-destination-in-cutover {true|false}] - Is the Volume a Target of a Move Which Is Currently in Cutover (privilege: advanced)

If this parameter is specified, the command displays whether or not the volume is a move destination that is currently in cutover.

[-snaplock-type {non-snaplock|compliance|enterprise}] - SnapLock Type

If this parameter is specified, the command displays information only about volumes that match the specified snaplock-type.

[-vserver-dr-protection {protected|unprotected}] - Vserver DR Protection

If this parameter is specified, the command displays information only about the volumes having the specified Vserver SnapMirror protection.

[-has-optimized-sparse-directories {true|false}] - Volume Has Sparse Directories in an Optimized Format (privilege: advanced)

If this parameter is specified, the command displays information only about volumes that have sparse directories in an optimized format.

[-encrypt {true|false}] - Enable or Disable Encryption

If this parameter is specified, the command displays information only about the volumes that are encrypted (true) or not encrypted (false).

[-is-encrypted {true|false}] - Is Volume Encrypted

If this parameter is specified, the command displays information only about the volumes that are encrypted (true) or unencrypted (false).

[-encryption-state {none|full|partial|converting_to_encrypted|converting_to_plaintext|rekeying}] - Encryption State

If this parameter is specified, the command displays information only about the volumes that have the specified encryption state. The possible values are *none*, *full*, *partial*, *converting_to_encrypted*, *converting_to_plaintext*, and *rekeying*. The value *partial* is used for FlexGroup volumes, which indicates that some constituents are encrypted and some are not.

[-key-id <text>] - Encryption Key ID

If this parameter is specified, the command displays information only about the volume whose encryption key-id matches the specified key-id.

[-key-creation-time <MM/DD/YYYY HH:MM:SS>] - Encryption Key Creation Time

If this parameter is specified, the command displays information only about volumes having an encryption key with a creation time that matches, is greater than or less than the provided key-creation-time value. For example, the `volume show -key-creation-time >1/1/2022 13:12:11` command will show all volumes that have an encryption key created later than 1/1/2022 13:12:11

[-application <text>] - Application

Selects the volumes that are part of an application that matches the parameter value.

[-is-protocol-access-fenced {true|false}] - Is Fenced for Protocol Access

If this parameter is specified, the command displays information only about the volumes that are fenced for protocol access. Only FlexGroup volume constituents and volumes in SnapMirror Synchronous relationships can be fenced for protocol access.

[-protocol-access-fenced-by {none|coordinated_snaprestore|coordinated_redirection|snapmirror_synchronous|vserver_migrate}] - Protocol Access Fence Owner

This field indicates the owner of the protocol access fence when the volume's protocol access is fenced. Only FlexGroup volume constituents and volumes in SnapMirror Synchronous relationships can be fenced for protocol access.

[-single-instance-data-logging {off|on}] - Is SIDL enabled

If this parameter is specified, the command displays whether Single Instance Data Logging feature is enabled on the specified volume.

[-over-provisioned {<integer>[KB|MB|GB|TB|PB]}] - Over Provisioned Size

If this parameter is specified, the command displays information only about the volume or volumes that have the specified size unavailable in the aggregate. This applies only for a "none" guaranteed volume when its unused size exceeds available space in the aggregate. This value is always zero for "volume" guaranteed volumes. For FlexGroup volumes, the value is not the direct sum of the constituents' overprovisioned size. Instead the value represents the total overprovisioned space across all aggregates for the FlexGroup volume. For each aggregate, the value is computed as the sum of the unused size of the thin provisioned constituents of the FlexGroup volumes that exceed available space in the aggregate.

[-snapshot-reserve-available {<integer>[KB|MB|GB|TB|PB]}] - Available Snapshot Reserve Size

If this parameter is specified, the command displays information only about the volume or volumes that have the specified size available for Snapshot copies within the Snapshot reserve. This value is zero if Snapshot spill is present. For 'none' guaranteed volumes, this may get reduced due to less available space

in the aggregate.

[-logical-used {<integer>[KB|MB|GB|TB|PB]}] - Logical Used Size

If this parameter is specified, the command 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.

[-logical-used-percent <percent_no_limit>] - Logical Used Percentage

If this parameter is specified, the command displays information only about the volume or volumes that have the specified logical used percentage.

[-logical-available {<integer>[KB|MB|GB|TB|PB]}] - Logical Available Size

If this parameter is specified, the command displays information only about the volume or volumes that have the specified logical available size. This value is the amount of free space currently available considering space saved by the storage efficiency features as being used. This does not include Snapshot reserve.

[-logical-used-by-afs {<integer>[KB|MB|GB|TB|PB]}] - Logical Size Used by Active Filesystem

If this parameter is specified, the command displays information only about the volume or volumes that have the specified logical size used by the active file system. This value differs from *logical-used* by the amount of Snapshot spill that exceeds Snapshot reserve. This parameter is not supported on FlexGroup volumes.

[-logical-used-by-snapshots {<integer>[KB|MB|GB|TB|PB]}] - Logical Size Used by All Snapshots

If this parameter is specified, the command displays information only about the volume or volumes that have the specified logical size used across all Snapshot copies. This value differs from *size-used-by-snapshots* by the space saved by the storage efficiency features across the Snapshot copies. This parameter is not supported on FlexGroup volumes.

[-is-space-reporting-logical {true|false}] - Logical Space Reporting

If this parameter is specified, the command displays information only about the volumes that have logical space reporting enabled or disabled as specified. When space is reported logically, ONTAP reports the volume space such that all the physical space saved by the storage efficiency features are also as reported as used.

[-is-space-enforcement-logical {true|false}] - Logical Space Enforcement

If this parameter is specified, the command displays information only about the volumes that have logical space enforcement enabled or disabled as specified. When space is enforced logically, ONTAP enforces volume settings such that all the physical space saved by the storage efficiency features will be calculated as used.

[-tiering-policy <Tiering Policy>] - Volume Tiering Policy

If this parameter is specified, the command displays information only about the volumes whose tiering policy matches the specified value. Tiering policies determine whether the user data blocks of a volume in a FabricPool will be tiered to the cloud tier when they become cold. FabricPool combines Flash (performance tier) with an object store (cloud tier) into a single aggregate. The temperature of a volume block increases if it is accessed frequently and decreases when it is not.

The available tiering policies are:

- snapshot-only - Only the volume Snapshot copies not associated with the active file system are tiered to the cloud tier.
- auto - Both Snapshot copy data and active file system user data are tiered to the cloud tier.
- none - No volume blocks are tiered to the cloud tier.
- all - Both Snapshot copy data and active file system user data are tiered to the cloud tier as soon as possible without waiting for a cooling period. On DP volumes all transferred user data blocks start in the cloud tier.

[`-cloud-retrieval-policy` {`default`|`on-read`|`never`|`promote`}] - Volume Cloud Retrieval Policy (privilege: advanced)

This optional parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier.

The available cloud retrieval policies are:

- default - This policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot-only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved.
- on-read - This policy retrieves tiered data for all client driven data reads.
- never - This policy never retrieves tiered data.
- promote - This policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot-only'. If the tiering policy is 'snapshot-only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud.

[`-tiering-minimum-cooling-days` <integer>] - Volume Tiering Minimum Cooling Days (privilege: advanced)

This parameter displays the minimum number of days that user data blocks of the volume must be cooled before they can be considered cold and tiered out to the cloud tier. For volumes hosted on FabricPools, this parameter is used for tiering purposes and does not affect the reporting of inactive data. For volumes hosted on non-FabricPools, this parameter affects the cooling window used for reporting inactive data. The default minimum cooling period for the *snapshot-only* policy is 2 days and for the *auto* policy is 31 days.

[`-performance-tier-inactive-user-data` {<integer>[`KB`|`MB`|`GB`|`TB`|`PB`]}] - Performance Tier Inactive User Data

If this parameter is specified, the command displays the amount of inactive user data stored in the performance tier that could be tiered out to a cloud tier if the volume is in a FabricPool and for which the *auto* tiering policy has been specified. For more information see the `tiering-policy` parameter.

[`-performance-tier-inactive-user-data-percent` <percent>] - Performance Tier Inactive User Data Percent

If this parameter is specified, the command displays the percentage of inactive user data in the performance tier.

[`-tiering-object-tags` <text>,...] - Tags to be Associated with Objects Stored on a FabricPool

This optional parameter specifies the tiering object tags to be associated with objects stored on a FabricPool.

[`-needs-object-retagging {true|false}`] - Does the Object Tagging Scanner Need to Run on This Volume

This optional parameter specifies if an object tagging scanner needs to run again for the volume.

[`-is-analytics-supported {true|false}`] - Is File System Analytics Supported

If this parameter is specified, the command displays information only about volumes that are supported (true) or not supported (false). This field indicates whether or not file system analytics is supported on the volume. If file system analytics is not supported, the reason will be specified in the *analytics_unsupported_reason* field.

[`-analytics-unsupported-reason <text>`] - Reason File System Analytics is not Supported

If this parameter is specified, the command displays information only about volumes that are not supported with the specified reason. If file system analytics is not supported on the volume, this field provides the reason why.

[`-analytics-state <Analytics State>`] - File System Analytics State

If this parameter is specified, the command displays information only about volumes with file system analytics *on*, *off* or *initializing*. If this value is *on*, ONTAP collects extra file system analytics information for all directories on the volume. There will be a slight impact to I/O performance to collect this information. If this value is *off*, file system analytics information is not collected and not available to be viewed. If this value is *initializing*, that means file system analytics was recently turned on, and the initialization scan to gather information for all existing files and directories is currently running.

[`-analytics-scan-progress <percent>`] - File System Analytics Scan Progress

If this parameter is specified, the command displays information only about the volume or volumes that have the specified file system analytics progress percentage. This value will only be set if a file system analytics initialization scan is in progress.

[`-analytics-files-scanned <integer>`] - File System Analytics Files Scanned Progress

If this parameter is specified, the command displays information only about the volume or volumes that have the number of files scanned in the file system. This value will only be set if a file system analytics initialization scan is in progress.

[`-analytics-total-files <integer>`] - File System Analytics Total Files

If this parameter is specified, the command displays information only about the volume or volumes that have the total files in the file system analytics progress percentage. This value will only be set if a file system analytics initialization scan is in progress.

[`-activity-tracking-state <Activity Tracking State>`] - Activity Tracking State

If this parameter is specified, the command displays information only about volumes with volume activity tracking *on* or *off*. If this value is *on*, ONTAP tracks volume activity in real time to provide a more detailed view. There will be a slight impact to I/O performance to collect this information. If this value is *off*, activity is not tracked and not available to be viewed.

[`-is-activity-tracking-supported {true|false}`] - Is Activity Tracking Supported

If this parameter is specified, the command displays information only about volumes that are supported (true) or not supported (false). This field indicates whether or not volume activity tracking is supported on the volume. If volume activity tracking is not supported, the reason will be specified in the *activity_tracking_unsupported_reason* field.

[-activity-tracking-unsupported-reason <text>] - Reason Activity Tracking Is Not Supported

If this parameter is specified, the command displays information only about volumes that are not supported with the specified reason. If volume activity tracking is not supported on the volume, this field provides the reason why.

[-is-smbc-master {true|false}] - Is SnapMirror Active Sync Master

If this parameter is specified, it displays if the volume is acting as a SnapMirror active sync master.

[-is-smbc-failover-capable {true|false}] - Is SnapMirror Active Sync Failover Capable

This parameter specifies if volume is capable of SnapMirror active sync failover.

[-smbc-consensus {Awaiting-consensus|Consensus|No-consensus}] - SnapMirror Active Sync Consensus

This parameter specifies the SnapMirror active sync consensus value.

Consensus indicated if cluster has to serve IO. Possible consensus values:

- Consensus Awaiting: Allow IO
- Consensus: Allow IO
- No consensus: Do not allow IO

[-anti-ransomware-state {disabled|enabled|dry-run|paused|dry-run-paused|enable-paused|disable-in-progress}] - Anti-ransomware State

If this parameter is specified, the command displays information only about the volumes that have the specified Anti-ransomware-state. The possible values are *disabled*, *enabled*, *dry-run*, *dry-run-paused*, *enable-paused* or *disable-in-progress*. The value *disabled* is used for volumes which do not have Anti-ransomware feature enabled. The value is *enabled* when Anti-ransomware feature is enabled on the volume. If the value is *dry-run*, the volume is in Anti-ransomware evaluation mode. The value is *dry-run-paused* when the Anti-ransomware feature is paused from evaluation mode on the volume. The value is *enable-paused* when the Anti-ransomware feature is paused on the volume. The value is *disable-in-progress* when disable work is in progress on the volume.

[-granular-data {disabled|basic}] - Granular data

If this parameter is specified, the command displays information only about volumes that match the specified value for whether data storage on the volume is granular or not.

[-atime-update-period <integer>] - Access Time Update Period (Seconds) (privilege: advanced)

"If this parameter is specified, the command displays information only about volumes that match the specified value for the atime update period."

[-snapshot-locking-enabled {true|false}] - Enable Snapshot Copy Locking

"If this parameter is specified, the command displays information only about volumes that match the specified `-snapshot-locking-enabled` value. A volume with locked Snapshot copies cannot be deleted until the `-expiry-time` has passed."

[-expiry-time <text>] - Expiry Time

If this parameter is specified, the command displays all the volumes that match the specified `-expiry-time` value.

[`-compliance-clock-time <text>`] - ComplianceClock Time

If this parameter is specified, the command displays all the volumes that match the specified `-compliance-clock-time` value.

[`-is-large-size-enabled {true|false}`] - Are Large Size Volumes and Files Enabled

If this parameter is specified, the command displays all the volumes that match the specified `-is-large-size-enabled` value.

[`-is-preserve-unlink-enabled {true|false}`] - Is Preserve Unlink Enabled (privilege: advanced)

If this parameter is specified, the command displays information only about volumes that match the specified `-is-preserve-unlink-enabled` value.

[`-is-cloud-write-enabled {true|false}`] - Is Cloud Write Enabled (privilege: advanced)

If this parameter is specified, the command displays all the volumes that match the specified `-is-cloud-write-enabled` value.

[`-aggressive-readahead-mode {none|file_prefetch}`] - Aggressive readahead mode (privilege: advanced)

If this parameter is specified, the command displays all the volumes that match the specified `-aggressive-readahead-mode` value.

[`-in-consistency-group {true|false}`] - If this Volume is part of a Consistency Group

This parameter specifies if this volume is associated with a consistency group.

[`-total-metadata {<integer>[KB|MB|GB|TB|PB]}`] - Total Metadata Size

If this parameter is specified, the command displays information only about the volume or volumes that have the specified amount of space used by metadata in the volume. This field will not include the metadata of the volume which is residing inside the aggregate.

[`-total-metadata-footprint {<integer>[KB|MB|GB|TB|PB]}`] - Total Metadata Footprint Used in Aggregate

If this parameter is specified, the command displays information only about volume or volumes that have the specified amount of space used by volume metadata footprint in the aggregate. The value of this field will be higher than the field 'total-metadata' from the [volume show-space](#) command output since this includes additional metadata of the volume in the aggregate.

Examples

The following example displays information about all volumes on the Vserver named vs1:

```
cluster1::> volume show -vserver vs1
Vserver   Volume      Aggregate   State    Type    Size   Available
Used%
-----
vs1       vol1        aggr1      online   RW      2GB    1.9GB
5%
vs1       vol1_dr     aggr0_dp   online   DP      200GB  160.0GB
```

```

20%
vs1      vol2      aggr0    online   RW       150GB    110.3GB
26%
vs1      vol2_dr   aggr0_dp online   DP       150GB    110.3GB
26%
vs1      vol3      aggr1    online   RW       150GB    120.0GB
20%
vs1      vol3_dr   aggr1_dp online   DP       150GB    120.0GB
20%
vs1      vol4      aggr1    online   RW       200GB    159.8GB
20%
vs1      vol4_dr   aggr1_dp online   DP       200GB    159.8GB
20%
vs1      vol5      aggr2    online   RW       200GB    102.3GB
48%
vs1      vol5_dr   aggr2_dp online   DP       200GB    102.3GB
48%
vs1      vol6      aggr2    online   RW       150GB    117.2GB
21%
vs1      vol6_dr   aggr2_dp online   DP       150GB    117.2GB
21%
vs1      vol7      aggr3    online   RW       150GB    118.5GB
20%
vs1      vol7_dr   aggr3_dp online   DP       150GB    118.5GB
20%
vs1      vol8      aggr3    online   RW       150GB    90.03GB
39%
vs1      vol8_dr   aggr3_dp online   DP       150GB    90.03GB
39%
vs1      vol9      aggr4    online   RW       150GB    43.67GB
70%
vs1      vol9_dr   aggr4_dp online   DP       150GB    43.67GB
70%
vs1      vol10     aggr4    online   RW       150GB    108.7GB
27%
vs1      vol10_dr  aggr4_dp online   DP       150GB    108.7GB
27%
vs1      vol11     aggr5    online   RW       250GB    45.65GB
81%
vs1      vol11_dr  aggr5_dp online   DP       250GB    45.65GB    81%
22 entries were displayed.

```

The following example displays detailed information about a volume named vol1 on an SVM named vs1:

```
cluster1::*> volume show -vserver vs1 -volume vol1
```

Vserver Name: vs1

Volume Name: voll
Aggregate Name: aggr1
Volume Size: 30MB
Volume Data Set ID: 1026
Volume Master Data Set ID: 2147484674
Volume State: online
Volume Type: RW
Volume Style: flex
Is Cluster Volume: true
Is Constituent Volume: false
Export Policy: default
User ID: root
Group ID: daemon
Security Style: mixed
Unix Permissions: ---rwx-----
Junction Path: -
Junction Path Source: -
Junction Active: -
Junction Parent Volume: -
Comment:
Available Size: 23.20MB
Filesystem Size: 30MB
Total User-Visible Size: 28.50MB
Used Size: 5.30MB
Used Percentage: 22%
Volume Nearly Full Threshold Percent: 95%
Volume Full Threshold Percent: 98%
Maximum Autosize (for Flexvol volumes only): 8.40GB
Minimum Autosize: 30MB
Autosize Grow Threshold Percentage: 85%
Autosize Shrink Threshold Percentage: 50%
Autosize Mode: off
Autosize Enabled (for Flexvol volumes only): false
Total Files (for user-visible data): 217894
Files Used (for user-visible data): 98
Space Guarantee Style: volume
Space Guarantee In Effect: true
Snapshot Directory Access Enabled: true
Space Reserved for Snapshot Copies: 5%
Snapshot Reserve Used: 98%
Snapshot Policy: default
Creation Time: Mon Jul 08 10:54:32
2013
Language: C.UTF-8
Clone Volume: false

```

Node name: cluster-1-01
NVFAIL Option: off
Force NVFAIL on MetroCluster Switchover: off
Is File System Size Fixed: false
Extent Option: off
Reserved Space for Overwrites: 0B
Fractional Reserve: 100%
Primary Space Management Strategy: volume_grow
Read Reallocation Option: space-optimized
Inconsistency in the File System: false
Is Volume Quiesced (On-Disk): false
Is Volume Quiesced (In-Memory): false
Transition Operation State: none
Copied for Transition: false
Transitioned: true
Volume Contains Shared or Compressed Data: false
Efficiency Policy: default
UUID of the Efficiency Policy: b0f36cd7-e7bc-11e2-
9994-123478563412
Space Saved by Storage Efficiency: 0B
Percentage Saved by Storage Efficiency: 0%
Space Saved by Deduplication: 0B
Percentage Saved by Deduplication: 0%
Space Shared by Deduplication: 0B
Space Saved by Compression: 0B
Percentage Space Saved by Compression: 0%
Volume Size Used by Snapshot Copies: 1.48MB
Block Type: 64-bit
Is Volume Moving: false
Flash Pool Caching Eligibility: read-write
Flash Pool Write Caching Ineligibility Reason: -
Managed By Storage Service: -
Enable Object Store: -
Constituent Volume Role: -
Is cft precommit: false
QoS Policy Group Name: -
Caching Policy Name: auto
Is Volume Move in Cutover Phase: false
Number of Snapshot Copies in the Volume: 10
VBN_BAD may be present in the active filesystem: false
Is Eligible for Auto Balance Aggregate: -
Is Volume on a hybrid aggregate: false
Total Physical Used Size: 4.55MB
Physical Used Percentage: 14%
FlexGroup volume Master Data Set ID: -
FlexGroup volume Index: -

```

```

        UUID of the FlexGroup volume: -
Maximum size of a FlexGroup volume Constituent: -
        Inofile Version: 3
        List of Nodes: -
        Is Volume Flexgroup: false
        SnapLock Type: -
        Vserver DR Protection: -
        Healthy: true
        Unhealthy Reason:
        Is Fenced for Protocol Access: false
        Protocol Access Fence Owner: -
        Is SnapMirror active sync Master: false
        Is SnapMirror active sync Failover Capable: false
        SnapMirror active sync Consensus: -

```

Related Links

- [volume show-space](#)

volume size

Set/Display the size of the volume.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume size` command allows the user to set or display the volume size. If *new-size* is not specified then the current volume size is displayed.

Parameters

-vserver <vserver name> - Vserver Name

This parameter can be used to specify the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This parameter specifies the volume for which the user wants to set or display the size.

[-new-size {<integer>[KB|MB|GB|TB|PB] }] - [+|-]<New Size>

This optional parameter specifies the size of the volume. It can be used to set the volume size to a particular number or grow/shrink the size by a particular amount. The size is specified as a number (preceded with a sign for relative growth/shrinkage) followed by a unit designation: k (kilobytes), m (megabytes), g (gigabytes), or t (terabytes). If the unit designation is not specified, bytes are used as the unit, and the specified number is rounded up to the nearest 4 KB. The minimum size for a flexible volume is 20 MB, and the maximum size depends on hardware platform and free space in the containing aggregate. If the volume's space guarantee is currently disabled, its size cannot be increased.

Examples

The following example shows the size of a volume called vol1.

```
cluster1::> vol size vol1
(volume size)
vol size: Flexible volume 'vs1:vol1' has size 2g.
```

The following example sets the size of a volume called vol1 to 1GB.

```
cluster1::> vol size vol1 1g
(volume size)
vol size: Flexible volume 'vs1:vol1' size set to 1g.
```

The following example increases the size of a volume called vol1 by 500MB.

```
cluster1::> vol size vol1 +500m
(volume size)
vol size: Flexible volume 'vs1:vol1' size set to 1.49g.
```

The following example decreases the size of a volume called vol1 by 250MB.

```
cluster1::> vol size vol1 -250m
(volume size)
vol size: Flexible volume 'vs1:vol1' size set to 1.24g.
```

volume transition-prepare-to-downgrade

Verifies that there are no volumes actively transitioning from 7-mode to clustered Data ONTAP, and configures the transition feature for downgrade.

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `volume transition-prepare-to-downgrade` command is used to verify that a volume is not currently being transitioned from 7-Mode to clustered Data ONTAP. This check must be done before reverting or downgrading a node.

Parameters

Examples

The following example executes the `volume transition-inactive-verify` command during active

volume transitions.

```
cluster1::*> volume transition-prepare-to-downgrade
```

Error: command failed: Downgrade cannot proceed because one or more volumes are transitioning from 7-Mode to clustered Data ONTAP.

You must complete or cancel all transitions. View the transitioning volumes that are preventing the downgrade by using the (privilege: advanced) command "lun transition show -status active -fields vservers, volume".

To complete a transition, break the SnapMirror transition relationship using the command "snapmirror break -destination-path <destination-path>".

For detailed information about transitioning, refer to the "7-Mode Data Transition Using SnapMirror" guide.

To cancel a transition, contact technical support.

volume unmount

Unmount a volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume unmount` command unmounts a volume from its parent volume. The volume can be remounted at the same or a different location by using the [volume mount](#) command.

Parameters

-vservers <vservers name> - Vserver Name

This specifies the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This specifies the volume that is to be unmounted.

Examples

The following example unmounts a volume named vol2 on a Vserver named vs0:

```
node::> volume unmount -vservers vs0 -volume vol2
```

Related Links

- [volume mount](#)

volume activity-tracking commands

volume activity-tracking off

Disable activity tracking for a volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume activity-tracking off` command disables volume activity tracking on a volume.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver in which the volume is located.

-volume <volume name> - Volume Name

This specifies the name of the volume on which volume activity tracking is being disabled.

[-foreground <true>] - Run Operation in the Foreground

If true, the command doesn't complete until the volume activity tracking `state` is `off`. If false, the command returns immediately, and the volume activity tracking `state` can be tracked using the [volume activity-tracking show](#) command.

Examples

The following examples disable volume activity tracking on a volume.

```
cluster::*> volume activity-tracking off -vserver vs1 -volume fv1
```

```
cluster::*> volume activity-tracking off -vserver vs1 -volume fg1
```

Related Links

- [volume activity-tracking show](#)

volume activity-tracking on

Enable activity tracking for a volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume activity-tracking on` command enables volume activity tracking on a volume.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver in which the volume is located.

-volume <volume name> - Volume Name

This specifies the name of the volume on which volume activity tracking is being enabled.

[-foreground <true>] - Run Operation in the Foreground

If true, the command doesn't complete until the volume activity tracking *state* is *on* . If false, the command returns immediately, and the volume activity tracking *state* can be tracked using the [volume activity-tracking show](#) command.

Examples

The following examples enable volume activity tracking on a volume.

```
cluster::*> volume activity-tracking on -vserver vs1 -volume fv1
```

```
cluster::*> volume activity-tracking on -vserver vs1 -volume fg1
```

Related Links

- [volume activity-tracking show](#)

volume activity-tracking show

Display activity tracking information

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume activity-tracking show` command displays the volume activity tracking information for a volume. By default, the command displays the following information:

- Vserver name
- Volume name
- Activity tracking state

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[`-vserver <vserver name>`] - Vserver Name

If this parameter and the `-volume` parameter are specified, the command displays volume activity tracking information about the specified volume. If this parameter is specified by itself, the command displays information about volumes on the specified Vserver.

[`-volume <volume name>`] - Volume Name

If this parameter and the `-vserver` parameter are specified, the command displays volume activity tracking information about the specified volume. If this parameter is specified by itself, the command displays information about all volumes matching the specified name.

[`-state <Activity Tracking State>`] - State

If this parameter is specified, the command displays information only about the volume or volumes that have the specified volume activity tracking state.

Examples

The following examples display the volume activity tracking information for a volume.

```
cluster::*> volume activity-tracking show -vserver vs1 -volume fv1
Vserver Name: vs1
  Volume Name: fv1
    State: on
```

```
cluster::*> volume activity-tracking show -vserver vs1 -volume fg1
Vserver Name: vs1
  Volume Name: fg1
    State: on
```

volume analytics commands

volume analytics off

Disable collection of file system analytics for a volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume analytics off` command disables file system analytics on a volume.

Parameters

`-vserver <vserver name>` - Vserver Name

This specifies the Vserver in which the volume is located.

-volume <volume name> - Volume Name

This specifies the name of the volume on which file system analytics is being disabled.

[-foreground <true>] - Run Operation in the Foreground

If true, the command doesn't complete until the `analytics-state` of the volume is `off`. If false, the command returns immediately, and the `analytics-state` can be tracked using the [volume analytics show](#) command.

Examples

The following example disables file system analytics on a flexible volume named "flexvol" in Vserver "vs1":

```
cluster::*> volume analytics off -vserver vs1 -volume flexvol
```

Related Links

- [volume analytics show](#)

volume analytics on

Enable collection of file system analytics for a volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume analytics on` command enables file system usage analytics on a volume.

Parameters**-vserver <vserver name> - Vserver Name**

This specifies the Vserver in which the volume is located.

-volume <volume name> - Volume Name

This specifies the name of the volume on which file system usage analytics is being enabled.

[-foreground <true>] - Run Operation in the Foreground

If true, the command doesn't complete until the `analytics-state` of the volume is `on`. If false, the command returns immediately, and the `analytics-state` and `scan-progress` can be tracked using the [volume analytics show](#) command.

Examples

The following example enables file system usage analytics on a flexible volume named "flexvol" in Vserver "vs1":

```
cluster::*> volume analytics on -vserver vs1 -volume flexvol
```

Related Links

- [volume analytics show](#)

volume analytics show

Display file system analytics information

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume analytics show` command displays the file system analytics information for a volume. By default, the command displays the following information:

- Vserver name
- Volume name
- File system analytics state

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Name

If this parameter and the `-volume` parameter are specified, the command displays file system analytics information about the specified volume. If this parameter is specified by itself, the command displays information about volumes on the specified Vserver.

[-volume <volume name>] - Volume Name

If this parameter and the `-vserver` parameter are specified, the command displays file system analytics information about the specified volume. If this parameter is specified by itself, the command displays information about all volumes matching the specified name.

[-state <Analytics State>] - State

If this parameter is specified, the command displays information only about the volume or volumes that have the specified file system analytics state.

[-scan-progress <percent>] - Scan Progress

If this parameter is specified, the command displays information only about the volume or volumes that have the specified file system analytics progress percentage. This value will only be set if a file system analytics initialization scan is in progress.

[-tracked-dirs <integer>] - Tracked Directories

If this parameter is specified, the command displays information only about the volume or volumes that have the specified number of tracked directories.

[`-tracked-files <integer>`] - Tracked Files

If this parameter is specified, the command displays information only about the volume or volumes that have the specified number of tracked files.

[`-tracked-items <integer>`] - Tracked Items

If this parameter is specified, the command displays information only about the volume or volumes that have the specified number of tracked items. An item is either a file or a directory.

[`-top-level-dirs <integer>`] - Top Level Directories

If this parameter is specified, the command displays information only about the volume or volumes that have the specified number of top level directories.

[`-average-files-per-dir <integer>`] - Average Files Per Directory

If this parameter is specified, the command displays information only about the volume or volumes that have the specified average number of files per directory.

[`-files-scanned <integer>`] - Number of Files Scanned

If this parameter is specified, the command displays information only about the volume or volumes that have the specified number of files scanned. This value will only be set if a file system analytics initialization scan is in progress.

[`-total-files <integer>`] - Total Number of Files

If this parameter is specified, the command displays information only about the volume or volumes that have the specified total number of files. This value will only be set if a file system analytics initialization scan is in progress.

Examples

The following example displays the file system analytics information on a flexible volume named "flexvol" in Vserver "vs1":

```
cluster::*> volume analytics show -vserver vs1 -volume flexvol
Vserver Name: vs1
  Volume Name: flexvol
    State: on
  Scan Progress: -
```

volume analytics initialization pause

Pause the analytics file system scan

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume analytics initialization pause` command pauses the file system analytics scan.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver in which the volume is located.

-volume <volume name> - Volume Name

This specifies the name of the volume on which file system analytics initialization is being paused.

Examples

The following example pauses the file system analytics scan on a volume named "vol1" in Vserver "vs1":

```
cluster::*> volume analytics initialization pause -vserver vs1 -volume  
vol1
```

volume analytics initialization resume

Resume the analytics file system scan

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume analytics initialization resume` command resumes the file system analytics scan.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver in which the volume is located.

-volume <volume name> - Volume Name

This specifies the name of the volume on which file system analytics initialization is being resumed.

Examples

The following example resumes the file system analytics scan on a volume named "vol1" in Vserver "vs1":

```
cluster::*> volume analytics initialization resume -vserver vs1 -volume  
vol1
```

volume analytics initialization show

Display file system analytics information

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume analytics initialization show` command displays the file system analytics initialization information for a volume. By default, the command displays the following information:

- Vserver name
- Volume name
- File system analytics initialization state

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Name

If this parameter is specified, the command displays information only about the volumes in vservers that match the specified value for `-vserver`

[-volume <volume name>] - Volume Name

If this parameter is specified, the command displays information only about the volumes that match the specified value for `-volume`

[-state <Analytics Initialization State>] - State

If this parameter is specified, the command displays information only about the volume or volumes that have the specified file system analytics initialization state.

Examples

The following example displays the file system analytics initialization information on a volume named "vol1" in Vserver "vs1":

```
cluster::*> volume analytics initialization show -vserver vs1 -volume vol1
Vserver Name: vs1
  Volume Name: vol1
    State: Running
```

volume clone commands

volume clone create

Create a FlexClone volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume clone create` command creates a FlexClone volume on the aggregate containing the specified parent volume. This command is supported for flexible volumes or FlexGroups. The maximum volume clone hierarchy depth is 500 and the default depth is 60. You can optionally specify the following attributes for the new FlexClone volume:

- Vserver on which the parent volume resides
- Name of the FlexClone parent snapshot
- Junction path where FlexClone volume should be mounted
- State of the junction path
- Space guarantee style (none, volume or file)
- Comment
- Whether the `volume clone create` command runs as a foreground or background process
- Key Manager Attribute

Parameters

-vserver <vserver name> - Vserver Name

This parameter specifies the Vserver on which the FlexClone volume is to be created. If only one data Vserver exists, you do not need to specify this parameter.

-flexclone <volume name> - FlexClone Volume

This parameter specifies the name of the FlexClone volume. The name must be unique within the hosting Vserver.

[-type {RW|DP}] - FlexClone Type

This parameter specifies the type of FlexClone volume. A read-only FlexClone volume is created if you specify the `type` as DP; otherwise a read-write FlexClone volume is created.

[-parent-vserver <vserver name>] - FlexClone Parent Vserver

This parameter specifies the name of the Vserver to which the FlexClone parent volume belongs. If it is different from the Vserver on which the FlexClone volume is to be created, then the FlexClone volume inherits the export policies from the residing Vserver, and not from the FlexClone parent volume.

-b, -parent-volume <volume name> - FlexClone Parent Volume

This parameter specifies the name of parent volume from which the FlexClone clone volume is derived.

[-parent-snapshot <snapshot name>] - FlexClone Parent Snapshot

This specifies the name of the parent snapshot from which the FlexClone clone volume is derived.

[-junction-path <junction path>] - Junction Path

This specifies the junction path at which the new FlexClone clone volume should be mounted.

[-junction-active {true|false}] - Junction Active

This optionally specifies whether the volume's junction path is active. The default setting is `true`. If the junction path is inactive, the volume does not appear in the Vserver's namespace. This parameter is available only at the advanced privilege level and higher.

[-s, -space-guarantee {none|volume}] - Space Guarantee Style

This optionally specifies the space guarantee style for the FlexClone volume. A value of *volume* reserves space on the aggregate for the entire volume. A value of *none* reserves no space on the aggregate, meaning that writes can fail if the aggregate runs out of space. The default setting is inherited from the parent volume.

[-comment <text>] - Comment

This optionally specifies a comment for the FlexClone volume.

[-foreground {true|false}] - Foreground Process

This optionally specifies whether the FlexClone volume create operation runs as a foreground process. The default setting is *true* (that is, the operation runs in the foreground).

{ [-qos-policy-group <text>] - QoS Policy Group Name

This parameter optionally specifies which QoS policy group to apply to the FlexClone volume. The policy group defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated. If you do not assign a policy group to the FlexClone volume, the system does not monitor and control the traffic to the volume.

| [-qos-adaptive-policy-group <text>] - QoS Adaptive Policy Group Name }

This optionally specifies which QoS adaptive policy group to apply to the volume. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs) that adjust based on the volume allocated space or used space. This parameter is not supported on FlexGroups.

[-caching-policy <text>] - Caching Policy Name

This optionally specifies the caching policy to apply to the volume. A caching policy defines how the system caches this volume's data in Flash Cache modules. If a caching policy is not assigned to this volume, the system uses the caching policy that is assigned to the containing Vserver. If a caching policy is not assigned to the containing Vserver, the system uses the default cluster-wide policy. The available caching policies are:

- *none* - Does not cache any user data or metadata blocks.
- *auto* - Read caches all metadata and randomly read user data blocks, and write caches all randomly overwritten user data blocks.
- *meta* - Read caches only metadata blocks.
- *random_read* - Read caches all metadata and randomly read user data blocks.
- *random_read_write* - Read caches all metadata, randomly read and randomly written user data blocks.
- *all_read* - Read caches all metadata, randomly read, and sequentially read user data blocks.
- *all_read_random_write* - Read caches all metadata, randomly read, sequentially read and randomly written user data.
- *all* - Read caches all data blocks read and written. It does not do any write caching.
- *noread-random_write* - Write caches all randomly overwritten user data blocks. It does not do any read caching.
- *meta-random_write* - Read caches all metadata and write caches randomly overwritten user data blocks.
- *random_read_write-random_write* - Read caches all metadata, randomly read and randomly written user data blocks. It also write caches randomly overwritten user data blocks.

- `all_read-random_write` - Read caches all metadata, randomly read, and sequentially read user data blocks. It also write caches randomly overwritten user data blocks.
- `all_read_random_write-random_write` - Read caches all metadata, randomly read, sequentially read and randomly written user data. It also write caches randomly overwritten user data blocks.
- `all-random_write` - Read caches all data blocks read and written. It also write caches randomly overwritten user data blocks.

Note that in a caching-policy name, a hyphen (-) separates read and write policies. Default caching-policy is `auto`.

`[-vserver-dr-protection {protected|unprotected}] - Vserver DR Protection`

This optionally specifies whether the volume should be protected by Vserver level SnapMirror. This parameter is applicable only if the Vserver is the source of a Vserver level SnapMirror relationship. By default the clone volume will inherit this value from the parent volume.

`[-uid <integer>] - Volume-Level UID`

This parameter optionally specifies a volume-level user ID (UID). All files and directories in a FlexClone volume will inherit this UID.

`[-gid <integer>] - Volume-Level GID`

This parameter optionally specifies a volume-level group ID (GID). All files and directories in a FlexClone volume will inherit this GID.

`[-snaplock-type {non-snaplock|compliance|enterprise}] - SnapLock Type`

This parameter optionally specifies the SnapLock type for the FlexClone volume. The `snaplock-type` can be `compliance`, `enterprise` or `non-snaplock`. By default, the `snaplock-type` will be same as the parent volume.

`[-key-manager-attribute <text>] - Key Manager Attribute`

This optional parameter specifies an additional key manager attribute that should be an identifier-value pair separated by '=', ex. `CRN=my-unique-value`

The following identifiers are currently supported:

- CRN

Examples

The following command creates a FlexClone volume `fc_vol_1` from parent volume `fv2` on Vserver `vs1` and the job runs as a foreground process.

```

cluster1::> volume clone create -vserver vs1 -flexclone fc_vol_1 -parent
-volume fv2 -junction-active true -foreground true -comment "Testing
FlexClone creation"
cluster1::> volume clone show fc_vol_1 -vserver vs1
Vserver Name: vs1
    FlexClone Volume: fc_vol_1
    FlexClone Parent Volume: fv2
    FlexClone Parent Snapshot: clone_fc_vol_1.0
        Junction Path: -
        Junction Active: -
    Space Guarantee Style: volume
    Space Guarantee In Effect: true
    FlexClone Aggregate: test_aggr
    FlexClone Data Set ID: 1046
    FlexClone Master Data Set ID: 2147484694
        FlexClone Size: 19MB
            Used Size: 108KB
        Split Estimate: 0.00B
    Inodes processed: -
        Total Inodes: -
    Percentage complete: -
        Blocks Scanned: -
        Blocks Updated: -
        Comment: Testing FlexClone creation

```

volume clone show

Display a list of FlexClones

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume clone show` command displays information about FlexClone clone volumes. This command is only supported for flexible volumes. By default, the command displays the following information about all FlexClone volume clones:

- Vserver name
- FlexClone volume name
- Parent volume name
- Parent snapshot name
- Whether a FlexClone volume is online or offline

To display detailed information about all FlexClone volumes, run the command with the `-instance` parameter.

Parameters

{ [-fields <fieldname>, ...]

Selects the fields to be displayed.

| [-estimate]

Displays an estimate of the free disk space required in the aggregate to split the indicated clone volume from its underlying parent volume. The value reported may differ from the space actually required to perform the split, especially if the clone volume is changing when the split is being performed.

| [-instance] }

Displays detailed information about FlexClone volumes. If `-flexclone` is also specified, the command displays detailed information about the FlexClone volume.

[-vserver <vserver name>] - Vserver Name

Selects summary information for the FlexClone volumes on the specified Vserver. If `-flexclone` is also specified, the command displays detailed information about the specified FlexClone volume.

[-flexclone <volume name>] - FlexClone Volume

Selects summary information for the specified FlexClone volume. If `-vserver` is also specified, the command displays detailed information about the specified FlexClone volume.

[-type {RW|DP}] - FlexClone Type

Selects information for the specified type of FlexClone volume. The type can be specified as either read-only (DP) or read-write (RW).

[-parent-vserver <vserver name>] - FlexClone Parent Vserver

Selects summary information for the FlexClone volumes that are clone volumes in the specified parent Vserver.

[-b, -parent-volume <volume name>] - FlexClone Parent Volume

Selects summary information for the FlexClone volumes that are clones of the specified parent volume.

[-parent-snapshot <snapshot name>] - FlexClone Parent Snapshot

Selects summary information for the FlexClone volumes that are clones of the parent volume to which the specified snapshot belongs.

[-state {online|restricted|offline|force-online|force-offline|mixed}] - FlexClone Volume State

Selects summary information for the FlexClone volumes that are in the specified state.

[-junction-path <junction path>] - Junction Path

Selects summary information for the FlexClone volumes that have the specified junction path.

[-junction-active {true|false}] - Junction Active

Selects summary information for the FlexClone volumes that have the specified junction path status.

[-s, -space-guarantee {none|volume}] - Space Guarantee Style

If this parameter is specified, the command displays information only about the volumes that have the specified space guarantee style.

[-space-guarantee-enabled {true|false}] - Space Guarantee In Effect

Selects summary information for the FlexClone volumes that have the specified space-guarantee setting.

[-aggregate <aggregate name>] - FlexClone Aggregate

Selects summary information for the FlexClone volumes that reside on the specified storage aggregate.

[-dsid <integer>] - FlexClone Data Set ID

Selects summary information for the FlexClone volumes that have the specified Data Set ID.

[-msid <integer>] - FlexClone Master Data Set ID

Selects summary information for the FlexClone volumes that have the specified Master Data Set ID.

[-size {<integer>[KB|MB|GB|TB|PB]}] - FlexClone Size

Selects summary information for the FlexClone volumes that have the specified size.

[-used {<integer>[KB|MB|GB|TB|PB]}] - Used Size

Selects summary information for the FlexClone volumes that have the specified amount of used space.

[-split-estimate {<integer>[KB|MB|GB|TB|PB]}] - Split Estimate

Selects summary information for the FlexClone volumes that require the specified amount of free disk space for splitting from the parent.

[-blocks-scanned <integer>] - Blocks Scanned

Selects summary information for the FlexClone volumes that have the specified number of blocks scanned for splitting the FlexClone volume from its parent volume.

[-blocks-updated <integer>] - Blocks Updated

Selects summary information for the FlexClone volumes that have the specified number of blocks updated for after splitting the FlexClone volume from its parent volume.

[-comment <text>] - Comment

Selects summary information for the FlexClone volumes that have the specified comment.

[-qos-policy-group <text>] - QoS Policy Group Name

Selects summary information for the FlexClone volumes that have the specified QoS policy group.

[-qos-adaptive-policy-group <text>] - QoS Adaptive Policy Group Name

Selects summary information for the FlexClone volumes that have the specified QoS adaptive policy group.

[-caching-policy <text>] - Caching Policy Name

This optionally specifies the caching policy to apply to the volume. A caching policy defines how the system caches this volume's data in Flash Cache modules. If a caching policy is not assigned to this volume, the system uses the caching policy that is assigned to the containing Vserver. If a caching policy is not assigned to the containing Vserver, the system uses the default cluster-wide policy. The available caching policies are:

- none - Does not cache any user data or metadata blocks.
- auto - Read caches all metadata and randomly read user data blocks, and write caches all randomly overwritten user data blocks.

- meta - Read caches only metadata blocks.
- random_read - Read caches all metadata and randomly read user data blocks.
- random_read_write - Read caches all metadata, randomly read and randomly written user data blocks.
- all_read - Read caches all metadata, randomly read, and sequentially read user data blocks.
- all_read_random_write - Read caches all metadata, randomly read, sequentially read and randomly written user data.
- all - Read caches all data blocks read and written. It does not do any write caching.
- noread-random_write - Write caches all randomly overwritten user data blocks. It does not do any read caching.
- meta-random_write - Read caches all metadata and write caches randomly overwritten user data blocks.
- random_read_write-random_write - Read caches all metadata, randomly read and randomly written user data blocks. It also write caches randomly overwritten user data blocks.
- all_read-random_write - Read caches all metadata, randomly read, and sequentially read user data blocks. It also write caches randomly overwritten user data blocks.
- all_read_random_write-random_write - Read caches all metadata, randomly read, sequentially read and randomly written user data. It also write caches randomly overwritten user data blocks.
- all-random_write - Read caches all data blocks read and written. It also write caches randomly overwritten user data blocks.

Note that in a caching-policy name, a hyphen (-) separates read and write policies. Default caching-policy is auto.

[-parent-vol-type <volAccessType>] - Parent volume type (privilege: advanced)

Selects summary information for the FlexClone volumes that are clones of the parent volumes with the specified type.

[-flexclone-used-percent <percent>] - FlexClone Used Percentage

Selects summary information for the FlexClone volumes that have the specified percentage of used space.

[-vserver-dr-protection {protected|unprotected}] - Vserver DR Protection

Selects summary information for the FlexClone volumes that have the specified type of Vserver SnapMirror protection. This parameter is applicable only if the Vserver is the source of a Vserver level SnapMirror relationship.

[-block-percentage-complete <integer>] - Percentage Complete

Selects summary information for the FlexClone volumes that have specified percentage of Blocks processed for splitting the FlexClone volume from its parent volume.

[-uid <integer>] - Volume-Level UID

Selects summary information for the FlexClone volumes that are created with the specified volume-level UID.

[-gid <integer>] - Volume-Level GID

Selects summary information for the FlexClone volumes that are created with the specified volume-level GID.

[`-flexgroup-uuid <UUID>`] - UUID of the FlexGroup

If this parameter is specified, the command displays information only about the FlexGroup volume or FlexGroup volume constituents that have the specified FlexGroup volume UUID. This parameter applies to FlexGroup volumes and FlexGroup volume constituents only.

[`-flexgroup-msid <integer>`] - FlexGroup Master Data Set ID

If this parameter is specified, the command displays information only about the FlexGroup volume or FlexGroup volume constituents that have the specified FlexGroup volume master data-set ID. This parameter applies to FlexGroup volumes and FlexGroup volume constituents only.

[`-flexgroup-index <integer>`] - FlexGroup Index

If this parameter is specified, the command displays information only about the FlexGroup volume constituents that have the specified FlexGroup volume index. This parameter applies to FlexGroup volume constituents only.

[`-max-constituent-size {<integer>[KB|MB|GB|TB|PB]}`] - Maximum size of a FlexGroup Constituent

If this parameter is specified, the command displays information only about the FlexGroup volume or FlexGroup volumes that have the specified maximum constituent size. This parameter applies to FlexGroup volumes only.

[`-constituent-role <Constituent Roles>`] - Constituent Volume Role

If this parameter is specified, the command displays information only about the constituent volume or volumes that are of the specified role. This parameter applies to FlexGroup volumes only.

[`-is-flexgroup-constituent-active {true|false}`] - Is Active FlexGroup Constituent

If this parameter is specified, the command displays information only about the volume or volumes that are either active constituent of the FlexGroup or not, depending on the value provided.

[`-is-constituent {true|false}`] - Is Constituent Volume

If this parameter is specified, the command displays information only about volumes that either are or are not constituents of a FlexGroup, depending on the value provided.

[`-is-flexgroup {true|false}`] - Is Volume a FlexGroup

If this parameter is specified, the command displays information only about the volume or volumes that are either FlexGroups or not, depending on the value provided.

[`-volume-style-extended {flexvol|flexgroup|flexgroup-constituent}`] - Extended Volume Style

If this parameter is specified, the command displays information only about the volumes that are configured with the specified extended style. Possible values are *flexvol* for FlexVol volumes, *flexgroup* for FlexGroups and *flexgroup-constituent* for FlexGroup constituents.

[`-snaplock-type {non-snaplock|compliance|enterprise}`] - SnapLock Type

Selects summary information for the FlexClone volumes that are created with the specified SnapLock type.

Examples

The following example displays detailed information about all FlexClone volumes on Vserver *vs0* :

```

cluster1::> volume clone show -vserver vs0
(volume clone show)
Vserver    FlexClone    Parent-Volume    Parent-Snapshot
-----
vs0        fc_vol_1     test_vol         clone_fc_vol_1.0
          fc_vol_2     test_vol2        clone_fc_vol_2.0
          fc_vol_3     tv9              clone_fc_vol_3.0
          tv8         tv7              clone_tv8.0
          tv9         test_vol2        clone_tv9.0
5 entries were displayed.

```

The following example displays detailed information about FlexClone volume *fc_vol_2* on Vserver *vs0*:

```

cluster1::> volume clone show -vserver vs0 -flexclone fc_vol_2
Vserver Name: vs0
    FlexClone Volume: fc_vol_2
    FlexClone Parent Volume: test_vol2
    FlexClone Parent Snapshot: clone_fc_vol_2.0
    Junction Path: -
    Junction Active: -
    Space Guarantee Style: volume
    Space Guarantee In Effect: true
    FlexClone Aggregate: test_aggr
    FlexClone Data Set ID: 1038
FlexClone Master Data Set ID: 2147484686
    FlexClone Size: 47.50MB
    Used Size: 128KB
    Split Estimate: 0.00B
    Inodes processed: -
    Total Inodes: -
    Inode Percentage complete: -
    Blocks Scanned: -
    Blocks Updated: -
    Comment:
    Qos Policy Group Name: pg1
FlexClone Parent Volume Type: RW
    Block Percentage complete: -

```

The following example displays summary information about all FlexClone volumes residing on Vserver *vs0* along with the fields *msid*, *dsid*, *state* and *parent-volume*.

```

cluster1::> volume clone show -vserver vs0 -fields msid, dsid, state,
parent-volume
vserver flexclone parent-volume state dsid msid
-----
vs0      fc_vol_1  test_vol      online 1037 2147484685
vs0      fc_vol_3  tv9           online 1039 2147484687
vs0      flex_clone1
                fc_vol_1      online 1041 2147484689
vs0      fv_2      fc_vol_1      online 1043 2147484691
vs0      tv9      test_vol2     online 1036 2147484684
5 entries were displayed.

```

The following example displays summary information about all FlexClone volumes residing on Vserver *vs0* along with `space-guarantee-enabled` and `space-guarantee` information about each FlexClone volume.

```

cluster1::> vol clone show -vserver vs0 -fields space-guarantee-enabled,
space-guarantee
(volume clone show)
vserver flexclone space-guarantee space-guarantee-enabled
-----
vs0      fc_vol_1  volume        true
vs0      fc_vol_3  volume        true
vs0      flex_clone1
                volume        true
vs0      fv_2      volume        true
vs0      tv9      volume        true
5 entries were displayed.

```

volume clone sharing-by-split show

Show the split flexclone volumes with shared physical blocks

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `volume clone sharing-by-split show` command displays the split volumes with shared physical blocks. This command is only supported for flexible volumes. By default, this command displays the following information:

- Node Name
- Vserver Name
- Volume Name
- Aggregate Name

- Volume State

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-node {<nodename>|local}] - Node Name (privilege: advanced)

This parameter selects information about the split volumes with shared physical blocks on this node.

[-vserver <Vserver Name>] - Vserver Name (privilege: advanced)

This parameter selects information about the split volumes with shared physical blocks on this Vserver.

[-volume <volume name>] - Volume Name (privilege: advanced)

This parameter selects information about shared physical blocks for this volume.

[-aggregate <aggregate name>] - Aggregate Name (privilege: advanced)

This parameter specifies the aggregate associated with the given volume.

Examples

The following example displays the split volumes with shared physical blocks in the node:

```
cluster1::> volume clone sharing-by-split show -node node1
Node           Vserver    Volume           Aggregate
-----
node1          vs1        vol_clone1       aggr1
```

The following example displays information about volume `vol_clone1` residing on vserver `vs1`:

```
cluster1::> volume clone sharing-by-split show -node node1 -vserver vs1
-volume vol_clone1 -instance
Node Name: node1
          Vserver Name: vs1
          Volume Name: vol_clone1
          Aggregate Name: aggr1
```

volume clone sharing-by-split undo show

Show the status of volume clone undo-sharing operations in-progress

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `volume clone sharing-by-split undo show` command displays the progress information of undo-sharing in the split volumes with shared physical blocks. This command is only supported for flexible volumes. By default, the command displays the following information:

- Vserver name
- Volume name
- Total number of blocks scanned for undo sharing
- Total number of blocks present
- Percentage of blocks processed

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <Vserver Name>] - Vserver Name (privilege: advanced)

This parameter selects information about the ongoing undo-sharing scan for all volumes on this Vserver.

[-volume <volume name>] - Volume Name (privilege: advanced)

This parameter selects information about the ongoing undo-sharing scan on this volume.

[-blocks-scanned <integer>] - Scanned Blocks (privilege: advanced)

This parameter selects information about the total number of blocks scanned by undo-sharing in the given volume.

[-blocks-total <integer>] - Total Blocks (privilege: advanced)

This parameter selects information about the total number of blocks for the undo-sharing to scan in the given volume.

[-blocks-percentage-complete <integer>] - Blocks Percentage Complete (privilege: advanced)

This parameter selects information about the percentage of block processing completed by undo-sharing in the given volume.

Examples

The following example displays information about all the ongoing undo-sharing scan in the cluster:

```
cluster1::> volume clone sharing-by-split undo show
                Blocks      Blocks      Blocks
Vserver   Volume      Scanned      Total  % Complete
-----
vs1       vol_clone1      0           1260      0
```

The following example displays information about volume `vol_clone1` residing on vserver `vs1` :

```
cluster1::> volume clone sharing-by-split undo show -vserver vs1 -volume
vol_clone1 -instance
Vserver Name: vs1
    Volume Name: vol_clone1
        Blocks Scanned: 0
            Blocks Total: 1260
Block Percentage complete: 0
```

volume clone sharing-by-split undo start-all

Undo the physical block sharing in split FlexClone volumes across the cluster

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `volume clone sharing-by-split undo start-all` command starts a scan to undo the shared physical blocks in all the relevant volumes across the cluster. The volumes will be available for the duration of the undo-sharing operation. You can monitor the current progress of the scan using the [volume clone sharing-by-split undo show](#) command. This command is supported for flexible volumes that were split from their parent volumes.

Examples

The following example starts the scan to undo the physical block sharing in all volumes across the cluster:

```
cluster1::> volume clone sharing-by-split undo start-all
```

Related Links

- [volume clone sharing-by-split undo show](#)

volume clone sharing-by-split undo start

Undo the physical block sharing in split FlexClone volumes

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume clone sharing-by-split undo start` command starts a scan to undo the shared physical blocks in the given volume. The volume will be available for the duration of the undo-sharing operation. After the scan starts, you can stop it using the [volume clone sharing-by-split undo stop](#) command. You can monitor the current progress of the scan using the [volume clone sharing-by-split undo show](#) command. This command is supported for flexible volumes that were split from their parent volumes.

Parameters

-vserver <Vserver Name> - Vserver Name (privilege: advanced)

This parameter specifies the vserver that the volume exists on.

-volume <volume name> - Volume Name (privilege: advanced)

This parameter specifies the split volume with shared physical blocks, in which the sharing will be undone.

Examples

The following example starts the scan to undo the physical block sharing in volume `vol_clone1` on vserver `vs1`:

```
cluster1::> volume clone sharing-by-split undo start -vserver vs1 -volume
vol_clone1
```

Related Links

- [volume clone sharing-by-split undo stop](#)
- [volume clone sharing-by-split undo show](#)

volume clone sharing-by-split undo stop

Stop an ongoing undo-sharing operation

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume clone sharing-by-split undo stop` command stops the process of reverting the shared physical blocks from the split volume. If you restart the undo-sharing operation, scan begins from the beginning because no information about previously achieved progress is saved, but previously unshared blocks are not processed again. This command is only supported for flexible volumes.

Parameters

-vserver <Vserver Name> - Vserver Name (privilege: advanced)

This parameter specifies the vserver that the volume exists on.

-volume <volume name> - Volume Name (privilege: advanced)

This parameter specifies the volume whose unsharing of blocks will be stopped.

Examples

The following example stops an ongoing undo-sharing scan for volume `vol_clone1` on vserver `vs1`:

```
cluster1::> volume clone sharing-by-split undo stop -vserver vs1 -volume
vol_clone1
```


volume clone split estimate

Estimates the space required by the containing-aggregate to split the FlexClone volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume clone split estimate` command displays an estimate of the free disk space required in the aggregate to split the indicated clone volume from its underlying parent volume. The value reported might differ from the space actually required to perform the split, especially if the clone volume is changing when the split is being performed. This command is only supported for flexible volumes.

Parameters

[-vserver <vserver name>] - Vserver Name

This specifies the estimates for free disk space required for splitting FlexClone volumes residing on this Vserver. If the `-flexclone` option is also specified, then the command displays the free disk space estimate only for the specified FlexClone volume residing on the specified Vserver.

[-flexclone <volume name>] - FlexClone Volume

This specifies the free disk space estimate for splitting this FlexClone volume.

[-type {RW|DP}] - FlexClone Type

This parameter specifies the type of FlexClone volume. A read-only FlexClone volume is created if you specify the `type` as `DP`; otherwise a read-write FlexClone volume is created.

[-parent-vserver <vserver name>] - FlexClone Parent Vserver

This specifies the free disk space estimates for splitting the FlexClone volumes that are clones in the specified parent Vserver.

[-b, -parent-volume <volume name>] - FlexClone Parent Volume

This specifies the free disk space estimates for splitting the FlexClone volumes cloned off this parent volume.

[-parent-snapshot <snapshot name>] - FlexClone Parent Snapshot

This specifies the free disk space estimates for splitting the FlexClone volumes cloned off this parent snapshot.

[-state {online|restricted|offline|force-online|force-offline|mixed}] - FlexClone Volume State

This specifies the free disk space estimates for splitting the FlexClone volumes with the specified state.

[-junction-path <junction path>] - Junction Path

This specifies the free disk space estimates for splitting the FlexClone volumes mounted at this junction path.

[-junction-active {true|false}] - Junction Active

If this specified, the command displays the free disk space estimate for splitting the FlexClone volumes with the specified junction path status.

[-s, -space-guarantee {none|volume}] - Space Guarantee Style

This specifies the free disk space estimates for splitting the FlexClone volumes with the specified type of space guarantee.

[-space-guarantee-enabled {true|false}] - Space Guarantee In Effect

This specifies the free disk space estimates for splitting the FlexClone volumes with the specified state of space guarantee.

[-aggregate <aggregate name>] - FlexClone Aggregate

This specifies the free disk space estimates for splitting the FlexClone volumes residing on the specified aggregate.

[-dsid <integer>] - FlexClone Data Set ID

This specifies the free disk space estimates for splitting the FlexClone volume with the specified DSID (data set ID).

[-msid <integer>] - FlexClone Master Data Set ID

This specifies the free disk space estimates for splitting the FlexClone volumes with the specified MSID (master data set ID).

[-size {<integer>[KB|MB|GB|TB|PB]}] - FlexClone Size

This specifies the free disk space estimates for splitting FlexClone volumes with the specified size.

[-used {<integer>[KB|MB|GB|TB|PB]}] - Used Size

This specifies the free disk space estimates for splitting the FlexClone volumes with the specified amount of used disk space.

[-split-estimate {<integer>[KB|MB|GB|TB|PB]}] - Split Estimate

This specifies the free disk space estimates for splitting the FlexClone volumes which match with the specified free disk space estimate for splitting.

[-blocks-scanned <integer>] - Blocks Scanned

This specifies the free disk space estimates for splitting the FlexClone volumes for which the specified number of blocks have been scanned.

[-blocks-updated <integer>] - Blocks Updated

This specifies the free disk space estimates for splitting the FlexClone volumes for which the specified number of blocks have been updated.

[-comment <text>] - Comment

This specifies the free disk space estimates for splitting the FlexClone volumes that have the specified value for the comment field.

[-qos-policy-group <text>] - QoS Policy Group Name

This parameter optionally specifies which QoS policy group to apply to the FlexClone volume. The policy group defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated. If you do not assign a policy group to the FlexClone volume, the system does not monitor and control the traffic to the volume.

[-qos-adaptive-policy-group <text>] - QoS Adaptive Policy Group Name

This optionally specifies which QoS adaptive policy group to apply to the volume. This policy group defines measurable service level objectives (SLOs) and Service Level agreements (SLAs) that adjust based on the volume allocated space or used space. This parameter is not supported on FlexGroups.

[-caching-policy <text>] - Caching Policy Name

This specifies the free disk space estimates for splitting the FlexClone volumes that are clones with the specified caching policy.

[-parent-vol-type <volAccessType>] - Parent volume type (privilege: advanced)

This specifies the free disk space estimates for splitting the FlexClone volumes that are clones of the parent volumes with the specified type.

[-flexclone-used-percent <percent>] - FlexClone Used Percentage

This specifies the free disk space estimates for splitting the FlexClone volumes that are clones with the specified percentage of used space.

[-vserver-dr-protection {protected|unprotected}] - Vserver DR Protection

This specifies the free disk space estimates for splitting the FlexClone volumes that are clones with the specified Vserver SnapMirror protection.

[-block-percentage-complete <integer>] - Percentage Complete

This specifies the free disk space estimates for splitting the FlexClone volumes for which the specified percentage of Block processing has been completed.

[-uid <integer>] - Volume-Level UID

This specifies the free disk space estimates for splitting the FlexClone volumes that are clones with the specified volume-level UID.

[-gid <integer>] - Volume-Level GID

This specifies the free disk space estimates for splitting the FlexClone volumes that are clones with the specified volume-level GID.

[-flexgroup-uuid <UUID>] - UUID of the FlexGroup

If this parameter is specified, the command displays information only about the FlexGroup volume or FlexGroup volume constituents that have the specified FlexGroup volume UUID. This parameter applies to FlexGroup volumes and FlexGroup volume constituents only.

[-flexgroup-msid <integer>] - FlexGroup Master Data Set ID

If this parameter is specified, the command displays information only about the FlexGroup volume or FlexGroup volume constituents that have the specified FlexGroup volume master data-set ID. This parameter applies to FlexGroup volumes and FlexGroup volume constituents only.

[-flexgroup-index <integer>] - FlexGroup Index

If this parameter is specified, the command displays information only about the FlexGroup volume constituents that have the specified FlexGroup volume index. This parameter applies to FlexGroup volume constituents only.

[-max-constituent-size {<integer>[KB|MB|GB|TB|PB]}] - Maximum size of a FlexGroup Constituent

If this parameter is specified, the command displays information only about the FlexGroup volume or

FlexGroup volumes that have the specified maximum constituent size. This parameter applies to FlexGroup volumes only.

[-constituent-role <Constituent Roles>] - Constituent Volume Role

If this parameter is specified, the command displays information only about the constituent volume or volumes that are of the specified role. This parameter applies to FlexGroup volumes only.

[-is-flexgroup-constituent-active {true|false}] - Is Active FlexGroup Constituent

If this parameter is specified, the command displays information only about the volume or volumes that are either active constituent of the FlexGroup or not, depending on the value provided.

[-is-constituent {true|false}] - Is Constituent Volume

If this parameter is specified, the command displays information only about volumes that either are or are not constituents of a FlexGroup, depending on the value provided.

[-is-flexgroup {true|false}] - Is Volume a FlexGroup

If this parameter is specified, the command displays information only about the volume or volumes that are either FlexGroups or not, depending on the value provided.

[-volume-style-extended {flexvol|flexgroup|flexgroup-constituent}] - Extended Volume Style

If this parameter is specified, the command displays information only about the volumes that are configured with the specified extended style. Possible values are *flexvol* for FlexVol volumes, *flexgroup* for FlexGroups and *flexgroup-constituent* for FlexGroup constituents.

[-snaplock-type {non-snaplock|compliance|enterprise}] - SnapLock Type

If this parameter is specified, the command displays information only about the volumes that are configured with the specified SnapLock type. Possible values are *compliance* for SnapLock Compliance volumes, *enterprise* for SnapLock Enterprise volumes and *non-snaplock* for non-SnapLock volumes.

Examples

The following example displays the FlexClone split free disk space estimates for the FlexClone volumes residing on Vserver *vs0*.

```
cluster1::> volume clone split estimate -vserver vs0
(volume clone split estimate)

```

Vserver	FlexClone	Split Estimate
vs0	fc_vol_1	851.5MB
	fc_vol_3	0.00B
	flex_clone1	350.3MB
	fv_2	47.00MB
	tv9	0.00B

5 entries were displayed.

volume clone split show

Show the status of FlexClone split operations in-progress

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume clone split show` command displays the progress information of all the active FlexClone volume splitting jobs. This command is only supported for flexible volumes. By default, this command displays the following information:

- Vserver name
- FlexClone volume name
- Percentage of blocks processed
- Total number of blocks scanned for clone splitting
- Total number of blocks updated for clone splitting

If the `-instance` option is also specified, detailed information about all splitting jobs is displayed.

Parameters

{ [-fields <fieldname>,...]

This specifies the fields to be displayed, for all the ongoing FlexClone splitting jobs.

| [-instance] }

This specifies the command to display detailed information about the ongoing FlexClone volume splitting jobs.

[-vserver <vserver name>] - Vserver Name

Selects information about the ongoing FlexClone volume splitting jobs for all FlexClone volumes on this Vserver.

[-flexclone <volume name>] - FlexClone Volume

Selects information about ongoing FlexClone volume splitting jobs for this FlexClone volume.

[-block-percentage-complete <integer>] - Percentage Complete

Selects information about all the ongoing FlexClone splitting jobs that have the specified percentage of Block processing completed.

[-blocks-scanned <integer>] - Blocks Scanned

Selects information about all the ongoing FlexClone splitting jobs that have the specified number of blocks scanned.

[-blocks-updated <integer>] - Blocks Updated

Selects information about all the ongoing FlexClone splitting jobs that have the specified number of blocks updated.

Examples

The following example displays information about all the ongoing FlexClone splitting jobs in the cluster.

```
cluster1::> volume clone split show
(volume clone split show)

```

Vserver	FlexClone	Blocks		% Complete
		Scanned	Updated	
vs1	fc_vol_1	229	2	0

The following example displays information about FlexClone volume `fc_vol_2` residing on Vserver `vs0`.

```
cluster1::> volume clone split show -vserver vs0 -flexclone fc_vol_2
-instance
(volume clone split show)
Vserver Name: vs0
FlexClone Volume: fc_vol_2
Percentage Complete: 0
Blocks Scanned: 229
Blocks Updated: 2
```

volume clone split start

Split a FlexClone from the parent volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume clone split start` command starts a job to separate the FlexClone volume from the underlying parent volume. Both, the parent and the FlexClone volumes will be available for the duration of the split operation. After the job starts, you can stop it using the `volume clone split stop` command. You can also stop the job using the `job stop` command. You can monitor the current progress of the job using the `volume clone split show` and `job show` commands. This command is only supported for flexible volumes. This command is not supported on volumes that are being protected as part of a Vserver level SnapMirror.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver that the FlexClone volume exists on.

-flexclone <volume name> - FlexClone Volume

This specifies the FlexClone volume that will be split from its parent volume.

[`-foreground <true>`] - Foreground Process

This specifies whether the clone splitting job will run as a foreground job. The default value of this option is `true`.

[`-match-parent-storage-tier {true|false}`] - Match Split Data to Parent Storage Tier

This specifies whether the FlexClone volume splits the data blocks by matching its parent storage tier. This option is applicable only if the tiering policy and the tiering minimum cooling days of the parent volume and the FlexClone volume are the same. The default value of this option is `false`.

Examples

The following example starts splitting FlexClone volume `fc_vol_1` on Vserver `vs1` as a foreground job.

```
cluster1::> volume clone split start -vserver vs1 -flexclone fc_vol_1
-foreground true
```

Related Links

- [volume clone split stop](#)
- [job stop](#)
- [volume clone split show](#)
- [job show](#)

volume clone split stop

Stop an ongoing FlexClone split job

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume clone split stop` command stops the process of separating the FlexClone volume from its underlying parent volume, but does not lose any of the progress achieved while the split process was active. That is, all the clone volume blocks already separated from the parent volume remain separated. If you restart the split operation, splitting process begins from the beginning because no information about previously achieved progress is saved, but previously split blocks are not re-split. This command is only supported for flexible volumes.

Parameters

`-vserver <vserver name>` - Vserver Name

This specifies the Vserver that the FlexClone volume exists on.

`-flexclone <volume name>` - FlexClone Volume

This specifies the FlexClone volume whose separation from the parent volume will be stopped.

Examples

The following example stops an ongoing clone splitting job for FlexClone volume `fc_vol_1` on Vserver `vs1`.

```
cluster1::> volume clone split stop -vserver vs1 -flexclone fc_vol_1
```

volume conversion commands

volume conversion start

Convert a volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume conversion start` command converts a volume.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

This specifies the Vserver in which the volume is located.

-volume <volume name> - Volume Name (privilege: advanced)

This specifies the name of the volume to convert.

[-check-only <true>] - Validate the Conversion Only (privilege: advanced)

If true, this specifies whether the command will only validate that the conversion can be done, and that the volume will not be converted. The default value is false.

[-foreground <true>] - Foreground Process (privilege: advanced)

This specifies whether the operation runs in the foreground. The default setting is `true` (the operation runs in the foreground). When set to true, the command will not return until the operation completes.

Examples

The following example validates and converts a flexible volume named "flexvol" in Vserver "vs1" to a FlexGroup:

```
cluster::*> volume conversion start -vserver vs1 -volume flexvol
```

The following example validates whether flexible volume named "flexvol" in Vserver "vs1" can be converted to a FlexGroup:

```
cluster::*> volume conversion start -vserver vs1 -volume flexvol -check  
-only
```


volume conversion validation show

Show the result of volume conversion pre checks.

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `volume conversion validation show` command displays the results of volume conversion validation for volumes.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver (privilege: advanced)

If this parameter is specified, the command displays information only about the conversion operations that are running on volumes on the specified Vserver.

[-volume <volume name>] - Volume (privilege: advanced)

If this parameter is specified, the command displays information only about the conversion operations on the specified volume.

[-precheck <Volume conversions pre check>] - Volume Feature (privilege: advanced)

If this parameter is specified, the command displays information only about the specified conversion pre checks.

[-status <Feature status for volume conversion>] - Feature Status (privilege: advanced)

If this parameter is specified, the command displays information only about the conversion pre checks with the specified status. The supported statuses are:

- supported - The volume can be converted.
- warn - The volume can be converted, though there may be an issue highlighted by a warning in the `error` value.
- block - The volume cannot be converted until the specified `action` is taken
- unsupported - The volume cannot be converted.
- error - There was an error performing the pre check.

[-error <text>] - Feature Error (privilege: advanced)

If this parameter is specified, the command displays information only about the conversion pre checks with the specified error message.

[-action <text>] - Feature Action (privilege: advanced)

If this parameter is specified, the command displays information only about the onversion pre checks with

the specified resolution action.

Examples

The following example displays all conversion pre checks

```
cluster::*> volume conversion validation show
```

The following example displays all conversion pre checks that block conversion until an action is taken, for volume 'vol1' in Vserver 'vs0'

```
cluster::*> volume conversion validation show -vserver vs0 -volume vol1  
-status block
```

volume efficiency commands

volume efficiency check

Scrub efficiency metadata of a volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

This command verifies and updates the fingerprint database for the specified volume. This command is not supported on FlexGroups or Infinite Volumes that are managed by storage services.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

Specifies the Vserver on which the volume is located.

{ -volume <volume name> - Volume Name (privilege: advanced)

Specifies the volume on which the verify operation needs to be started.

| -path </vol/volume> - Volume Path (privilege: advanced) }

Specifies the volume path on which the verify operation needs to be started.

[-d, -delete-checkpoint {true|false}] - Delete Checkpoint (privilege: advanced)

Deletes existing checkpoint.

Examples

The following example runs `volume efficiency check` with delete checkpoint option turned on.

```
cluster1::*> volume efficiency check -vserver vs1 -volume voll -delete
-checkpoint true
```

volume efficiency modify

Modify the efficiency configuration of a volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver on which the volume is located.

{ -volume <volume name> - Volume Name

This specifies the volume on which efficiency options need to be modified.

| -path </vol/volume> - Volume Path }

This specifies the volume path on which efficiency options need to be modified.

{ [-schedule <text>] - Schedule

This option is used to set and modify the schedule.

schedule is [day_list][@hour_list] or [hour_list][@day_list] or - or auto or manual

The *day_list* specifies the days of the week that an efficiency operation should run. It is a list of the first three letters of the day (sun, mon, tue, wed, thu, fri, sat), separated by a comma. Day ranges such as mon-fri can also be used. The default *day_list* is sun-sat. The names are not case sensitive.

The *hour_list* specifies the hours of each scheduled day that an efficiency operation should run. The *hour_list* is from 0 to 23, separated by a comma. Hour ranges such as 8-17 are allowed. Step values can be used in conjunction with ranges (For example, 0-23/2 means every two hours in a day). The default *hour_list* is 0, i.e. at midnight of each scheduled day.

When efficiency is enabled on a volume for the first time, an initial schedule is assigned to the volume. This initial schedule is sun-sat@0, which means run once every day at midnight.

If "-" is specified, no schedule is set on the volume. The auto schedule string triggers an efficiency operation depending on the amount of new data written to the volume. The manual schedule string prevents SIS from automatically triggering any operations and disables change-logging. This schedule string can only be used on SnapVault destination volumes. The use of this schedule is mainly desirable when inline compression is enabled on a SnapVault destination volume and background processing is not necessary.

Note that schedule and policy are mutually exclusive options.

| [-policy <text>] - Efficiency Policy Name }

This option is used to set an efficiency policy. The policy cannot be changed to the predefined *inline-only* and *auto* policy when there is an active background operation on the volume. The policies *inline-only* and *none* are not supported on Capacity optimized Flash with QAT supported platforms.

Note that schedule and policy are mutually exclusive options.

{ [-compression {true|false}] - Compression

This option is used to enable and disable compression. The default value is determined based on the platform.

[-inline-compression {true|false}] - Inline Compression

This option is used to enable and disable inline compression. Inline compression can be enabled only if `compression` is enabled. The default value is determined based on the platform.

You can use the *inline-only* predefined efficiency policy to run inline compression without the need of any background efficiency operations.

[-compression-type {none|secondary|adaptive}] - Compression Type (privilege: advanced)

This option is used to specify the size of compression group on the volume. The default value is determined based on the platform.

[-storage-efficiency-mode {default|efficient}] - Storage Efficiency Mode }

This option is used to modify the storage efficiency mode. The default mode sets 8k adaptive compression on the volume. The efficient mode sets auto adaptive compression and attempts to enable inline deduplication, cross volume inline deduplication and cross volume background deduplication on the volume. There is no difference between default and efficient modes on QAT supported platforms and auto adaptive compression is set irrespective of the modes.

The available efficiency modes are:

- default
- efficient

[-inline-dedupe {true|false}] - Inline Dedupe

This option is used to enable and disable inline deduplication. The default value is determined based on the platform.

You can use the *inline-only* predefined efficiency policy to run inline deduplication without the need of any background efficiency operations.

[-data-compaction {true|false}] - Data Compaction

This option is used to enable and disable data compaction. The default value is determined based on the platform.

[-cross-volume-inline-dedupe {true|false}] - Cross Volume Inline Deduplication

This option is used to enable and disable cross volume inline deduplication. The default value is determined based on the platform.

[-cross-volume-background-dedupe {true|false}] - Cross Volume Background Deduplication

This option is used to enable and disable cross volume background deduplication. The default value is determined based on the platform.

Examples

The following examples modify efficiency options on a volume.

```
cluster1::> volume efficiency modify -vserver vs1 -volume voll -schedule
sun-sat@12
```

```
cluster1::> volume efficiency modify -vserver vs1 -volume voll -policy
policy1
```

```
cluster1::> volume efficiency modify -vserver vs1 -volume voll
-compression true -inline-compression true -inline-dedupe true -data
-compaction true -cross-volume-inline-dedupe true -cross-volume-background
-dedupe true
```

volume efficiency off

Disables efficiency on a volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume efficiency off` command disables efficiency on a volume. Disabling efficiency is not supported on Capacity optimized Flash with QAT supported platforms.

Parameters

-vserver <vserver name> - Vserver Name

Specifies the Vserver on which the volume is located.

{ -volume <volume name> - Volume Name

Specifies the name of the volume on which efficiency needs to be disabled.

| -path </vol/volume> - Volume Path }

Specifies the volume path on which efficiency needs to be disabled.

Examples

The following examples disable efficiency on a volume:

```
cluster1::> volume efficiency off -vserver vs1 -volume voll
```

```
cluster1::> volume efficiency off -vserver vs1 -path /vol/voll
```

volume efficiency on

Enable efficiency on a volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume efficiency on` command enables efficiency on a volume. The specified volume must be online. Efficiency operations will be started periodically according to a per volume schedule or policy. The [volume efficiency modify](#) command can be used to modify schedule and the [volume efficiency policy modify](#) command can be used to modify policy. You can also manually start an efficiency operation with the [volume efficiency start](#) command.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver on which the volume is located.

{ -volume <volume name> - Volume Name

This specifies the name of the volume on which efficiency needs to be enabled.

| -path </vol/volume> - Volume Path }

This specifies the volume path on which efficiency needs to be enabled.

Examples

The following examples enable efficiency on a volume.

```
cluster1::> volume efficiency on -vserver vs1 -volume vol1
```

```
cluster1::> volume efficiency on -vserver vs1 -path /vol/vol1
```

Related Links

- [volume efficiency modify](#)
- [volume efficiency policy modify](#)
- [volume efficiency start](#)

volume efficiency prepare-to-downgrade

Identify any incompatible volumes or Snapshot copies before downgrade

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume efficiency prepare-to-downgrade` command updates efficiency configurations and

metadata to be compatible with prior releases as specified. This command also disables the use of incompatible efficiency features. This command is not supported on FlexGroups.

Parameters

[-disable-feature-set <downgrade version>] - Data ONTAP Version (privilege: advanced)

This parameter specifies the ONTAP version that introduced new volume efficiency feature set.

Examples

The following example disables the features introduced in Data ONTAP 8.3.1

```
cluster1::*> volume efficiency prepare-to-downgrade -disable-feature-set
8.3.1
```

The following example disables the features introduced in ONTAP 9.6.0.

```
cluster1::*> volume efficiency prepare-to-downgrade -disable-feature-set
9.6.0
```

The following example ignores offline volumes while disabling the features introduced in ONTAP 9.6.0 .

```
cluster1::*> volume efficiency prepare-to-downgrade -disable-feature-set
9.6.0 -skip-offline-volumes true
```

The following example ignores offline volumes while disabling the features introduced in Data ONTAP 8.3.1 .

```
cluster1::*> volume efficiency prepare-to-downgrade -disable-feature-set
8.3.1 -skip-offline-volumes true
```

volume efficiency promote

Add a volume to the preferred set of volumes for efficiency processing

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

Use the `volume efficiency promote` command to promote a volume from deprioritized state back to auto state.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

This specifies the Vserver on which the volume is located.

{ -volume <volume name> - Volume Name (privilege: advanced)

This specifies the name of the volume on which auto scheduling needs to be restarted.

| -path </vol/volume> - Volume Path (privilege: advanced) }

This specifies the volume path on which auto scheduling needs to be restarted.

Examples

The following example promotes a volume from deprioritized state back to auto state.

```
cluster1::*> volume efficiency promote -vserver vs1 -volume voll
```

volume efficiency revert-to

Reverts volume efficiency metadata

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume efficiency revert-to` command reverts the format of volume efficiency metadata for the volume to the given version of ONTAP. This command is not supported on FlexGroups.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

This specifies the Vserver on which the volume is located.

{ -volume <volume name> - Volume Name (privilege: advanced)

This specifies the name of the volume for which volume efficiency metadata needs to be reverted.

| -path </vol/volume> - Volume Path (privilege: advanced) }

This specifies the volume path for which volume efficiency metadata needs to be reverted.

[-version <revert version>] - Revert to Version (privilege: advanced)

Specifies the version of ONTAP to which the volume efficiency metadata needs to be formatted.

[-d, -delete {true|false}] - Delete Existing Metafile on Revert (privilege: advanced)

If set to `true`, this parameter specifies that the volume efficiency metadata be deleted instead of reverting its format. By default this parameter is set to `false`.

[-c, -clean-up {true|false}] - Delete Previously Downgraded Metafiles (privilege: advanced)

If set to `true`, this parameter specifies that the volume efficiency metadata already reverted using `volume efficiency revert-to` be deleted. By default this parameter is set to `false`.

[-revert-adaptive-compression {true|false}] - Downgrade to minor version (privilege: advanced)

If set to `true`, this parameter specifies that the volume efficiency metadata needs to be reverted to a minor

version of ONTAP. By default this parameter is set to `false`.

[`-check-snapshot {true|false}`] - Revert ignore snapshots (privilege: advanced)

If set to `false`, this parameter specifies that the volume efficiency revert will not check for Snapshot copies created by previous releases of ONTAP. By default this parameter is set to `true`.

Examples

The following examples reverts volume efficiency metadata on a volume named `vol1` located in vserver `vs1` to version 9.8.

```
cluster1::*> volume efficiency revert-to -vserver vs1 -volume vol1
-version 9.8
```

```
cluster1::*> volume efficiency revert-to -vserver vs1 -path /vol/vol1
-version 9.8
```

volume efficiency show

Display a list of volumes with efficiency

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume efficiency show` command displays the information about storage efficiency of volumes. The command output depends on the parameter or parameters specified. If no parameters are specified, the command displays the following information for all volumes with efficiency:

- **Vserver:** Vserver the volume belongs to.
- **Volume:** Name of the volume.
- **State:** Current state of efficiency on the volume (Enabled, Disabled, or Mixed).
- **Status:** Status of the efficiency on the volume. Following are the possible values:
 - *Active*: An efficiency operation is currently running.
 - *Idle*: There are no efficiency operations running.
 - *Initializing*: An efficiency operation is being initialized.
 - *Undoing*: Efficiency is being undone on the volume.
 - *Pending*: An efficiency operation is queued.
 - *Downgrading*: An efficiency operation necessary to downgrade the efficiency metafiles to a previous ONTAP release is active.
 - *Disabled*: Efficiency is disabled on the volume.
- **Progress:** The progress of the current efficiency operation with information as to which stage of the efficiency process is currently in progress and how much data is processed for that stage. For example:

"25 MB Scanned", "20 MB Searched", "500 KB (2%) Compressed", "40 MB (20%) Done", "30 MB Verified".

To display detailed information, run the command with the `-l`` or `-instance` parameter. The detailed view provides all information in the previous list and the following additional information:

- Path: Volume Path.
- Compression: Current state of compression on the volume (Enabled or Disabled).
- Inline Compression: Current state of inline compression on the volume (Enabled or Disabled).
- Type: Type of volume (Regular or SnapVault).
- Schedule: The schedule of efficiency operation for the volume.
- Policy: Efficiency policy for the volume.
- Minimum Blocks Shared: The minimum number of adjacent blocks in a file that can be shared.
- Blocks Skipped Sharing: Blocks skipped sharing because of the minimum block share value.
- Last Operation State: Status of the last operation (Success or Failure).
- Last Successful Operation Begin: The time and date at which the last successful operation began.
- Last Successful Operation End: The time and date at which the last successful operation ended.
- Last Operation Begin: The time and date at which the last operation began.
- Last Operation End: The time and date at which the last operation ended.
- Last Operation Size: The size of the last operation.
- Last Operation Error: The error encountered by the last operation.
- Change Log Usage: The percentage of the change log that is used.
- Logical Data: The total logical data in the volume, and how much is reached compared to the deduplication logical data limit.
- Queued Job: The job that is queued. Following are the possible values:
 - `-` : There are no queued jobs.
 - `scan` : A job to process existing data is queued.
 - `start` : A job to process newly added data is queued.
 - `check` : A job to eliminate stale data from the fingerprint database is queued.
 - `downgrading` : An efficiency operation necessary to downgrade the efficiency metafiles to a previous ONTAP release is queued.
- Stale Fingerprints: The percentage of stale entries in the fingerprint database. If this is greater than 20 percent a subsequent [volume efficiency start](#) operation triggers the verify operation, which might take a long time to complete.
- Inline Dedupe: Current state of inline deduplication on the volume (Enabled or Disabled).
- Cross Volume Inline Deduplication: Current state of cross volume inline deduplication on the volume (Enabled or Disabled).
- Cross Volume Background Deduplication: Current state of cross volume background deduplication on the volume (Enabled or Disabled).
- Extended Compressed Data: Is there extended compressed data present on the volume.
- Inline Adaptive Data Compaction: Whether Inline Adaptive Data Compaction is enabled or disabled on the volume. When enabled, ONTAP combines data fragments to reduce on-disk block consumption.

You can specify additional parameters to display information that matches only those parameters. For example, to display information only about volumes with efficiency in Vserver vs1, run the command with the `-vserver vs1` parameter.

Parameters

{ [-fields <fieldname>,...]

This specifies the fields that need to be displayed. The fields Vserver and volume name are the default fields.

| [-1]

This option displays detailed information about the volumes with efficiency.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Name

Displays information only for those volumes that match the specified Vserver.

{ [-volume <volume name>] - Volume Name

Displays information only for those volumes that match the specified volume.

| [-path </vol/volume>] - Volume Path }

Displays information only for those volumes that match the specified volume path.

[-state {Disabled|Enabled|Mixed}] - State

Displays information only for those volumes that match the specified state.

[-op-status <Efficiency status>] - Status

Displays information only for those volumes that match the specified operation status.

[-progress <text>] - Progress

Displays information only for those volumes that match the specified progress.

[-type {Regular|SnapVault}] - Type

Displays information only for those volumes that match the specified type of volume.

[-schedule <text>] - Schedule

Displays information only for those volumes that match the specified schedule.

[-policy <text>] - Efficiency Policy Name

Displays information only for those volumes that match the specified policy.

[-blks-skipped-sharing <integer>] - Blocks Skipped Sharing

Displays information only for those volumes that match the specified blocks skipped sharing.

[-last-op-state <text>] - Last Operation State

Displays information only for those volumes that match the specified last operation state.

[-last-success-op-begin <Date>] - Last Success Operation Begin

Displays information only for those volumes that match the specified last successful operation begin time.

[-last-success-op-end <Date>] - Last Success Operation End

Displays information only for those volumes that match the specified last successful operation end time.

[-last-op-begin <Date>] - Last Operation Begin

Displays information only for those volumes that match the specified last operation begin time.

[-last-op-end <Date>] - Last Operation End

Displays information only for those volumes that match the specified last operation end time.

[-last-op-size {<integer>[KB|MB|GB|TB|PB]}] - Last Operation Size

Displays information only for those volumes that match the specified last operation size.

[-last-op-error <text>] - Last Operation Error

Displays information only for those volumes that match the specified last operation error.

[-changelog-usage <percent_no_limit>] - Changelog Usage

Displays information only for those volumes that match the specified change log usage.

[-logical-data-size {<integer>[KB|MB|GB|TB|PB]}] - Logical Data Size

Displays information only for those volumes that match the specified logical data size.

[-logical-data-limit {<integer>[KB|MB|GB|TB|PB]}] - Logical Data Limit

Displays information only for those volumes that match the specified logical data limit.

[-logical-data-percent <percent_no_limit>] - Logical Data Percent

Displays information only for those volumes that match the specified logical data percentage.

[-queued-job <text>] - Queued Job

Displays information only for those volumes that match the specified number of queued jobs.

[-stale-fingerprint-percentage <integer>] - Stale Fingerprint Percentage

Displays information only for those volumes that match the specified stale fingerprint percentage.

[-compression {true|false}] - Compression

Displays information only for those volumes that match the specified compression setting.

[-inline-compression {true|false}] - Inline Compression

Displays information only for those volumes that match the specified inline compression setting.

[-compression-type {none|secondary|adaptive}] - Compression Type (privilege: advanced)

Displays information about the type of compression on the volume[adaptive or secondary].

[-storage-efficiency-mode {default|efficient}] - Storage Efficiency Mode

Displays information only for those volumes that match the specified storage efficiency mode. The default mode sets 8k adaptive compression on the volume. The efficient mode sets auto adaptive compression on the volume.

The available efficiency modes are:

- default
- efficient

[`-is-constituent {true|false}`] - Constituent Volume

Displays information only for those volumes that either are or are not constituents of a FlexGroup, depending on the value provided.

[`-inline-dedupe {true|false}`] - Inline Dedupe

Displays information only for those volumes that match the specified inline deduplication setting.

[`-data-compaction {true|false}`] - Data Compaction

Displays information only for those volumes that match the specified data compaction setting.

[`-cross-volume-inline-dedupe {true|false}`] - Cross Volume Inline Deduplication

Displays information only for those volumes that match the specified cross volume inline deduplication setting.

[`-cross-volume-background-dedupe {true|false}`] - Cross Volume Background Deduplication

Displays information only for those volumes that match the specified cross volume background deduplication setting.

[`-extended-compressed-data {true|false}`] - Extended Compressed Data

Displays information only for those volumes that match the specified extended compressed data value. Extended compressed data is enabled on a volume when both adaptive compression configured with application IO size 8K and data compaction are enabled. Once enabled, extended compressed data can only be disabled by using the `volume efficiency revert-to` command.

[`-extended-auto-adaptive-compression {true|false}`] - Volume Has Extended Auto Adaptive Compression (privilege: advanced)

Displays information only for those volumes that match the specified extended auto adaptive compression setting.

[`-active-data-extended-auto-adaptive-compression {true|false}`] - Volume Has Active Data Extended Auto Adaptive Compression (privilege: advanced)

Displays information only for those volumes that match the specified active data extended auto adaptive compression setting.

Examples

The following example displays information about all volumes with efficiency on the Vserver named vs1:

```
cluster1::> volume efficiency show -vserver vs1
Vserver      Volume              State   Status   Progress
-----
vs1          vol1                Enabled Idle     Idle for 22:37:53
vs1          vol2                Enabled Idle     Idle for 22:37:53
vs1          vol3                Enabled Idle     Idle for 22:37:49
vs1          vol4                Enabled Idle     Idle for 22:37:53
vs1          vol5                Enabled Idle     Idle for 22:37:53
vs1          volham              Enabled Idle     Idle for 22:37:53
vs1          volham1             Enabled Idle     Idle for 22:37:53
7 entries were displayed.
```

The following example displays detailed information about a volume named vol1 on a Vserver named vs1:

```
cluster1::> volume efficiency show -vserver vs1 -volume voll
      Vserver Name: vs1
      Volume Name: voll
      Volume Path: /vol/voll
      State: Enabled
      Status: Idle
      Progress: Idle for 02:14:28
      Type: Regular
      Schedule: sun-sat@0
Efficiency Policy Name: -
Blocks Skipped Sharing: 0
  Last Operation State: Success
Last Success Operation Begin: Wed Jul 14 03:08:44 2021
  Last Success Operation End: Wed Jul 14 03:08:46 2021
  Last Operation Begin: Wed Jul 14 03:08:44 2021
    Last Operation End: Wed Jul 14 03:08:46 2021
  Last Operation Size: 2.52MB
  Last Operation Error: -
    Changelog Usage: 0%
      Logical Data Size: 20.48MB
      Logical Data Limit: 640TB
  Logical Data Percent: 0%
    Queued Job: -
Stale Fingerprint Percentage: 0
  Compression: true
    Inline Compression: true
Storage Efficiency Mode: -
  Constituent Volume: false
    Inline Dedupe: false
    Data Compaction: false
Cross Volume Inline Deduplication: false
ross Volume Background Deduplication: false
  Extended Compressed Data: false
```

Related Links

- [volume efficiency start](#)

volume efficiency start

Starts efficiency operation on a volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

Use the `volume efficiency start` command to start an efficiency operation. The volume must be online and have efficiency enabled. If there is an efficiency operation already active on the volume, this command fails.

When the `volume efficiency start` command is issued, a checkpoint is created at the end of each stage or sub-stage, or on an hourly basis in the gathering phase. If at any point the `volume efficiency start` operation is stopped, the system can restart the efficiency operation from the execution state saved in the checkpoint. The `delete-checkpoint` parameter can be used to delete the existing checkpoint and restart a fresh efficiency operation. The checkpoint corresponding to gathering has a validity period of 24 hours. If the user knows that significant changes have not been made on the volume, then such a gathering checkpoint whose validity has expired can be used with the help of the `use-checkpoint` parameter. There is no time restriction for checkpoints of other stages.

When the volume is configured to use the `inline-only` efficiency policy, the system will stop monitoring changes to the data for the purpose of running background efficiency operations. The background deduplication operations will be disabled. The user can still execute compression specific efficiency operation with `-scan-old-data` and `-compression` parameters to compress the existing data on the volume.

Parameters

-vserver <vserver name> - Vserver Name

Specifies the Vserver on which the volume is located.

{ -volume <volume name> - Volume Name

Specifies the name of the volume.

| -path </vol/volume> - Volume Path }

Specifies the complete path of the volume.

[-s, -scan-old-data <true>] - Scan Old Data

This option scans the file system and processes all existing data. It prompts for user confirmation before proceeding. Use the force option to suppress this confirmation.

{ [-p, -use-checkpoint <true>] - Use Checkpoint (if scanning old data)

Use the checkpoint when scanning existing data. Valid only if `scan-old-data` parameter is true.

| [-d, -delete-checkpoint <true>] - Delete Checkpoint }

Deletes the existing checkpoint and restarts a new `volume efficiency start` operation.

[-qos-policy <sis_qos>] - QoS Policy

Specifies the qos-policy, which indicates how the efficiency operations are throttled. This option can be configured to be `background` or `best-effort`. Default value is `best-effort`. If `background` is specified, the efficiency operations are run with minimum or no impact on the data serving client operations. If `best-effort` is specified, the efficiency operations might have some impact on the data serving client operations.

[-C, -compression <true>] - Start Compression (if scanning old data) (privilege: advanced)

Compresses existing data. Deduplication is not run unless the dedupe option is also specified. Valid only if `scan-old-data` parameter is true.

[-D, -dedupe <true>] - Start Deduplication (if scanning old data) (privilege: advanced)

Deduplicates existing data on disk. Similarly, compression is not run unless the compression option is also specified. Valid only if `scan-old-data` parameter is true.

[-P, -compaction <true>] - Start Compaction (if scanning old data) (privilege: advanced)

Compacts existing data on disk. Valid only if `scan-old-data` parameter is true.

[-m, -build-metadata <true>] - Build metadata without sharing(if scanning old data)

Builds deduplication metadata by scanning the entire file system. You will not achieve any space savings with this option. Once the metadata is built, existing data can be shared with newly written data on subsequent deduplication runs.

[-o, -scan-all <true>] - Scan all the data without shared block optimization(if scanning old data)

Scans the entire file system and processes the shared blocks also. You may be able to achieve additional space savings using this option. Where as, by default the option `-scan-old-data` saves some time by skipping the shared blocks.

[-a, -shared-blocks <true>] - Compress Shared Blocks (if scanning old data) (privilege: advanced)

Compresses the Compression Groups that have shared blocks created by deduplication or cloning data. Valid only if `scan-old-data` parameter is true.

[-b, -snapshot-blocks <true>] - Compress Blocks In Snapshots (if scanning old data) (privilege: advanced)

Compresses data blocks locked in a Snapshot copy. Valid only if `scan-old-data` parameter is true.

[-q, -queue <true>] - Operation Should Be Queued

Queues an efficiency operation. It will be queued only if an operation is already in progress. Valid only if `scan-old-data` is false.

[-f, -force <true>] - Force Operation

Suppresses all confirmation messages.

[-z, -skip-zero-replacement <true>] - Skip Zero block detection and replacement (privilege: advanced)

Skip the zero block detection and replacement during the gatherer scan. Valid only if `scan-old-data` parameter is true.

Examples

The following examples start efficiency on a volume:

```
cluster1::> volume efficiency start -volume vol1 -vserver vs1
```

```
cluster1::> volume efficiency start -scan-old-data -volume vol1 -vserver vs1
```

```
cluster1::> volume efficiency start -volume voll -vserver vs1 -queue
-delete-checkpoint
```

volume efficiency stat

Show volume efficiency statistics

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume efficiency stat` command displays efficiency statistics. The output depends on the parameters specified with the command. If no parameters are specified, the command displays the following efficiency statistics fields for all the volumes:

- Vserver: The Vserver that the volume belongs to.
- Volume Name: Name of the volume.
- Inline Compression Attempts: Number of inline compression attempts done.
- Inline Incompressible CGs: Number of compression groups that cannot be compressed by inline compression.

To display detailed information, run the command with `-instance` parameter.

Parameters

{ [-fields <fieldname>,...]

This specifies the fields that need to be displayed. The Vserver and volume name are the default fields.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Name

Displays statistics only for those volume(s) that match the specified Vserver.

{ [-volume <volume name>] - Volume Name

Displays statistics only for those volume(s) that match the specified volume name.

| [-path </vol/volume>] - Volume Path }

Displays statistics only for those volume(s) that match the specified volume path.

[-b <>true>] - Display In Blocks

Displays usage size in 4k block counts.

[-num-compressed-inline <integer>] - Inline Compression Attempts

Displays statistics only for those volume(s) that match the specified number of Compression Groups attempted inline.

[*-inactive-blocks-compressed* <integer>] - Number of Inactive Blocks Compressed

Displays statistics only for those volumes that match the specified number of cold blocks on which compression was done.

Examples

The following example displays default efficiency statistics for all the volumes.

```
cluster1::> volume efficiency stat
Vserver:                vs1
Volume:                 vol2
Inline Compression Attempts: 0
Inline Incompressible CGs: 0

Vserver:                vs1
Volume:                 vol3
Inline Compression Attempts: 0
Inline Incompressible CGs: 0
```

volume efficiency stop

Stop efficiency operation on a volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

Use the `volume efficiency stop` command to stop an efficiency operation.

Parameters

`-vserver <vserver name>` - Vserver Name

This specifies the Vserver on which the volume is located.

`{ -volume <volume name>` - Volume Name

This specifies the name of the volume on which efficiency operation needs to be stopped.

`| -path </vol/volume>` - Volume Path }

This specifies the volume path on which efficiency operation needs to be stopped.

`[-a, -all <true>]` - Stop All Operations

This specifies both active and queued efficiency operations to be aborted.

Examples

The following examples stop efficiency on a volume.

```
cluster1::> volume efficiency stop -vserver vs1 -volume vol1
```

```
cluster1::> volume efficiency stop -vserver vs1 -volume vol1 -all
```

volume efficiency undo

Undo efficiency on a volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The command `volume efficiency undo` removes volume efficiency on a volume by undoing compression, undoing compaction and removing all the block sharing relationships, and cleaning up any volume efficiency specific data structures. Any efficiency operations on the volume must be disabled before issuing this command. The volume efficiency configuration is deleted when the undo process completes. The command is used to revert a volume to an earlier version of ONTAP where some of the efficiency features are not supported. During this revert not all efficiencies needs to be undone but only those gained by that particular feature (for example, compaction), which is not supported in the earlier version.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

This specifies the Vserver on which the volume is located.

{ -volume <volume name> - Volume Name (privilege: advanced)

This specifies the volume name.

| -path </vol/volume> - Volume Path (privilege: advanced) }

This specifies the volume path.

[-C, -compression <>true>] - Decompress Data in the Volume (privilege: advanced)

Undo the effects of compression. This requires efficiency to be disabled (by performing [volume efficiency off](#)).

[-D, -dedupe <>true>] - Undo Block Sharing in the Volume (privilege: advanced)

Undo the effects of deduplication. This requires efficiency to be disabled (by performing [volume efficiency off](#)).

[-i, -inode <integer>] - Inode Number to Undo Sharing (privilege: advanced)

Remove the block sharings from a specified inode.

[-t, -undo-type {all|wrong}] - Selective Undo (privilege: advanced)

This specifies to remove either all or only invalid block sharing. When *all* is used, all block sharings are removed. When *wrong* is used, only invalid sharings present in the volume are removed. When used along with *log* option, it logs information about all or wrong block sharings without sharing removal.

[-d, -log <true>] - Only Log Incorrect Savings (privilege: advanced)

If specified, information about invalid block sharing relationships will only be logged. Invalid sharings will not be removed. This parameter is only valid when the parameter `-undo-type` is specified as *wrong*.

[-P, -data-compaction <true>] - Undo Data Compaction in the Volume (privilege: advanced)

Undo the effects of data compaction.

[-A, -cross-volume-dedupe <true>] - Undo Cross Volume Deduplication (privilege: advanced)

Undo the effects of cross volume deduplication.

[-X, -extended-compression <true>] - Extended Compression (privilege: advanced)

Undo the effects of extended compression. This removes the compression savings for data that requires more resources to compress.

[-c, -auto-adaptive-compression <true>] - Auto Adaptive Compression (privilege: advanced)

Undo the effects of auto adaptive compression.

[-Z, -undo-compression-algorithm-list {lzopro|zstd|zlib}] - Undo Compression Algorithm List (privilege: advanced)

Undo the effects of the specified compression algorithms.

[-s, -extended-auto-adaptive-compression <true>] - Undo Extended Auto Adaptive Compression (privilege: advanced)

Undo the effects of extended auto adaptive compression.

[-a, -active-data-extended-auto-adaptive-compression <true>] - Undo Active Data Extended Auto Adaptive Compression (privilege: advanced)

Undo the effects of active data extended auto adaptive compression.

Examples

The following are examples of how to use efficiency undo.

To undo deduplication and compression savings, but not compaction savings in a volume name vol1 on a Vserver named vs1:

```
cluster1::*> volume efficiency undo -vserver vs1 -volume vol1
```

To rewrite compressed blocks and undo compression savings in a volume name vol1 on a Vserver named vs1:

```
cluster1::*> volume efficiency undo -vserver vs1 -volume vol1 -compression
```

To rewrite compressed and deduped blocks without any efficiency in a volume name vol1 on a Vserver named vs1:

```
cluster1::*> volume efficiency undo -vserver vs1 -volume vol1 -dedup
-compression
```

To rewrite compacted blocks in a volume name vol1 on a Vserver named vs1:

```
cluster1::*> volume efficiency undo -vserver vs1 -volume vol1 -data
-compaction
```

Related Links

- [volume efficiency off](#)

volume efficiency inactive-data-compression modify

Modify volume inactive data compression configuration of a volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume efficiency inactive-data-compression modify` command is used to modify the state of inactive data compression on a volume.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

This specifies the Vserver of the volume.

-volume <volume name> - Volume Name (privilege: advanced)

This specifies the volume on which the inactive data compression state should be modified.

[-progress <text>] - Progress (privilege: advanced)

Progress of the scan.

[-status <text>] - Status (privilege: advanced)

Operation on the scan succeeded or failed.

[-failure-reason <text>] - Failure Reason (privilege: advanced)

If scan failed, the reason of the failure.

[-total-blocks <integer>] - Total Blocks to be Processed (privilege: advanced)

Total number of blocks which must be processed by inactive data compression scanner.

[-total-processed <integer>] - Total Blocks Processed (privilege: advanced)

Total number of blocks which are already processed by inactive data compression scanner.

[-percentage <percent>] - Progress (privilege: advanced)

Scan progress in percentage.

[-is-enabled {true|false}] - State of Inactive Data Compression on the Volume (privilege: advanced)

Enable or disable inactive data compression on the volume. Disabling inactive data compression is not allowed on Capacity optimized Flash with QAT supported platforms.

[-threshold-days <integer>] - Inactive data compression scan threshold days value (privilege: advanced)

Threshold days value for inactive data compression scan. Valid values for this option range from `-threshold-days-min` to `-threshold-days-max`. This field is not supported on QAT supported platforms.

[-threshold-days-min <integer>] - Inactive data compression scan threshold minimum allowed value. (privilege: advanced)

Minimum allowed value in threshold days for inactive data compression scan. Valid values for this option range from `1` to `-threshold-days-max`. This field is not supported on QAT supported platforms.

[-threshold-days-max <integer>] - Inactive data compression scan threshold maximum allowed value. (privilege: advanced)

Maximum allowed value in threshold days for inactive data compression scan. Valid values for this option range from `-threshold-days-min` to `60`. This field is not supported on QAT supported platforms.

[-read-history-window-size <integer>] - Time window(in days) for which client reads data is collected for tuning. (privilege: advanced)

Client read history window. Valid values for this option range from `1` to `60`. This field is not supported on QAT supported platforms.

[-tuning-enabled {true|false}] - State of auto-tuning of Inactive data compression scan on volume. (privilege: advanced)

Auto-tuning state of inactive data compression scan. This field is not supported on QAT supported platforms.

[-compression-algorithm {lzopro|zstd|zlib}] - Inactive data compression algorithm (privilege: advanced)

Inactive data compression algorithm. The only supported algorithm is ZLIB on QAT supported platforms.

Examples

The following example displays information for modifying the state of inactive data compression scan on volume "vol1":

```
cluster:::> volume efficiency inactive-data-compression modify -vserver
vs1 -volume vol1 -is-enabled true
Inactive data compression scan enabled on volume vol1 in Vserver vs1.
```

volume efficiency inactive-data-compression show

Display volume inactive data compression progress

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume efficiency inactive-data-compression show` command is used to show details of inactive data compression on a volume.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Name (privilege: advanced)

This specifies the vserver of the volume.

[-volume <volume name>] - Volume Name (privilege: advanced)

This specifies the volume on which inactive data compression should be started.

[-progress <text>] - Progress (privilege: advanced)

Progress of the scan.

[-status <text>] - Status (privilege: advanced)

Operation on the scan succeeded or failed.

[-failure-reason <text>] - Failure Reason (privilege: advanced)

If scan failed, the reason of the failure.

[-total-blocks <integer>] - Total Blocks to be Processed (privilege: advanced)

Total number of blocks which must be processed by inactive data compression scanner.

[-total-processed <integer>] - Total Blocks Processed (privilege: advanced)

Total number of blocks which are already processed by inactive data compression scanner.

[-percentage <percent>] - Progress (privilege: advanced)

Scan progress in percentage.

[-cold-blocks-found <integer>] - Number of Cold Blocks encountered (privilege: advanced)

Number of cold blocks encountered by inactive data compression scanner.

[-compression-done <integer>] - Number of Compression Done Blocks (privilege: advanced)

Number of blocks on which compression was done.

[-vol-overwrites <integer>] - Number of Vol-Overwrites (privilege: advanced)

count of logical blocks overwritten in compressed data extents.

[-is-enabled {true|false}] - State of Inactive Data Compression on the Volume (privilege: advanced)

Is scan enabled on the volume.

[-last-op-start-time <integer>] - Time since Last Inactive Data Compression Scan started (privilege: advanced)

Time since the last inactive data compression scan op started.

[-last-op-end-time <integer>] - Time since Last Inactive Data Compression Scan ended (privilege: advanced)

Time since the last inactive data compression scan op ended.

[-last-succ-op-start-time <integer>] - Time since Last Successful Inactive Data Compression Scan started (privilege: advanced)

Time since the last successful inactive data compression scan op started.

[-last-succ-op-end-time <integer>] - Time since Last Successful Inactive Data Compression Scan ended (privilege: advanced)

Time since the last successful inactive data compression scan op ended.

[-average-run-time <integer>] - Average run time for Inactive Data Compression Scan (privilege: advanced)

Average runtime of inactive data compression.

[-threshold-days <integer>] - Inactive data compression scan threshold days value (privilege: advanced)

Threshold days value for inactive data compression scan. Valid values for this option range from `-threshold-days-min` to `-threshold-days-max`.

[-threshold-days-min <integer>] - Inactive data compression scan threshold minimum allowed value. (privilege: advanced)

Minimum allowed value in threshold days for inactive data compression scan. Valid values for this option range from `1` to `-threshold-days-max`.

[-threshold-days-max <integer>] - Inactive data compression scan threshold maximum allowed value. (privilege: advanced)

Maximum allowed value in threshold days for inactive data compression scan. Valid values for this option range from `-threshold-days-min` to `60`.

[-read-history-window-size <integer>] - Time window(in days) for which client reads data is collected for tuning. (privilege: advanced)

Client read history window. Valid values for this option range from `1` to `60`.

[-tuning-enabled {true|false}] - State of auto-tuning of Inactive data compression scan on volume. (privilege: advanced)

Auto-tuning state of inactive data compression scan.

[-compression-algorithm {lzopro|zstd|zlib}] - Inactive data compression algorithm (privilege: advanced)

Inactive data compression algorithm.

[-scan-mode {default|compute_compression_savings|extended_recompression}] - Mode of Inactive data compression scan (privilege: advanced)

Mode of the Inactive data compression scanner.

[-phase1-l1s-processed <integer>] - Phase1 level1s Processed (privilege: advanced)

Phase1 number of L1s processed by inactive data compression scanner.

[-phase1-lns-skipped <text>, ...] - Phase1 lns skipped (privilege: advanced)

Phase1 number of indirect blocks skipped by inactive data compression scanner.

[-phase2-total-blocks <integer>] - Phase2 Total Blocks to be Processed (privilege: advanced)

Phase2 number of total blocks which must be processed by inactive data compression scanner.

[-phase2-blocks-processed <integer>] - Phase2 Blocks Processed (privilege: advanced)

Phase2 number of blocks processed by inactive data compression scanner.

[-repacked-blocks <integer>] - Number of Repacked Blocks (privilege: advanced)

Number of cold blocks repacked by inactive data compression scanner in extended-recompression mode.

Examples

The following example displays details about the inactive data compression on volume "vol1":

```

cluster:::> volume efficiency inactive-data-compression show -vserver vs1
-volume vol1
Volume: vol1

          Vserver: vs1
          Is Enabled: true

                                     Progress: IDLE

          Status: SUCCESS
          Compression Algorithm: lzopro
          Failure Reason: -

                                     Total Blocks: -
                                     Total Blocks Processed: -
                                     Phase1 Lls Processed: -
                                     Phase1 Lns Skipped: -
                                     Phase2 Total Blocks: -
                                     Phase2 Blocks Processed: -

          Percentage: -
          Number of Cold Blocks Encountered: -
          Number of Repacked Blocks: -
          Number of Compression Done Blocks: -
                                     Number of Vol-Overwrites: -
          Time since Last Inactive Data Compression Scan Started(sec): -
          Time since Last Inactive Data Compression Scan Ended(sec): -
          Time since Last Successful Inactive Data Compression Scan Started(sec): -
          Time since Last Successful Inactive Data Compression Scan Ended(sec): -
          Average runtime of Inactive Data Compression(sec): -
          Tuning Enabled: true
          Threshold: 14
          Threshold Upper Limit: 21
          Threshold Lower Limit: 14
          Client Read history window: 14

```

volume efficiency inactive-data-compression start

Start inactive data compression on a volume.

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume efficiency inactive-data-compression start` command is used to start inactive data compression on a volume.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

This specifies the vservers of the volume.

-volume <volume name> - Volume Name (privilege: advanced)

This specifies the volume on which inactive data compression should be started.

[-s, -inactive-days <integer>] - Compression Days Count (privilege: advanced)

Data older than 'inactive-days' days is compressed.

[-m, -scan-mode {default|compute_compression_savings|extended_recompression}] - scanner mode (privilege: advanced)

This specifies in which mode inactive data compression scanner should be started. Three modes available 'default', 'compute_compression_savings' and 'extended_recompression'. 'default' scanner will start compressing the cold data in volume. 'compute_compression_savings' scanner will calculate the auto adaptive compression savings on the volume. 'extended_recompression' scanner will attempt to re-write existing cold data to reduce internal fragmentation.

Examples

The following example displays information for starting inactive data compression scan on volume "vol1":

```
cluster:::> volume efficiency inactive-data-compression start -vserver vs1
-volume vol1
Inactive data compression scan started on volume vol1 in Vserver vs1.
```

volume efficiency inactive-data-compression stop

Stop inactive data compression on a volume.

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume efficiency inactive-data-compression stop` command is used to stop inactive data compression on a volume.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

This specifies the vserver of the volume.

-volume <volume name> - Volume Name (privilege: advanced)

This specifies the volume on which inactive data compression should be stopped.

Examples

The following example displays information for stopping inactive data compression scan on volume "vol1":

```
cluster:::> volume efficiency inactive-data-compression stop -vserver vs1
-volume vol1
Inactive data compression scan stopped on volume vol1 in Vserver vs1.
```

volume efficiency inactive-data-reallocation modify

Updates volume inactive data reallocation policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume efficiency inactive-data-reallocation modify` command is used to modify the inactive data reallocation policy on a volume.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

This specifies the Vserver of the volume.

-volume <volume name> - Volume Name (privilege: advanced)

This specifies the volume on which the inactive data compression state should be modified.

[-policy {none|default|on-read|all}] - Policy for Data Reallocation (privilege: advanced)

This specifies policy for inactive data reallocation on the volume. Inactive to active data reallocation is not supported on QAT supported platforms.

Following are the possible values:

- *default* : Reallocates data for random client driven data reads on AFS.
- *none* : Disables inactive to active data reallocation.
- *on-read* : Reallocates data for client driven data reads including sequential reads on AFS.

Currently unsupported policy:

- *all* : Reallocates data for all client driven data reads on AFS and snapshots. Also includes highly compressed data which was compressed to one single 4k block.

Examples

The following example displays information for modifying the inactive data reallocation policy on volume "vol1":

```
cluster:::> volume efficiency inactive-data-reallocation modify -vserver  
vs1 -volume vol1 -policy default
```

volume efficiency inactive-data-reallocation show

Display volume inactive data reallocation policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume efficiency inactive-data-reallocation show` command is used to show details of inactive data reallocation policy on a volume.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Name (privilege: advanced)

This specifies the vservers of the volume.

[-volume <volume name>] - Volume Name (privilege: advanced)

This specifies the volume on which inactive data compression should be started.

[-policy {none|default|on-read|all}] - Policy for Data Reallocation (privilege: advanced)

This specifies policy for inactive data reallocation on the volume.

Following are the possible values:

- *default*: Reallocates data for random client driven data reads on AFS.
- *none*: Disables inactive to active data reallocation.
- *on-read*: Reallocates data for client driven data reads including sequential reads on AFS.

Currently unsupported policy:

- *all*: Reallocates data for all client driven data reads on AFS and snapshots. Also includes highly compressed data which was compressed to one single 4k block.

Examples

The following example displays details about the inactive data reallocation on volume "vol1":

```
cluster:::> volume efficiency inactive-data-reallocation show -vserver vs1
-volume vol1
Volume: vol1
Vserver: vs1
Policy: default
```

volume efficiency policy create

Create an efficiency policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume efficiency policy create` creates an efficiency policy.

Parameters

-vserver <vserver name> - Vserver

Specifies the Vserver on which the volume is located.

-policy <text> - Efficiency Policy Name

This specifies the policy name.

[-type {threshold|scheduled}] - Policy Type

This specifies the policy type. The policy type defines when the volume using this policy will start processing a changelog. There are two possible values:

- *threshold* means changelog processing occurs when the changelog reaches a certain percentage.
- *scheduled* means changelog processing will be triggered by time.

The default value is *scheduled*.

[-schedule <text>] - Job Schedule Name

This specifies the job schedule. Use `job schedule` commands to manage job schedules. Only cron job schedules are supported.

[-duration <text>] - Duration

This specifies the duration that an efficiency operation can run (in hours). The possible values are "-" or a number between 1 and 999 inclusive. Default value is "-", which means no duration.

[-start-threshold-percent <percent>] - Threshold Percentage

The percentage at which the changelog will be processed. The percentage is checked on an hourly basis. The default value is 20. Valid only if `-type` parameter is set as *threshold*.

[-qos-policy {background|best_effort}] - QoS Policy

This specifies how the efficiency operations are throttled. This option can be configured to be *background* or *best-effort*. Default value is *best-effort*. If *background* is specified, the efficiency operations are run with minimum or no impact on the data serving client operations. If *best-effort* is specified, the efficiency operations might have some impact on the data serving client operations.

[-enabled {true|false}] - Enabled

This specifies whether the policy is enabled or not. The policy is enabled by default.

[-comment <text>] - Comment

User specified comment.

Examples

The following example creates an efficiency policy.

```
cluster1::> volume efficiency policy create -vserver vs1 -policy policy1
-schedule daily -duration 100
```

volume efficiency policy delete

Delete an efficiency policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume efficiency policy delete` command deletes an efficiency policy. An efficiency policy can be deleted only when it is not associated with any volume. The pre-defined policies *default* and *inline-only* cannot be deleted.

Parameters

-vserver <vserver name> - Vserver

This specifies the Vserver on which the volume is located.

-policy <text> - Efficiency Policy Name

This specifies the policy name.

Examples

The following example deletes an efficiency policy:

```
cluster1::> volume efficiency policy delete -vserver vs1 -policy policy1
```

volume efficiency policy modify

Modify an efficiency policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume efficiency policy modify` command can be used to modify the policy attributes.

The attributes of the *inline-only* predefined policy cannot be modified.

Parameters

-vserver <vserver name> - Vserver

This specifies the Vserver on which the volume is located.

-policy <text> - Efficiency Policy Name

This specifies the policy name.

[-type {threshold|scheduled}] - Policy Type

This specifies the policy type. The policy type defines when the volume using this policy will start processing a changelog. There are two possible values:

- *threshold* means changelog processing occurs when the changelog reaches a certain percentage.
- *scheduled* means changelog processing will be triggered by time.

The default value is *scheduled*.

[-schedule <text>] - Job Schedule Name

This specifies the job schedule. Use [job schedule show](#) to show all the jobs.

[-duration <text>] - Duration

This specifies the duration that an efficiency operation can run in hours. The possible value is between 1 and 999 inclusive.

[-start-threshold-percent <percent>] - Threshold Percentage

The percentage at which the changelog will be processed. The percentage is checked on an hourly basis. The default value is 20. Valid only if `-type` parameter is set as *threshold*.

[-qos-policy {background|best_effort}] - QoS Policy

This specifies how the efficiency operations are throttled. This option can be configured to be *background* or *best-effort*. Default value is *best-effort*. If *background* is specified, the efficiency operations are run with minimum or no impact on the data serving client operations. If *best-effort* is specified, the efficiency operations might have some impact on the data serving client operations.

[-enabled {true|false}] - Enabled

This specifies whether the policy is enabled or not. Default value is true.

[-comment <text>] - Comment

User specified comment.

Examples

The following example modifies efficiency policy.

```
cluster1::> volume efficiency policy modify -policy policy1 -schedule hourly
```

Related Links

- [job schedule show](#)

volume efficiency policy show

Show efficiency policies

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume efficiency policy show` command displays information about efficiency policies. By default, the command displays the following information about all policies:

- **Vserver:** Name of the Vserver that the policy belongs to.
- **Policy Name:** Efficiency policy name.
- **Job Schedule:** Job schedule name.
- **Duration (Hours):** The duration in hours that the efficiency operation can run.
- **Enable:** Whether the policy is enabled or not.
- **Comment:** User specified comment.

You can specify additional parameters to select the displayed information. For example, to display efficiency policies only with duration 5 hours, run the command with the `-duration 5` parameter.

The pre-defined policies *default* and *inline-only* are available when all the nodes in the cluster are running Data ONTAP version 8.3 or later.

The *inline-only* pre-defined policy must be used when the user wants to use the inline compression feature without any regularly scheduled or manually started background storage efficiency operations. When a volume is configured to use the *inline-only* efficiency policy, the system will stop monitoring changes to the data for running the background efficiency operations on the volume. Volumes cannot be configured with the *inline-only* policy if there is a currently active background efficiency operation.

Parameters

{ [-fields <fieldname>,...]

Selects the fields to be displayed. Vserver and policy are the default fields (see example).

| [-instance] }

If this parameter is specified, the command displays information about all entries.

[-vserver <vserver name>] - Vserver

Selects information about the policies that match the specified Vserver.

[-policy <text>] - Efficiency Policy Name

Selects information about the policies that match the specified policy name.

[-type {threshold|scheduled}] - Policy Type

Selects information about the policies that match the specified policy type. There are two possible values - *threshold* and *scheduled*.

[-schedule <text>] - Job Schedule Name

Selects information about the policies that match the specified schedule.

[-duration <text>] - Duration

Selects information about the policies that match the specified duration hours.

[-start-threshold-percent <percent>] - Threshold Percentage

Selects information about the policies that match the specified start-threshold-percent. Valid only if `-type` parameter is set as `threshold`.

[-qos-policy {background|best_effort}] - QoS Policy

Selects information about the policies that match the specified throttling method. The values can be `background` or `best-effort`.

[-enabled {true|false}] - Enabled

Selects information about the policies that have the specified enabled setting.

[-comment <text>] - Comment

Selects information about the policies that match the specified comment.

[-policy-owner {cluster-admin|vserver-admin}] - Owner of the Policy

Selects information about the policies that match the specified owner. The values can be `cluster-admin` or `vserver-admin`.

Examples

The following example shows all the efficiency policies with the matching Vserver vs1.

```
cluster1::> volume efficiency policy show -vserver vs1
      Policy      Job      Duration
Vserver Name      Schedule (Hours) QoS Policy Enabled Comment
-----
vs1      default    daily    -      best_effort true    Default
policy
vs1      inline-only -      -      -      -      Inline-Only
policy
vs1      policy1     daily    -      best_effort true    user-
defined
3 entries were displayed.
```

The following example shows all the policies with the following fields - Vserver (default), policy (default) and duration.

```
cluster1::> volume efficiency policy show -fields duration
vserver policy      duration
-----
vs1      default      -
vs1      inline-only  -
vs1      policy1     -
3 entries were displayed.
```

volume encryption commands

volume encryption conversion pause

Pause a running volume encryption conversion operation

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume encryption conversion pause` command pauses the running encryption conversion operation on a volume.

Parameters

-vserver <vserver name> - Vserver Name

This parameter specifies the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This parameter specifies the name of the volume being encrypted.

[-ignore-warning {true|false}] - Ignore Warning for Conversion Pause

If this parameter is set, the command ignores the confirmation message.

Examples

volume encryption conversion resume

Resume a paused volume encryption conversion operation

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume encryption conversion resume` command resumes the paused encryption conversion operation on a volume.

Parameters

-vserver <vserver name> - Vserver Name

This parameter specifies the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This parameter specifies the name of the volume being encrypted.

Examples**volume encryption conversion show**

Show status of a volume encryption conversion

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume encryption conversion show` command displays information about volume encryption conversion in the cluster. By default, with no parameters, it only shows volume encryption operations that have failed or are currently running. The command display output depends on the parameters passed. If `-vserver` and `-volume` are specified, the following information is displayed:

- Vserver Name: The Vserver on which the volume is located.
- Volume Name: The volume that is part of a completed or running volume move operation.
- Start Time: The date and time when the volume encryption operation was started.
- Status of Operation: The status of the operation.
- Percentage Completed: The amount of work to encrypt the volume completed thus far in terms of percentage.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Name

This parameter specifies the Vserver on which the volume is located.

[-volume <volume name>] - Volume Name

This parameter specifies the name of the volume being encrypted.

[-start-time <MM/DD/YYYY HH:MM:SS>] - Start Time

If this parameter is specified, the command displays encryption operations that match the specified date and time in the cluster time zone when the volume move operation started.

[-status <text>] - Status

If this parameter is specified, the command displays encryption operations that match the specified status of the encryption operation.

Examples

The following example shows a sample output for this command:

```
cluster1::> volume encryption conversion show

Vserver      Volume      Start Time      Status
-----
vs1          p2          9/18/2017 17:44:36  Phase 2 of 2 (redirect
scan) is in progress.
```

volume encryption conversion start

Start a volume encryption conversion operation

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume encryption conversion start` command converts a non-encrypted volume to encrypted volume.

Parameters

-vserver <vserver name> - Vserver Name

This parameter specifies the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This parameter specifies the name of the volume being encrypted.

[-ignore-warning {true|false}] - Ignore Warning for Conversion Start

If this parameter is set, the command ignores the confirmation message.

Examples

volume encryption rekey pause

Pause a running volume encryption rekey operation

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume encryption rekey pause` command pauses the running encryption rekey operation on a volume.

Parameters

-vserver <vserver name> - Vserver Name

This parameter specifies the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This parameter specifies the name of the volume being encrypted.

[-ignore-warning {true|false}] - Ignore Warning for Rekey Pause

If this parameter is set, the command ignores the confirmation message.

Examples

volume encryption rekey resume

Resume a paused volume encryption rekey operation

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume encryption rekey resume` command resumes the paused encryption rekey operation on a volume.

Parameters

-vserver <vserver name> - Vserver Name

This parameter specifies the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This parameter specifies the name of the volume being encrypted.

Examples

volume encryption rekey show

Show status of a volume encryption rekey

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume encryption rekey show` command displays information about volume encryption rekey in the cluster. By default, with no parameters, it only shows volume encryption rekey operations that have failed or are currently running. The command display output depends on the parameters passed. If `-vserver` and `-volume` are specified, the following information is displayed:

- Vserver Name: The Vserver on which the volume is located.
- Volume Name: The volume that is part of a completed or running volume move operation.
- Start Time: The date and time when the volume encryption operation was started.

- Status of Operation: The status of the operation.
- Percentage Completed: The amount of work to encrypt the volume completed thus far in terms of percentage.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Name

This parameter specifies the Vserver on which the volume is located.

[-volume <volume name>] - Volume Name

This parameter specifies the name of the volume being encrypted.

[-start-time <MM/DD/YYYY HH:MM:SS>] - Start Time

If this parameter is specified, the command displays encryption operations that match the specified date and time in the cluster time zone when the volume move operation started.

[-status <text>] - Status

If this parameter is specified, the command displays encryption operations that match the specified status of the encryption operation.

Examples

The following example shows a sample output for this command:

```
cluster1::> volume encryption rekey show

Vserver      Volume      Start Time      Status
-----
vs1          vol2        9/18/2017 17:51:41  Phase 2 of 2 (redirect
scan) is in progress.
```

volume encryption rekey start

Start a volume encryption rekey operation

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume encryption rekey start` command changes the encryption key of a volume.

Parameters

-vserver <vserver name> - Vserver Name

This parameter specifies the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This parameter specifies the name of the volume being rekeyed.

[-ignore-warning {true|false}] - Ignore Warning for Rekey Start

If this parameter is set, the command ignores the confirmation message.

Examples

volume encryption secure-purge abort

Abort secure deletion of trash in existing volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume encryption secure-purge abort` command aborts the secure purge operation on a volume.

Parameters

-vserver <vserver name> - Vserver (privilege: advanced)

This parameter specifies the Vserver on which the volume is located.

-volume <volume name> - Volume (privilege: advanced)

This parameter specifies the name of the volume being encrypted.

Examples

volume encryption secure-purge show

Show status of secure-purge operation on a volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume encryption secure-purge show` command displays information about volume encryption securepurge operation in the cluster.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

[`-instance`] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[`-vserver <vserver name>`] - Vserver (privilege: advanced)

This parameter specifies the Vserver on which the volume is located.

[`-volume <volume name>`] - Volume (privilege: advanced)

This parameter specifies the name of the volume being secure purged.

[`-status {invalid|initializing|snapshots-deleting|snapshots-deleted|zombies-draining|zombies-drained|batched-free-log-draining|batched-free-log-drained|finishing-trash-purge|finished-trash-purge|reencrypting|aborting|aborted|success|failure|blocks-reclaim-scan-in-progress|blocks-reclaim-scan-completed|updating-blocks-used}`] - Status (privilege: advanced)

This parameter displays the status of the secure purge operation.

Examples

The following example shows a sample output for this command:

```
cluster1::> volume encryption secure-purge show
Vserver      Volume      Secure Purge Phase
-----
vs1          voll        reencrypting
```

volume encryption secure-purge start

Start secure deletion of trash in existing volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume encryption secure-purge start` command performs secure purge of encrypted volume.

Parameters

`-vserver <vserver name>` - Vserver (privilege: advanced)

This parameter specifies the Vserver on which the volume is located.

`-volume <volume name>` - Volume (privilege: advanced)

This parameter specifies the name of the volume being encrypted.

{ `-delete-all-snapshots {true|false}` } - Delete all Snapshot Copies? (privilege: advanced)

This parameter specifies if it is required to delete all the snapshots in the volume.

[`-re-encryption-method` {`volume-move`|`in-place-rekey`}] - Method to Use for Re-keying the Volume (privilege: advanced)

Specifies the re-encryption method of the secure-purge operation. Possible values are:

- `volume-move`
- `in-place-rekey`

The `volume-move` method is faster compared to `in-place-rekey` method. The `volume-move` method requires additional space in the same aggregate. The `volume-move` method will cause unforced MetroCluster switchovers to be blocked while it is running. The `volume-move` operation involves cut-over phase in which client I/O operations are temporary blocked for short duration - less than 45 seconds. There is no cut-over phase for `in-place-rekey` method.

[`-prepare` <`true`>] - Method to Use for Preparing Secure Purge (privilege: advanced) }

This parameter specifies if it is required to prepare secure purge of encrypted volume.

Examples

volume file commands

volume file compact-data

Apply Adaptive Data Compaction to a Snapshot copy of a file

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `volume file compact-data` command applies the Adaptive Data Compaction feature to the Snapshot copy of a file such that partially filled blocks from that file will merge and consume less storage space.

Parameters

`-node` <`nodename`> - Node

This parameter indicates the node name that the AWA instance runs on.

`-vserver` <`vserver name`> - Vserver Name

This specifies the Vserver in which the target file is located.

`-file` </`vol`/`<volume name>`/`<file path>`>> - File Path

This specifies the complete file path. The Snapshot copy name can be specified as part of the path or by specifying the `-snapshot` parameter.

[`-volume` <`volume name`>] - Volume Name

This specifies the volume in which the targeted file is located.

[`-snapshot` <`snapshot name`>] - Snapshot Copy Name

This specifies the Snapshot copy name in which the file will be compacted.

Examples

The following command applies the Adaptive Data Compaction feature to the Snapshot copy *snap1* of the file */file1* in volume *vol1*:

```
cluster1:> volume file compact-data -vserver vs1 -volume vol1 -file
/vol/vol1/file1 -snapshot snap1
```

volume file modify

Manage the association of a QoS policy group with a file

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command adds and removes files from QoS policy groups. QoS policy groups define measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated. A QoS policy group associated with this file can be created, modified, and deleted. You cannot associate a file to a QoS policy group if a LUN was created from the file.

Parameters

-vserver <vserver name> - Vserver Managing Volume

This specifies the Vserver on which the volume (containing the file) resides.

-volume <volume name> - Volume Name

This specifies the name of the volume. The name must be unique within the hosting Vserver.

-file <text> - File Path

This specifies the actual path of the file with respect to the volume.

{ [-qos-policy-group <text>] - QoS Policy Group Name

This option associates the file with a QoS policy group. This policy group manages storage system resources to deliver your desired level of service. If you do not assign a policy to a file, the system will not monitor and control the traffic to it. To remove this file from a QoS policy group, enter the reserved keyword "none".

| [-qos-adaptive-policy-group <text>] - QoS Adaptive Policy Group Name }

This optional parameter specifies which QoS adaptive policy group to apply to the file. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs) that adjust based on the file's allocated space or used space. To remove this file from an adaptive policy group, enter the reserved keyword "none".

[-caching-policy <text>] - Caching Policy Name

This optionally specifies the caching policy to apply to the file. A caching policy defines how the system caches this volume's data in Flash Cache modules. If a caching policy is not assigned to this file, the system uses the caching policy that is assigned to the containing volume. If a caching policy is not assigned to the containing volume, the system uses the caching policy that is assigned to the containing Vserver. If a caching policy is not assigned to the containing Vserver, the system uses the default cluster-wide policy.

The available caching policies are:

- none - Does not cache any user data or metadata blocks.
- auto - Read caches all metadata and randomly read user data blocks, and write caches all randomly overwritten user data blocks.
- meta - Read caches only metadata blocks.
- random_read - Read caches all metadata and randomly read user data blocks.
- random_read_write - Read caches all metadata, randomly read and randomly written user data blocks.
- all_read - Read caches all metadata, randomly read and sequentially read user data blocks.
- all_read_random_write - Read caches all metadata, randomly read, sequentially read, and randomly written user data.
- all - Read caches all data blocks read and written. It does not do any write caching.

Default caching-policy is auto.

Examples

```
cluster1::> vol file modify -vserver vs0 -volume vs0_vol56 -file 1.txt
-qos-policy-group fast -cache all-read
```

Associates the file *1.txt* with the *fast* QoS policy group and *all-read* caching policy.

volume file privileged-delete

Perform a privileged-delete operation on unexpired WORM files on a SnapLock enterprise volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume file privileged-delete` command is used to perform a privileged-delete operation on unexpired WORM files on a SnapLock enterprise volume. The only built-in role that has access to the command is `"vsadmin-snaplock"`.

Parameters

-vserver <vserver name> - Vserver

Specifies the Vserver which hosts the SnapLock enterprise volume.

-file </vol/<volume name>/<file path>> - File Path

Specifies the absolute path of the file to be deleted. The value begins with `/vol/<volumename>`.

Examples

The following example deletes the unexpired WORM file `"/vol/vol1/wormfile"`. The file `wormfile` is stored in volume `vol1` under Vserver `vserver1`.

```
vserver1::> volume file privileged-delete -file /vol/vol1/wormfile
[Job 76] Job succeeded: Privileged-delete of File
"vs1:/vol/sle_vol1/wormfile" Completed.
```

volume file reservation

Get/Set the space reservation info for the named file.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume file reservation` command can be used to query the space reservation settings for the named file, or to modify those settings. With no further modifiers, the command will report the current setting of the space reservation flag for a file. This tells whether or not space is reserved to fill holes in the file and to overwrite existing portions of the file that are also stored in a snapshot. For symlinks, the link is followed and the command operates on the link target.

Parameters

-vserver <vserver name> - Vserver Name

Specifies the Vserver on which the volume is located. If only one data Vserver exists, you do not need to specify this parameter.

-path </vol/<volume name>/<file path>> - File Name

Specifies the complete file path for which we want to get/set the space reservation settings.

[-is-enabled <text>] - enable | disable

Specifying enable or disable will turn the reservation setting on or off accordingly for the file.

Examples

The following example enables the file reservation setting for the file named file1. The file file1 is stored in volume testvol on Vserver vs0.

```
node::> file reservation -vserver vs0 /vol/testvol/file1 enable
space reservations for file /vol/testvol/file1: on.
```

volume file show-disk-usage

Show disk usage of file

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command requires a path to a file in a volume and displays the following information:

- Vserver name
- Total bytes used by the file in kilobytes
- Full Path to the file

If not logged in as Vserver administrator, the command also requires a Vserver name.



The "-instance" option provides the same result as the default as there are no extra fields to display.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-h]

If this parameter is specified, the command displays total bytes used by the file in human readable form.

| [-k]

If this parameter is specified, the command displays total bytes used by the file in kilobytes.

| [-m]

If this parameter is specified, the command displays total bytes used by the file in megabytes.

| [-u]

If this parameter is specified, the command displays the unique bytes used by the file (bytes that are not shared with any other file in the volume due to deduplication or FlexClone files) in kilobytes.

| [-uh]

If this parameter is specified, the command displays the unique bytes used by the file in human readable form.

| [-uk]

If this parameter is specified, the command displays the unique bytes used by the file in kilobytes.

| [-um]

If this parameter is specified, the command displays the unique bytes used by the file in megabytes.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

-vserver <vserver name> - Vserver

This parameter is used to specify the Vserver that contains the file for which the command displays the total bytes used. It is required if not logged in as Vserver administrator.

-path </vol/<volume name>/<file path>> - Full Path

This required parameter is used to specify the path of the file for which the command displays the total bytes used.

[-r, -range [start offset>:<end offset]] - Block Range

If this parameter is specified, the command displays the total bytes used by the file in the specified block range.

Examples

The following example displays the disk-usage of the file `file1.txt` in volume `/vol/root_vs0`.

```
cluster1::> volume file show-disk-usage -vserver vs0 -path
/vol/root_vs0/file1.txt

Vserver          Total          Path
-----          -
vs0              1408KB        /vol/root_vs0/file1.txt
cluster1::> volume file show-disk-usage -m -vserver vs0 -path
/vol/root_vs0/file1.txt

Vserver          Total          Path
-----          -
vs0              1MB          /vol/root_vs0/file1.txt
vs0::> volume file show-disk-usage -um -path /vol/root_vs0/file1.txt

Vserver          Total          Unique          Path
-----          -
vs0              1MB          1MB          /vol/root_vs0/file1.txt
```

volume file show-filehandle

Show the file handle of a file

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

This command requires a path to a file in a volume and displays the file handle and volume information described below:

- Vserver name
- Path to the file
- File handle flags
- Snapshot ID of the file (snapid)
- File ID
- File handle generation number
- File system ID (fsid)

- Master data set ID (msid)
- Data set ID (dsid)
- Volume or constituent name

If not logged in as a Vserver administrator, the command also requires a Vserver name.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Managing Volume (privilege: advanced)

This specifies the Vserver where the file resides.

[-path <text>] - Path to File (privilege: advanced)

This specifies the path to the file.

Examples

The following example displays the file handle and volume information of a file named `file1.txt` in the volume `/vol/vol1`.

```
cluster1::> volume file show-filehandle -vserver vs0 -path
/vol/vol1/file1.txt
          Vserver                Path
-----
vs0                /vol/vol1/file1.txt
flags  snapid  fileid  generation  fsid      msid      dsid
-----
0x402  0x0     0       0x60       0x206b6   0x402     0x80000402
volume
-----
vol1__0003
```

volume file show-inode

Display file paths for a given inode

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays information about all the files having a given inode in a volume of a Vserver. If the `-snapshot-id` or `-snapshot-name` parameter is specified, the command displays file information from the Snapshot copy; otherwise, it displays the information from the active file system. The `-vserver`, `-volume` and `-inode-number` are mandatory parameters.

If no optional parameter is specified, the command displays the following fields for all the files having the given inode:

- Vserver Name
- Volume Name
- Inode Number
- File Path

The `volume file show-inode` command is only supported on flexible volumes and FlexGroup constituents.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields` parameter, the command output also includes the specified field or fields.

| [-snapshot]

If this parameter is specified, the command displays the following information:

- Vserver Name
- Volume Name
- Inode Number
- Snapshot Name
- Snapshot ID
- File Path

| [-instance] }

If this parameter is specified, the command displays detailed information about the files matching the specified inode number. The following information is displayed:

- Vserver Name
- Volume Name
- Inode Number
- File Path
- Snapshot Name
- Snapshot ID
- File Name

- Parent Inode Number
- Parent Directory Cookie

-vserver <vserver name> - Vserver Name

This specifies the Vserver in which the volume or Snapshot copy is located.

-volume <volume name> - Volume Name

This specifies the volume in which the inode number is located.

-inode-number <integer> - Inode Number

This specifies the inode number whose information has to be retrieved.

{ [-snapshot-name <snapshot name>] - Snapshot Name

If this parameter or `-snapshot-id` is specified, information about the files is retrieved from the Snapshot copy instead of the active file system.

| [-snapshot-id <integer>] - Physical Snapshot ID }

If this parameter or `-snapshot-name` is specified, information about the files is retrieved from the Snapshot copy instead of the active file system.

[-file-path <text>] - File Path

If this parameter is specified, the command displays information only about the files that match the specified file path.

[-file-name <text>] - File Name

If this parameter is specified, the command displays information only about the files that match the specified file name.

[-parent-inode-number <integer>] - Parent Inode Number

The inode number of the parent directory of the file associated with the inode. If this parameter is specified, the command displays information only about the files that match the specified parent inode number.

[-parent-dir-cookie <integer>] - Parent Directory Cookie

The index of the directory entry of the file in its parent directory tree. If this parameter is specified, the command displays information only about the files that match the specified parent directory cookie.

Examples

The following example displays information about all the files having the inode number 96 in the active file system of a volume named vol1 on a Vserver named vs1:

```
cluster1::> volume file show-inode -vserver vs1 -volume vol1 -inode-number 96
```

Vserver	Volume	Inode Number	File Path
vs1	vol1	96	/vol/vol1/file1
vs1	vol1	96	/vol/vol1/file2
vs1	vol1	96	/vol/vol1/A/file2

3 entries were displayed.

The following example displays information about all the files with inode number 96 in a Snapshot copy named mysnap. The Snapshot copy is present in a volume named vol1 on a Vserver named vs1:

```
cluster1::> volume file show-inode -vserver vs1 -volume vol1 -inode-number 96 -snapshot-name mysnap -snapshot
```

Vserver	Volume	Inode Number	Snapshot Name	Snapshot ID	File Path
vs1	vol1	96	mysnap	131	/vol/vol1/.snapshot/mysnap/file1
vs1	vol1	96	mysnap	131	/vol/vol1/.snapshot/mysnap/file2

2 entries were displayed.

The following example displays detailed information about all the files with inode number 96 in a Snapshot copy named mysnap. The Snapshot copy is present in a volume named vol1 on a Vserver named vs1:

```

cluster1::> volume file show-inode -vserver vs1 -volume vol1 -inode-number
96 -snapshot-name mysnap -instance
Vserver Name: vs1
    Volume Name: vol1
    Inode number: 96
    File Path: /vol/vol1/.snapshot/mysnap/file1
    Snapshot Name: mysnap
Physical Snapshot ID: 131
    File Name: file1
    Parent Inode Number: 64
    Parent Directory Cookie: 2
Vserver Name: vs1
    Volume Name: vol1
    Inode number: 96
    File Path: /vol/vol1/.snapshot/mysnap/file2
    Snapshot Name: mysnap
Physical Snapshot ID: 131
    File Name: file2
    Parent Inode Number: 64
    Parent Directory Cookie: 3
    2 entries were displayed.

```

volume file async-delete cancel

Cancel an async directory delete job

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume file async-delete cancel` command cancels an async delete of a directory.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

This specifies the Vserver on which the volume containing the directory resides.

-volume <volume name> - Volume Name (privilege: advanced)

This specifies the name of the volume.

-jobid <text> - Job ID (privilege: advanced)

This specifies the Job ID of the async directory delete job to cancel. Run [volume file async-delete show](#) to get the job ID.

Examples

The following example cancels a job with job ID "0:1" in volume "vol1" in Vserver "vs1".

```
cluster1::*> volume file async-delete cancel -vserver vs1 -volume vol1
-jobid 0:1
```

Related Links

- [volume file async-delete show](#)

volume file async-delete prepare-for-revert

Cancel all async directory delete jobs before starting revert

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume file async-delete prepare_for_revert` command cancels all async directory delete jobs.

Examples

The following example deletes all async directory delete jobs.

```
cluster1::*> volume file async-delete prepare-for-revert
```

volume file async-delete show

List ongoing async directory delete jobs

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The event log will contain events for async directory delete operations that have already completed. Use the `event log show -event async_delete` command to find completed async directory delete operations.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Name (privilege: advanced)

If this parameter is specified, ongoing async directory delete operations that match the specified Vserver are displayed.

[`-volume <volume name>`] - Volume Name (privilege: advanced)

If this parameter is specified, ongoing async directory delete operations that match the specified volume are displayed.

[`-jobids <text>,...`] - Job ID (privilege: advanced)

If this parameter is specified, ongoing async directory delete operations that match the specified job ID are displayed.

[`-paths <text>,...`] - Directory Path (privilege: advanced)

If this parameter is specified, ongoing async directory delete operations that match the specified path are displayed.

[`-deleted-files <integer>,...`] - Number of files deleted (privilege: advanced)

This parameter specifies the number of files deleted as part of the corresponding async directory delete job in a FlexVol.

[`-deleted-dirs <integer>,...`] - Number of directories deleted (privilege: advanced)

This parameter specifies the number of directories deleted as part of the corresponding async directory delete job in a FlexVol.

[`-deleted-bytes <integer>,...`] - Deleted size in bytes (privilege: advanced)

This parameter specifies total number of bytes deleted as part of the corresponding async directory delete job in a FlexVol.

Examples

The following example shows that there are two async directory deletes in progress; one in flexvol vol1 and another one in flexgroup fg.

```
cluster1::*> volume file async-delete show
  Vserver      Volume      Job ID      Path      Progress
              (#File  #Dir
Size (B) )
-----
vs1           vol1        0:2         d2/dd     1250      2
24560
vs2           fg          0:1         d1/d2     -         -
```

volume file async-delete start

Start async delete of a directory

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume file async-delete start` command starts async delete of a directory.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

This specifies the Vserver on which the volume containing the directory resides.

-volume <volume name> - Volume Name (privilege: advanced)

This specifies the name of the volume.

-path <text> - Directory Path (privilege: advanced)

This specifies the path of the directory.

[-throttle <integer>] - Throttle (privilege: advanced)

This specifies maximum number of directory delete operations per second.

Examples

The following example deletes directory "d2" under directory "d1" in volume "vol1" hosted in Vserver "vs1".

```
cluster1::*> volume file async-delete start -vserver vs1 -volume vol1
-path d1/d2
```

volume file async-delete client disable

Disable async delete of a directory from the client

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume file async-delete client disable` command disables async delete of a directory on the volume, from client applications.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

This specifies the Vserver on which the volume resides.

-volume <volume name> - Volume Name (privilege: advanced)

This specifies the name of the volume.

Examples

The following example disables async directory delete from the client in volume "vol1" hosted in Vserver "vs1".


```
cluster1::*> volume file async-delete client disable -vserver vs1 -volume
vol1
```

volume file async-delete client enable

Enable async delete of a directory from the client

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume file async-delete client enable` command enables async delete of a directory on the volume, from client applications. When async directory delete from the client is enabled, 'mv' or 'rename' of a directory, to the trashbin name, triggers async delete of the directory.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

This specifies the Vserver on which the volume resides.

-volume <volume name> - Volume Name (privilege: advanced)

This specifies the name of the volume.

[-trashbin <text>] - Trashbin Directory Name (privilege: advanced)

This specifies the name of the trashbin directory.

Examples

The following example enables async delete of directory in volume "vol1" hosted in Vserver "vs1" from client.

```
cluster1::*> volume file async-delete client enable -vserver vs1 -volume
vol1 -trashbin ntaptrash
```

volume file async-delete client show

Display the status of async delete of a directory from the client

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume file async-delete client show` command shows the configuration of async delete of a directory on a volume from client applications.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Name (privilege: advanced)

This specifies the Vserver on which the volume resides.

[-volume <volume name>] - Volume Name (privilege: advanced)

This specifies the name of the volume.

[-is-enabled {true|false}] - Is Client Async Delete Enabled (privilege: advanced)

This specifies whether async delete of a directory from client applications is enabled on the volume.

[-trashbin <text>] - Trashbin Directory Name (privilege: advanced)

This specifies the name of the trashbin directory.

[-inode <integer>] - Trashbin Directory Inode Number (privilege: advanced)

This specifies the inode number of the trashbin directory.

Examples

The following example shows the enable status of async directory delete from the client for volume "vol1" hosted in Vserver "vs1".

```
cluster1::*> volume file async-delete client show -vserver vs1 -volume
vol1
```

volume file clone autodelete

Enable/Disable autodelete

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume file clone autodelete` command enables or disables the automatic deletion of file, LUN or NVMe namespace clones. Newly created file, LUN and NVMe namespace clones are disabled for automatic deletion by default.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver on which the volume resides. If only one data Vserver exists, you do not need to specify this parameter.

[`-volume <volume name>`] - Volume Name

This specifies the name of the volume in which the file, LUN or NVMe namespace is present.

`-clone-path <text>` - Clone Path

This specifies the path where clone resides. If you use the volume parameter, then specify the relative path to the file, LUN or NVMe namespace clone. Otherwise, specify the absolute path.

`-enable {true|false}` - Enable or Disable Autodelete

This parameter enables or disables the autodelete feature for the file, LUN or NVMe namespace clones in a specified volume if the clones are already added for automatic deletion. If you set the parameter to true, the specified file, LUN or NVMe namespace clones get automatically deleted in the 'try' or 'disrupt' mode. If the value is false, the clones get automatically deleted only in the 'destroy' mode.

[`-force <true>`] - Force Enable or Disable Autodelete

If `-enable` is *true* then this parameter forces automatic deletion of a specified file, LUN or NVMe namespace, or a file, LUN or NVMe namespace clone. If `-enable` is *false* then specifying this parameter disables autodeletion on a file, LUN or NVMe namespace - or a file, LUN or NVMe namespace clone - even if `-commitment`destroy` is specified.

Examples

The following command enables for automatic deletion a LUN Clone named `lun_clone` contained in a volume named `volume1`. This volume is present on a Vserver named `vs1`.

```
cluster1::> volume file clone autodelete /vol/volume1/lun_clone -enable true -vserver vs1
```

The following command specifies the relative clone path when the volume parameter is specified in the command.

```
cluster1::> volume file clone autodelete lun_clone -enable true -vserver vs1 -volume volume1
```

volume file clone create

Create file or LUN full or sub file clone

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume file clone create` command creates a clone of a file, a LUN or an NVMe namespace. You can optionally specify the following parameters for the clone file creation process:

- Vserver in which the volume resides
- Name of the parent snapshot
- The range of blocks to be cloned

- The option to avoid space reservations for the new file or LUN clone
- The option to assign a QoS policy group to the new file or LUN clone
- The option to assign a caching policy to the new file or LUN clone
- The option to mark the new file, LUN or NVMe namespace clone created for auto deletion
- The option to overwrite an existing file, LUN or NVMe namespace clone

File, LUN or NVMe namespace clones create a duplicate copy of another file, LUN or NVMe namespace, but don't require copying the data itself. This allows the clone operation to occur in constant time, taking the same amount of time to complete no matter the size of the file being cloned. This also means that clones require only a small amount of additional storage space because the clone shares the data with the source file, LUN or NVMe namespace.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver in which the parent volume resides. If only one data Vserver exists, you do not need to specify this parameter.

[-volume <volume name>] - Volume

This specifies the name of volume in which a file, LUN or NVMe namespace is going to be cloned.

-source-path <text> - Source Path

This specifies the path to the file, LUN or NVMe namespace to be cloned relative to the specified volume.

-destination-path <text> - Destination Path

This specifies the path for the newly-created cloned file, LUN or NVMe namespace relative to the specified volume. If the file, LUN or NVMe namespace clone to be created is a whole file, LUN or NVMe namespace, the destination file, LUN or NVMe namespace must not exist. If the `range` parameter is specified, the destination file or LUN must exist. If the `snapshot-name` parameter is specified, this option is mandatory.

[-s, -snapshot-name <snapshot name>] - Source Snapshot

The name of the Snapshot copy to use as the source for the clone operation. If this value is not specified, the active filesystem will be used instead.

{ [-r, -range <text>, ...] - Block Range

This specifies the block range to be cloned. If the range is not specified, the entire file, LUN, or NVMe namespace is cloned. The block range should be specified in the format `s:d:n` where `s` is the source start block number, `d` is the destination start block number, and `n` is the length in blocks to be cloned. The range of `n` should be from 1 to 32768 or 1 to 16777216 in case of clone from Active File System or Snapshot copy respectively. If this parameter is used in the path provided by the `destination-path`, the parameter must refer to a file, LUN, or NVMe namespace which already exists. If either the source or destination is a LUN or NVMe namespace, then the block size is measured in LBA blocks. The source object block size and destination block size must be equal. If neither the source nor destination is a LUN or NVMe namespace, then the block size will be 4KB. If 512-byte sectors are used, the source and destination offsets must have the same offset within 4KB blocks. + This option is most likely to be used by external automated systems in managing virtual disk configurations and not by human administrators.

[-o, -no-reserve {true|false}] - Do not reserve clone }

If this option is used, the clone file or LUN will not be guaranteed space in the underlying aggregate. While this out-of-space condition persists, writes to the clone file or LUN would fail. This option may be useful if

few writes to the clone are expected to be needed, or to allow a file or LUN clone to be created under space-constrained conditions for recovery purposes. If this option is not specified the clone will inherit the space reservation properties from the source.

[-i, -ignore-streams {true|false}] - Ignore streams

This parameter specifies whether streams should be ignored during cloning of files, LUNs, or NVMe namespaces. If you set this parameter to false, the streams are ignored; otherwise, they are included in the clones. The default value is false.

[-k, -ignore-locks {true|false}] - Ignore locks

This parameter specifies whether byte-range locks and shared-mode locks on files, LUNs or NVMe namespaces should be ignored during cloning. If you set this parameter to true, the locks are ignored; otherwise, clone operation fails if locks are present on files, LUNs or NVMe namespaces. The default value is false.

[-d, -overwrite-destination {true|false}] - Overwrite Destination

Specify this parameter to overwrite the destination file, LUN, or NVMe namespace, if it exists. The default is not to overwrite the destination file. The command will fail if the destination file exists.

{ [-qos-policy-group <text>] - QoS Policy Group Name

This optionally specifies which QoS policy group to apply to the file or LUN. This policy group defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated. If you do not assign a policy group to a file or LUN, the system will not monitor and control the traffic to it. You cannot associate a file to a QoS policy group if a LUN was created from the file.

| [-qos-adaptive-policy-group <text>] - QoS Adaptive Policy Group Name }

This optionally specifies which QoS adaptive policy group to apply to the file or LUN. This policy group defines measurable service level objectives (SLOs) and Service Level Agreements (SLAs) that adjust based on the file or LUN's allocated space or used space.

[-caching-policy <text>] - Caching Policy Name

This optionally specifies the caching policy to apply to the file. A caching policy defines how the system caches this volume's data in Flash Cache modules. If a caching policy is not assigned to this file, the system uses the caching policy that is assigned to the containing volume. If a caching policy is not assigned to the containing volume, the system uses the caching policy that is assigned to the containing Vserver. If a caching policy is not assigned to the containing Vserver, the system uses the default cluster-wide policy. The available caching policies are:

- none - Does not cache any user data or metadata blocks.
- auto - Read caches all metadata and randomly read user data blocks, and write caches all randomly overwritten user data blocks.
- meta - Read caches only metadata blocks.
- random_read - Read caches all metadata and randomly read user data blocks.
- random_read_write - Read caches all metadata, randomly read and randomly written user data blocks.
- all_read - Read caches all metadata, randomly read and sequentially read user data blocks.
- all_read_random_write - Read caches all metadata, randomly read, sequentially read, and randomly written user data.
- all - Read caches all data blocks read and written. It does not do any write caching.

Default caching-policy is auto.

[`-autodelete {true|false}`] - Mark Clone for Autodeletion

This parameter marks the file, LUN or NVMe namespace clones created for auto deletion. When set to true, the file, LUN or NVMe namespace clones get automatically deleted when the volume runs out of space. The default value is false.

[`-bypass-throttle {true|false}`] - Bypass Throttle Checks (privilege: advanced)

This parameter specifies whether clone throttle checks should be skipped during clone creation. When set to true, clones are created without enforcing any clone throttle checks. The default value is false.

[`-nosplit-entry {true|false}`] - Mark Clone to keep Unsplit

This parameter specifies whether the clone should be kept unsplit. When set to true, clones created will not undergo split routine. It is applicable only for full-file clones created from Active File System. The default value is false.

[`-is-backup {true|false}`] - Is a Clone for Backup

This parameter is used to mark the destination file as a backup clone, where divergence is expected on the source file and no divergence is expected on the destination file. It is applicable only for full-file clones created from Active File System volumes. The default value is *false*.

[`-destination-volume <volume name>`] - Destination Volume

This specifies the name of the volume where the destination file resides. This can be different from `volume`, whereas parameter `volume` specifies the volume on which source file resides. This is an optional argument that applies only to a MetaWAFL volume where the source and destination volumes for the clone operation can be different. If this parameter is not given, the destination file will be created in the volume where `source_file` resides.

Examples

The following command creates a FlexClone file of the file named *myfile* contained in a volume named *vol*. The file *myfile* is located in the root directory of that volume. The cloned file *myfile_copy* resides in the root directory same volume.

```
cluster1::> volume file clone create -volume vol -source-path /myfile
-destination-path /myfile_copy
```

The following command optionally associates the FlexClone file named *myfile_copy* with the *fast* QoS policy group and the caching policy named *random-read*.

```
cluster1::> volume file clone create -volume vol -source-path /myfile
-destination-path /myfile_copy -qos-policy-group fast -caching-policy
random-read
```

volume file clone show-autodelete

Show the autodelete status for a file or LUN clone

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume file clone show-autodelete` command displays the autodelete details of a file, LUN or NVMe namespace clone. The command displays the following information about a file, LUN or NVMe namespace clone:

- Vserver Name
- Clone Path
- Whether auto deletion of file, LUN or NVMe namespace clone is enabled

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

-vserver <vserver name> - Vserver Name

This specifies the Vserver to which the file, LUN or NVMe namespace clone belongs.

-clone-path <text> - Clone Path

This specifies the path of the file, LUN or NVMe namespace clone.

[-autodelete-enabled {true|false}] - Autodelete Enabled

If this parameter is true, the file, LUN or NVMe namespace clone gets automatically deleted in the 'try' or 'disrupt' mode. If the value is false, the clones get automatically deleted only in the 'destroy' mode.

Examples

The following example displays the autodelete information about a file, LUN or NVMe namespace clone.

```
cluster1::> volume file clone show-autodelete -vserver vs1 -clone-path
/vol/v1/f1
Vserver Name: vs1
      Clone Path: /vol/v1/f1
Autodelete Enabled: true
```

volume file clone deletion add-extension

Add new supported file extensions to be deleted with clone delete

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume file clone deletion add-extension` command can be used to add new supported file extensions for clone delete.

Parameters

-vserver <vserver name> - Vserver Name

Name of the vserver.

-volume <volume name> - Volume Name

Name of the volume.

-extensions <text> - Supported Extensions for Clone Delete

List of supported file extensions for clone delete.

Examples

The following example adds the new supported *vmdk*, *vhd* file extensions to volume *vol1* of vserver *vs1*.

```
cluster1::> volume file clone deletion add-extension -vserver vs1 -volume  
vol1 -extensions vmdk,vhd
```

volume file clone deletion modify

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume file clone deletion modify` command can be used to change the required minimum clone file size of a volume for clone delete.

Parameters

-vserver <vserver name> - Vserver Name

Name of the vserver.

-volume <volume name> - Volume Name

Name of the volume.

[-minimum-size {<integer>[KB|MB|GB|TB|PB]}] - Minimum Size Required for Clone delete

Minimum clone file size required for clone delete.

Examples

The following example changes the required minimum file size to 100M for volume *vol1* of vserver *vs1*.

```
cluster1::> volume file clone deletion modify -volume vol1 -vserver vs1  
-minimum-size 100M
```


volume file clone deletion remove-extension

Remove unsupported file extensions for clone delete

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume file clone deletion remove-extension` command can be used to remove the existing file extensions that are no longer supported for clone delete.

Parameters

-vserver <vserver name> - Vserver Name

Name of the vservers.

-volume <volume name> - Volume Name

Name of the volume.

[-extensions <text>] - Unsupported Extensions for Clone Delete

List of unsupported file extensions for clone delete.

Examples

The following example removes the existing unsupported *vmdk*, *vhd* file extensions to volume *vol1* of vservers *vs1*.

```
cluster1::> volume file clone deletion remove-extension -vserver vs1  
-volume vol1 -extensions vmdk,vhd
```

volume file clone deletion show

Show the supported file extensions for clone delete

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume file clone deletion show` command displays the following information for clone delete:

- Vserver Name
- Volume Name
- Minimum File Size Required for Clone Delete
- List of Supported File Extensions for Clone Delete

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Name

Name of the vservers.

[-volume <volume name>] - Volume Name

Name of the volume.

[-extensions <text>,...] - Supported Extensions for Clone Delete

List of supported file extensions for Clone Delete.

[-minimum-size {<integer>[KB|MB|GB|TB|PB] }] - Minimum Size Required for Clone delete

Minimum file size required for Clone Delete.

Examples

The following example displays the clone deletion information for all volumes of all vservers.

```
cluster1::> volume file clone deletion show
      Vserver                Volume                Minimum      Extensions
      -----                -
      -----                -
      vs0                    testvol                100B         vmdk, vhd,
vhd, vswp
      vs0_root                0B                   -
      vs1                    testvol                100G         vmdk, vhd,
vhd, vswp
      vs1_root                0B                   -
```

The following example displays the clone deletion information for volume `vol1` of vservers `vs1`.

```
cluster1::> volume file clone deletion show -vserver vs0 -volume testvol
      Vserver Name: vs0
      Volume Name: testvol
      Supported Extensions for Clone Delete: vmdk, vhd, vhd, vswp
      Minimum Size Required for Clone delete: 100B
```

volume file clone split load modify

Modify maximum split load on a node

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume file clone split load modify` command can be used to change the maximum split load (file, LUN or NVMe namespace clones) of a node.

Parameters

-node {<nodename>|local} - Node Name

Node name on which the new maximum split load is being applied.

[-max-split-load {<integer>[KB|MB|GB|TB|PB]}] - Maximum Clone Split Load

This specifies the new maximum split load of a node. This is the amount of clone create load, the node can take at any point of time. If it crosses this limit, then the clone create requests will not be allowed, till the split load is less than maximum split load

Examples

The following example changes the new maximum limit to 10TB on node1.

```
cluster1::*> volume file clone split load*> modify -node clone-01 -max
-split-load 100KB
```

volume file clone split load show

Show split load on a node

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume file clone split load show` command displays the corresponding file, LUN or NVMe namespace clone split loads on nodes. If no parameters are specified, the command displays the following information:

- Node
- Max Split Load
- Current Split Load
- Token Reserved Load
- Allowable Split Load

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-node {<nodename>|local}] - Node Name

Node on which the file, LUN or NVMe namespace clone split load is displayed.

[-max-split-load {<integer>[KB|MB|GB|TB|PB]}] - Maximum Clone Split Load

This specifies the maximum allowable split load on the node.

[-current-split-load {<integer>[KB|MB|GB|TB|PB]}] - Current Clone Split Load

This specifies the current on going split load on the node.

[-token-reserved-load {<integer>[KB|MB|GB|TB|PB]}] - Load Reserved for Clone Creation

This specifies the reserved split load of the node using the tokens.

[-allowable-split-load {<integer>[KB|MB|GB|TB|PB]}] - Allowable Clone Split Load

This specifies the available split load of the node.

Examples

The following example displays the current and allowable file, LUN or NVMe namespace clone split load on a node.

```
cluster1::> volume file clone split load show
Node
Allowable
          Max          Current      Token
          Split Load Split Load Reserved Load Split
Load
-----
-----
clone-01          15.97TB          0B          100MB
15.97TB
clone-02          15.97TB          0B          100MB
15.97TB
2 entries were displayed.
cluster1::> volume file clone split load show -node clone-01 -instance
Node Name: clone-01
          Maximum Clone Split Load: 15.97TB
          Current Clone Split Load: 0B
Load Reserved for Clone Creation: 100MB
          Allowable Clone Split Load: 15.97TB
```

volume file fingerprint abort

Abort a file fingerprint operation

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume file fingerprint abort` command aborts an in-progress fingerprint operation. This command only aborts the fingerprint operations that have not yet completed. This command takes `session-id` as input and aborts the fingerprint operation that is associated with that particular `session-id`.

Parameters

-session-id <integer> - Session ID of Fingerprint Operation

Specifies the `session-id` of the fingerprint operation that needs to be aborted. It is a unique identifier for the fingerprint operation. This `session-id` is returned when the fingerprint operation is started on a file.

Examples

The following example aborts the fingerprint operation identified by `17039361`:

```
cluster1::> volume file fingerprint abort -session-id 17039361
```

volume file fingerprint dump

Display fingerprint of a file

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume file fingerprint dump` command displays the following information given the `-session-id` of the fingerprint operation:

- **Vserver:**

The Vserver on which the file exists.

* **Session-ID:**

A unique identifier for the fingerprint operation. This `session-id` is returned when the fingerprint operation is started on a file. The `session-id` of the fingerprint operation can be used to get the progress of an ongoing fingerprint operation as well as the complete fingerprint output for the file once the operation is completed.

* **Volume:**

The name of the volume on which the file resides.

* **Path:**

The absolute path of the file on which the fingerprint is calculated. The value begins with `/vol/<volumename>`.

* **Data Fingerprint:**

The digest value of data of the file. The fingerprint is base64 encoded. This field is not included if the scope is *metadata-only*.

* Metadata Fingerprint:

The digest value of metadata of the file. The metadata fingerprint is calculated for file size, file ctime, file mtime, file crtime, file retention time, file uid, file gid, and file type. The fingerprint is base64 encoded. This field is not included if the scope is *data-only*.

* Fingerprint Algorithm:

The digest algorithm which is used for the fingerprint computation. Fingerprint is computed using *md5* or *sha-256* digest algorithm.

* Fingerprint Scope:

The scope of the file which is used for the fingerprint computation. Fingerprint is computed over *data-only*, *metadata-only*, or *data-and-metadata*.

* Fingerprint Start Time:

The start time of the fingerprint computation in seconds since 1 January 1970 00:00:00 in GMT timezone.

* Formatted Fingerprint Start Time:

The start time of the fingerprint computation in a human-readable format <day> <month> <day of month> <hour>:<min>:<sec><year> in GMT timezone.

* Fingerprint Version:

The version of the fingerprint output format.

* SnapLock License:

The status of the SnapLock license.

* Vserver UUID:

A universal unique identifier for the Vserver on which the file exists.

* Volume MSID:

The mirror set identifier of the volume where the file resides.

* Volume DSID:

The data set identifier of the volume where the file resides.

* Hostname:

The name of the storage system where the fingerprint operation is performed.

* Filer ID:

The NVRAM identifier of the storage system.

* Volume Containing Aggregate:

The name of the aggregate in which the volume resides.

* Aggregate ID:

A universal unique identifier for the aggregate containing the volume.

* SnapLock System ComplianceClock:

The System ComplianceClock time in seconds since 1 January 1970 00:00:00 in GMT timezone if it is initialized.

* Formatted SnapLock System ComplianceClock:

The System ComplianceClock time in a human-readable format <day> <month> <day of month> <hour>:<min>:<sec><year> in GMT timezone if it is initialized.

* Volume SnapLock Type:

The type of the SnapLock volume. This value is only given for SnapLock volumes. Possible values are *compliance* and *enterprise*.

* Volume ComplianceClock:

The volume ComplianceClock time in seconds since 1 January 1970 00:00:00 in GMT timezone. This has a value only for SnapLock volumes.

* Formatted Volume ComplianceClock:

The volume ComplianceClock time in a human-readable format <day> <month> <day of month> <hour>:<min>:<sec><year> in GMT timezone. This has a value only for SnapLock volumes.

* Volume Expiry Date:

The expiry date of the SnapLock volume in seconds since 1 January 1970 00:00:00 in GMT timezone. The volume expiry date can be in wraparound format.

* Is Volume Expiry Date Wraparound:

The value is *true* if the volume expiry date is in wraparound format. The wraparound format indicates that dates after 19 January 2038 are mapped from 1 January 1970 through 31 December 2002 to 19 January 2038 through 19 January 2071.

* Formatted Volume Expiry Date:

The expiry date of the SnapLock volume in a human-readable format <day> <month> <day of month> <hour>:<min>:<sec><year> in GMT timezone. The volume expiry date can be in wraparound format.

* Filesystem ID:

The filesystem identifier of the volume on which the file resides.

* File ID:

A unique number within the filesystem identifying the file.

* File Type:

The type of the file. Possible values include: *worm*, *worm_appendable*, *worm_active_log*, *worm_log*, and *regular*.

* File Size:

The size of the file in bytes.

* Creation Time:

The creation time of the file in seconds since 1 January 1970 00:00:00 in GMT timezone.

* Formatted Creation Time:

The creation time of the file in a human-readable format <day> <month> <day of month> <hour>:<min>:<sec><year> in GMT timezone.

* Modification Time:

The last modification time of the file in seconds since 1 January 1970 00:00:00 in GMT timezone.

* Formatted Modification Time:

The last modification time of the file in a human-readable format <day> <month> <day of month> <hour>:<min>:<sec><year> in GMT timezone.

* Changed Time:

The last changed time of the file attributes in seconds since 1 January 1970 00:00:00 in GMT timezone. Time is taken from the system clock for regular files and from the volume ComplianceClock for WORM files when they are committed. The changed time can be in wraparound format.

* Is Changed Time Wraparound:

The value is *true* if the last changed time of the file attributes is in wraparound format. The wraparound format indicates that dates after 19 January 2038 are mapped from 1 January 1970 through 31 December 2002 to 19 January 2038 through 19 January 2071.

* Formatted Changed Time:

The last changed time of the file attributes in a human-readable format <day> <month> <day of month> <hour>:<min>:<sec><year> in GMT timezone. The changed time can be in wraparound format.

* Retention Time:

The retention time of the files committed to WORM on SnapLock volumes in seconds since 1 January 1970 00:00:00 in GMT timezone. The retention time can be in wraparound format.

* Is Retention Time Wraparound:

The value is *true* if the retention time of the file is in wraparound format. The wraparound format indicates that dates after 19 January 2038 are mapped from 1 January 1970 through 31 December 2002 to 19 January 2038 through 19 January 2071.

* Formatted Retention Time:

The retention time of the files protected by SnapLock in a human-readable format <day> <month> <day of month> <hour>:<min>:<sec><year> in GMT timezone. The retention time can be in wraparound format.

* Access Time:

The last access time of the regular files on SnapLock volumes and files on non-SnapLock volumes attributes in seconds since 1 January 1970 00:00:00 in GMT timezone.

* Formatted Access Time:

The last access time of the regular files on SnapLock volumes and files on non-SnapLock volumes attributes in a human-readable format <day> <month> <day of month> <hour>:<min>:<sec><year> in GMT timezone.

* Owner ID:

The integer identifier of the owner of the file.

* Group ID:

The integer identifier of the group owning the file.

* Owner SID:

The security identifier of the owner of the file when it has NTFS security style.

* Fingerprint End Time:

The end time of the fingerprint computation in seconds since 1 January 1970 00:00:00 in GMT timezone.

* Formatted Fingerprint End Time:

The end time of the fingerprint computation in a human-readable format <day> <month> <day of month> <hour>:<min>:<sec><year> in GMT timezone.

* Litigation Count:

The number of litigations on the file.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

-session-id <integer> - Session ID of Fingerprint Operation

Specifies the session-id of the fingerprint operation whose output is to be displayed. It is a unique identifier for the fingerprint operation. This session-id is returned when the fingerprint operation is started on a file.

Examples

The following example displays the fingerprint information of the fingerprint session identified by session-id 17039367:

```
cluster1::> volume file fingerprint dump -session-id 17039367
Vserver:vs1
                                Session-ID:17039367
                                Volume:nfs_slc
                                Path:/vol/nfs_slc/worm
                                Data
Fingerprint:MOFJVevxNSJm3C/4Bn5oEEYH51CrudOzZYK4r5Cfy1g=
                                Metadata
Fingerprint:8iMjqJXiNcqqXT5XuRhLiEwIrJEihDmwS0hrexnjgmc=
                                Fingerprint Algorithm:SHA256
                                Fingerprint Scope:data-and-metadata
                                Fingerprint Start Time:1460612586
                                Formatted Fingerprint Start Time:Thu Apr 14 05:43:06 GMT 2016
                                Fingerprint Version:3
                                SnapLock License:available
                                Vserver UUID:acf7ae64-00d6-11e6-a027-
0050569c55ae
                                Volume MSID:2152884007
                                Volume DSID:1028
                                Hostname:cluster1
                                Filer ID:5f18eda2-00b0-11e6-914e-
6fb45e537b8d
                                Volume Containing Aggregate:slc_aggr
                                Aggregate ID:c84634aa-c757-4b98-8f07-
eeef32565f67
                                SnapLock System ComplianceClock:1460610635
                                Formatted SnapLock System ComplianceClock:Thu Apr 14 05:10:35 GMT 2016
                                Volume SnapLock Type:compliance
                                Volume ComplianceClock:1460610635
```

```

Formatted Volume ComplianceClock:Thu Apr 14 05:10:35 GMT 2016
      Volume Expiry Date:1465880998
Is Volume Expiry Date Wraparound:false
      Formatted Volume Expiry Date:Tue Jun 14 05:09:58 GMT 2016
      Filesystem ID:1028
      File ID:96
      File Type:worm
      File Size:1048576
      Creation Time:1460612515
      Formatted Creation Time:Thu Apr 14 05:41:55 GMT 2016
      Modification Time:1460612515
Formatted Modification Time:Thu Apr 14 05:41:55 GMT 2016
      Changed Time:1460610598
Is Changed Time Wraparound:false
      Formatted Changed Time:Thu Apr 14 05:09:58 GMT 2016
      Retention Time:1465880998
Is Retention Time Wraparound:false
      Formatted Retention Time:Tue Jun 14 05:09:58 GMT 2016
      Access Time:-
      Formatted Access Time:-
      Owner ID:0
      Group ID:0
      Owner SID:-
      Fingerprint End Time:1460612586
Formatted Fingerprint End Time:Thu Apr 14 05:43:06 GMT 2016
      Litigation Count:0

```

volume file fingerprint show

Display fingerprint operation status

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume file fingerprint show` command returns information for one or several fingerprint operations. This command requires either `-session-id` or `-vserver` and `-volume`.

Parameters

```
{ [-fields <fieldname>,...]
```

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

```
| [-instance ] }
```

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-session-id <integer>] - Session ID of Fingerprint Operation

If this parameter is specified, the command returns the progress of the fingerprint operation of the specified session-id. The session-id is a unique identifier for the fingerprint operation that is returned when the fingerprint operation is started on a file.

[-vserver <vserver name>] - Vserver

If this parameter is specified, -volume must also be specified. When queried with -vserver and -volume, the command returns the progress of all the fingerprint operations running on that particular volume.

[-volume <volume name>] - Volume Name

If this parameter is specified, -vserver must also be specified. When queried with -vserver and -volume, the command returns the progress of all the fingerprint operations running on that particular volume.

[-file </vol/<volume name>/<file path>>] - File Path

If this parameter is specified, the command returns the progress of all fingerprint operations on the specified file.

[-operation-status {Unknown|In-Progress|Failed|Aborting|Completed}] - Operation Status

If this parameter is specified, the command returns the progress of all fingerprint operations with matching status value.

[-progress-percentage <integer>] - Progress Percentage

If this parameter is specified, the command returns the progress of all fingerprint operations with matching progress percentage value.

Examples

The following example displays the progress of all the fingerprint operations running on volume *nfs_slc*:

```

cluster1::> volume file fingerprint show -vserver vs0 -volume nfs_slc

Progress
File-Path                               Session-ID                               Status
Percentage
-----
-----
/vol/nfs_slc/worm                       17104897                               Completed
100
/vol/nfs_slc/worm_appedable             17104898                               Completed
100
/vol/nfs_slc/regular                    17104899                               In-Progress
30
3 entries were displayed.

```

volume file fingerprint start

Start a file fingerprint computation on a file

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume file fingerprint start` command starts the fingerprint computation on a file. The fingerprint computation is started on the file, and a session-id is returned. This session-id is a unique identifier for the fingerprint operation and can be used to get the progress of an ongoing fingerprint operation as well as the complete fingerprint output for the file once the operation is completed.

Parameters

-vserver <vserver name> - Vserver

Specifies the name of the vservers which owns the volume on which the file resides.

-file </vol/<volume name>/<file path>> - Path

Specifies the absolute path of the file on which fingerprint needs to be calculated. The value begins with `/vol/<volumename>`.

[-algorithm {MD5|SHA256}] - Fingerprint Algorithm

Specifies the digest algorithm which is used for the fingerprint computation.

Fingerprint can be computed using one of the following digest algorithms:

- md5
- sha-256

[-scope {data-and-metadata|data-only|metadata-only}] - Fingerprint Scope

Specifies the scope of the file which is used for the fingerprint computation.

Fingerprint can be computed using one of the following scope:

- data-only
- metadata-only
- data-and-metadata

Examples

The following example starts computing fingerprint over data and metadata for file `/vol/nfs_slc/worm` using `md5` hash algorithm. The file `/vol/nfs_slc/worm` is stored in volume `nfs_slc` on Vserver `vs0`.

```
cluster1::> volume file fingerprint start -vserver vs0 -scope data-and-
metadata -algorithm md5 -file /vol/nfs_slc/worm
File fingerprint operation is queued. Run "volume file fingerprint show
-session-id 16973825" to view the fingerprint session status.
```

volume file retention show

Display retention time of a file protected by SnapLock.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume file retention show` command displays the retention time of a file protected by SnapLock given `-vserver` and `-file`.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Name

This specifies the Vserver on which the volume is located having the file in a SnapLock volume. An exact value is required for this parameter.

[-file </vol/<volume name>/<file path>>] - File path

This specifies the absolute path of the file in a SnapLock volume. The value begins with `/vol/<volumename>`. An exact value is required for this parameter.

[-retention-time <integer>] - Retention Time of the File

If this parameter is specified, the command displays information only about the file protected by SnapLock that matches the specified `retention-time`. The value represents file retention time in seconds since 1 January 1970 00:00:00.

[-formatted-retention-time <text>] - Formatted Retention Period

If this parameter is specified, the command displays information only about the file protected by SnapLock that matches the specified `formatted-retention-time`. The expiry date format is `<day> <month> <day of month> <hour>:<min>:<sec><year>` in GMT timezone taking care of wraparound. A value of *infinite* indicates that this file has infinite retention time. A value of *indefinite* indicates that this file has indefinite retention time. A value of *unspecified* indicates that this file will be retained forever; however, the retention time can be changed to an absolute value.

[-is-wraparound {true|false}] - Is Retention Time Wraparound

If this parameter is specified, the command displays information only about the file protected by SnapLock that matches the specified `-is-wraparound`. The value is `true` if the date represented in retention time is in wraparound format. The wraparound format indicates that dates after 19 January 2038 are mapped from 1 January 1970 through 31 December 2002 to 19 January 2038 through 19 January 2071.

[-seconds-until-expiry <integer>] - Seconds until the File Expires

If this parameter is specified, the command displays information only about the file protected by SnapLock that matches the specified `-seconds-until-expiry`. The value represents the number of seconds until the expiration time.

[`-is-expired {true|false}`] - Has the File Expired

If this parameter is specified, the command displays information only about the file protected by SnapLock that matches the specified `-is-expired`. The value represents file expiration status. The value is `false` if the file is under active retention or `true` if the file is past its expiry time.

Examples

The following example displays the retention time of the file `/vol/nfs1/file1`:

```
cluster1::> volume file retention show -vserver vs0 -file /vol/nfs1/file1
Vserver : vs0
Path : /vol/nfs1/file1
Retention Time (Secs from Epoch) : 1439111404
Formatted Retention Time : Sun Aug 9 09:10:04 GMT 2015
Is Retention Time Wraparound : false
```

volume flexcache commands

volume flexcache config-refresh

Refresh FlexCache configuration for a peer volume

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `volume flexcache config-refresh` command is used to refresh the FlexCache configuration. It can be used to update the configuration on either the FlexCache or origin of a FlexCache cluster. This command only needs to be run if the automatic refresh failed and the corresponding EMS was generated.



This command must be run on the peer cluster. For example, refresh of origin of a FlexCache volume must be run from the FlexCache cluster. Refresh a FlexCache volume must be run from the of an origin of a FlexCache cluster.

Parameters

`-peer-vserver <vserver name>` - Peer Vserver (privilege: advanced)

Name of the Vserver for which the configuration is being refreshed.

`-peer-volume <volume name>` - Peer Volume (privilege: advanced)

Name of the volume for which the configuration is being refreshed.

-peer-endpoint-type {cache|origin} - Origin/Cache Volume (privilege: advanced)

The `peer-endpoint-type` specifies the FlexCache endpoint type of the peer volume. Possible values are `cache` for FlexCache volumes and `origin` for origin of a FlexCache volumes.

Examples

The following example triggers config-refresh on origin of a FlexCache volume "origin1".

```
cluster1::> flexcache config-refresh -peer-vserver vs34 -peer-volume  
origin1 -peer-endpoint-type origin  
          (volume flexcache config-refresh)
```

The following example triggers config-refresh on FlexCache volume "fc1".

```
cluster1::> flexcache config-refresh -peer-vserver vs34 -peer-volume fc1  
-peer-endpoint-type cache  
          (volume flexcache config-refresh)
```

The following example triggers config-refresh on FlexCache volume "fc1" with an incorrect peer-endpoint-type.

```
cluster1::> flexcache config-refresh -peer-vserver vs34 -peer-volume fc1  
-peer-endpoint-type origin  
          (volume flexcache config-refresh)  
Error: command failed: Failed to store the configuration for peer volume  
"fc1"  
          in Vserver "vs34" on cluster "cluster1". Check the  
FlexCache  
          configuration on the local cluster and retry the  
operation.
```

volume flexcache create

Create a new cache relationship

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

```
The `volume flexcache create` command is used to create a FlexCache  
volume. It also creates the relationship between the FlexCache volume and  
the specified origin volume.
```



If the `vserver` and `origin-vserver` are different, the Vservers must be in a peer relationship.

Parameters

-vserver <Vserver Name> - Vserver Name

This specifies the Vserver on which the FlexCache volume is to be created.

-volume <volume name> - Cache Volume Name

This specifies the name of the FlexCache volume that is to be created.

{ [-aggr-list <aggregate name>,...] - List of Aggregates for FlexGroup Constituents

Specifies an array of names of aggregates to be used for creating the FlexCache volume. Each entry in the list will create a constituent on the specified aggregate. The root constituent will always be placed on the first aggregate in the list, unless `optimize-aggr-list` is specified as `true`. An aggregate may be specified multiple times to have multiple constituents created on it. This parameter only applies to FlexGroups.

[-aggr-list-multiplier <integer>] - Aggregate List Repeat Count

Specifies the number of times to iterate over the aggregates listed with the `-aggr-list` parameter when creating a FlexGroup. The aggregate list will be repeated the specified number of times. Example:

```
-aggr-list aggr1,aggr2 -aggr-list-multiplier 2
```

will cause four constituents to be created in the order `aggr1`, `aggr2`, `aggr1`, `aggr2`.

+

The default value is 4.

[-optimize-aggr-list {true|false}] - Have the System Optimize the Order of the Aggregate List

Specifies whether to create the constituents of the FlexCache volume on the aggregates specified in the `aggr-list` in the order they are specified, or whether the system should optimize the ordering of the aggregates. If this value is `true`, the system will optimize the ordering of the aggregates specified in the `aggr-list`. If this value is `false` the order of the `aggr-list` will be unchanged. The default value is `false`. This parameter only applies to FlexGroups.

[-auto-provision-as <FlexGroup>] - Automatically Provision as Volume of Type

Use this parameter to automatically select existing aggregates for provisioning the FlexCache volume.

[-use-tiered-aggregate {true|false}] - Automatically Provision FlexGroup on FabricPool }

This parameter specifies whether or not FabricPool-enabled aggregates are selected when provisioning a FlexCache volume using the `auto-provision-as` parameter. Only FabricPool-enabled aggregates are used if this parameter is set to `true` and only non-FabricPool-enabled aggregates are used if this parameter is set to `false`. The default value is `false`. Tiering-policy is always `none` for FlexCache.

This parameter only applies to FlexCaches created using the `-auto-provision-as` parameter.

[`-size` {<integer>[KB|MB|GB|TB|PB]}] - Volume Size

This optionally specifies the size of the FlexCache volume. The size is specified as a number followed by a unit designation: k (kilobytes), m (megabytes), g (gigabytes), or t (terabytes). If the unit designation is not specified, bytes are used as the unit, and the specified number is rounded up to the nearest 4 KB. If the size parameter is not specified, it defaults to 10% of the origin volume size.

[`-origin-vserver` <vserver name>] - Origin Vserver Name

This specifies the name of the Vserver where the origin volume is located.

`-origin-volume` <volume name> - Origin Volume Name

This specifies the name of the origin volume.

[`-junction-path` <junction path>] - Cache Junction Path

This optionally specifies the FlexCache volume's junction path. The junction path name is case insensitive and must be unique within a Vserver's namespace.

[`-foreground` {true|false}] - Foreground Process

This specifies whether the operation runs in the foreground. The default setting is `true` (the operation runs in the foreground). When set to `true`, the command will not return until the operation completes.

[`-s`, `-space-guarantee` {none|volume}] - Space Guarantee Style

This optionally specifies the space guarantee style for the volume. A value of `volume` reserves space on the aggregates for the entire volume. A value of `none` reserves no space on the aggregates, meaning that writes can fail if an aggregate runs out of space. The default setting for the FlexCache volumes is `none`.

[`-is-smb-enabled` {true|false}] - SMB access Enabled/Disabled (privilege: advanced)

This parameter specifies CIFS access for the volume. The default value is `true`.

[`-preserve-msid` {true|false}] - Preserve Origin MSID on the Cache (privilege: advanced)

Create the cache with the same MSID as the origin. Set this parameter to `true` to create FlexCache DR. The default setting for the FlexCache volumes is `false` to avoid a situation where a FlexCache create operation using `-preserve-msid` fails on a Vserver because of collision with a previously used MSID.

[`-is-global-file-locking-enabled` {true|false}] - Global File Locking Mode Enabled/Disabled

This parameter specifies the global file locking mode for the volume. The default value of global file locking mode is `false`.

[`-is-nfsv4-enabled` {true|false}] - NFSv4 access Enabled/Disabled (privilege: advanced)

This parameter specifies NFSv4 access for the volume. The default value is `true`.

[`-is-writeback-enabled` {true|false}] - Is Writeback Enabled (privilege: advanced)

This parameter specifies whether writeback is enabled for the FlexCache volume. Writeback is a storage method where data is first written to the FlexCache volume and then written to the origin of a FlexCache volume. The default value is `false`.

[`-is-relative-size-enabled` {true|false}] - Is Relative Sizing Enabled/Disabled (privilege: advanced)

This parameter specifies whether the relative sizing is enabled for the FlexCache volume. Relative sizing is introduced as a part of follow the origin feature. When relative sizing is enabled, it will block any

modifications done manually in the absolute size of the FlexCache. The size of FlexCache will be calculated and entered automatically based on the size of the origin. The default value for create is *false*.

[~~-relative-size-percentage~~ <integer>] - Relative FlexCache Size Percentage (privilege: advanced)

This parameter specifies the percent size FlexCache should have relative to the total size of the origin. The default value for create is 10%.

[~~-override-encryption~~ {true|false}] - Override FlexCache Volume Encryption

Use this parameter to create a plaintext FlexCache volume of an encrypted origin volume. The default value is *false*.



This parameter cannot be modified after the FlexCache volume is created.

[~~-is-atime-scrub-enabled~~ {true|false}] - Is atime-Based Scrubbing Enabled/Disabled (privilege: advanced)

This parameter specifies whether automatic scrubbing of inactive files based on atime is enabled. When enabled, inactive files are automatically scrubbed based on specified atime duration. The scrubbing operation will only occur if `-atime-update` is enabled on the FlexCache volume. The default value for the atime scrubber is *false*.

[~~-atime-scrub-period~~ <integer>] - Days After Which An Inactive File Is Scrubbed (privilege: advanced)

This parameter specifies the number of days for inactivity on a file after which the file will be scrubbed out. The default value is 30 days.

[~~-is-cifs-change-notify-enabled~~ {true|false}] - Is CIFS CHANGE_NOTIFY Enabled/Disabled (privilege: advanced)

This parameter specifies whether the FlexCache volume sends change notifications to CIFS clients. The default value is *false*.

Examples

The following example triggers FlexCache volume create:

```
cluster1::> flexcache create -vserver vs34 -volume fc1 -aggr-list
aggr34,aggr43 -origin-volume origin1 -size 400m
(volume flexcache create)
[Job 894] Job succeeded: Successful
```

```
cluster1::> flexcache create -vserver vs34 -volume fc3 -auto-provision-as
flexgroup -origin-volume origin1 -size 400m
(volume flexcache create)
[Job 898] Job succeeded: Successful
```

```
cluster1::> flexcache create -vserver vs34 -volume fc4 -aggr-list
aggr34,aggr43 -origin-volume origin1 -size 400m -junction-path /fc4
(volume flexcache create)
[Job 903] Job succeeded: Successful
```

volume flexcache delete

Delete a cache relationship

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume flexcache delete` command deletes the specified FlexCache volumes and their relationships.

NOTE:

- FlexCache volumes must be offline (see [volume offline](#)) to be deleted.

Parameters

-vserver <Vserver Name> - Vserver Name

This specifies the name of the Vserver from which the FlexCache volume is to be deleted.

-volume <volume name> - Cache Volume Name

This specifies the name of the FlexCache volume that is to be deleted.

[-foreground {true|false}] - Foreground Process

This specifies whether the operation runs in the foreground. The default setting is *true* (the operation runs in the foreground). When set to *true*, the command will not return until the operation completes.

Examples

The following example deletes FlexCache volume "fc1":

```
cluster1::> flexcache delete -volume fc1 -vserver vs34
(volume flexcache delete)
```

```
Error: command failed: Volume fc1 in Vserver vs34 must be offline to be
deleted. Use "volume offline -vserver vs34 -volume fc1" command to
offline the volume
```

```
cluster1::> volume offline -vserver vs34 -volume fc1
Volume "vs34:fc1" is now offline.
```

```
cluster1::> flexcache delete -volume fc1 -vserver vs34
(volume flexcache delete)
[Job 891] Job succeeded: Successful
```

```
cluster1::> flexcache delete -volume fc1 -vserver vs34 -foreground false
(volume flexcache delete)
```

Related Links

- [volume offline](#)

volume flexcache prepare-to-downgrade

Disables FlexCache features not supported by a previous version

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command disables the FlexCache feature enabled for a version.

Parameters

[~~-disable-feature-set~~ <ClusterVersion>] - Cluster Version (privilege: advanced)

This specifies the version for which we are running this command.

Examples

The following example disables FlexCache features for cluster "cluster1":

```
cluster1::> flexcache prepare-to-downgrade -disable-feature-set 9.8.0
FlexCache features introduced in 9.8 are successfully disabled.
```

volume flexcache show

Display cache relationships

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume flexcache show` command displays information about FlexCache volumes. The command output depends on the parameter or parameters specified with the command. If no parameters are specified, the command displays the following information about all FlexCache volumes:

- Vserver name

- Volume name
- Size
- Space guarantee
- Origin Vserver
- Origin volume
- Origin cluster

To display detailed information about all FlexCache volumes, run the command with the `-instance` parameter.

Parameters

{ [-fields <fieldname>,...]

This specifies the fields that need to be displayed.

| [-instance] }

If this parameter is specified, the command displays information about all values.

[-vserver <Vserver Name>] - Vserver Name

If this parameter and the `-volume` parameter are specified, the command displays detailed information about the specified volume. If this parameter is specified by itself, the command displays information about FlexCache volumes on the specified Vserver.

[-volume <volume name>] - Cache Volume Name

If this parameter is specified, the command displays detailed information about the specified FlexCache volume.

[-aggr-list <aggregate name>,...] - List of Aggregates for FlexGroup Constituents

If this parameter is specified, the command displays information only about the FlexCache volume or volumes that are located on the specified list of storage aggregates.

[-size {<integer>[KB|MB|GB|TB|PB] }] - Volume Size

If this parameter is specified, the command displays information only about the FlexCache volume or volumes that have the specified size. Size is the maximum amount of space a volume can consume from its associated aggregate(s), including user data, metadata, Snapshot copies, and Snapshot reserve. Note that for volumes without a `-space-guarantee` of `volume`, the ability to fill the volume to this maximum size depends on the space available in the associated aggregate or aggregates.

[-flexgroup-msid <integer>] - Cache FlexGroup MSID (privilege: advanced)

If this parameter is specified, the command displays information only about the FlexCache volume or volumes that have the specified FlexGroup master data-set ID.

[-origin-vserver <vserver name>] - Origin Vserver Name

If this parameter is specified, the command displays information only about the FlexCache volume or volumes which have a relationship with the specified origin-vserver.

[-origin-vserver-uuid <UUID>] - Origin Vserver UUID (privilege: advanced)

If this parameter is specified, the command displays information only about the FlexCache volume or volumes which have a relationship with the origin-vserver UUID.

[-origin-volume <volume name>] - Origin Volume Name

If this parameter is specified, the command displays information only about the FlexCache volume or volumes that have a relationship with the specified origin-volume.

[-origin-volume-msid <integer>] - Origin Volume MSID (privilege: advanced)

If this parameter is specified, the command displays information only about the FlexCache volume or volumes that have a relationship with the specified origin master data set ID.

[-origin-cluster <Cluster name>] - Origin Cluster Name

If this parameter is specified, the command displays information only about the FlexCache volume or volumes that have a relationship with the specified origin-cluster.

[-junction-path <junction path>] - Cache Junction Path

If this parameter is specified, the command displays information only about the FlexCache volume or volumes that have the specified junction path.

[-create-time <Date>] - FlexCache Create Time

If you specify this parameter, the command displays information only about the FlexCache volume or volumes for which the `create-time` option matches the specified input.

[-s, -space-guarantee {none|volume}] - Space Guarantee Style

If this parameter is specified, the command displays information only about the FlexCache volume or volumes that have the specified space guarantee style.

[-relationship-type {Intravserver|Intervserver|Intercluster}] - Relationship Type (privilege: advanced)

If this parameter is specified, the command displays the FlexCache volumes matching the relationship type.

[-is-smb-enabled {true|false}] - SMB access Enabled/Disabled (privilege: advanced)

If this parameter is specified, the command displays the FlexCache volumes matching the value for CIFS access.

[-preserve-msid {true|false}] - Preserve Origin MSID on the Cache (privilege: advanced)

If this parameter is specified, the command displays information only about the FlexCache volume or volumes that match the given value for the `-preserve-msid`. If this parameter is `true`, then FlexCache DR is enabled for the FlexCache volume.

[-is-global-file-locking-enabled {true|false}] - Global File Locking Mode Enabled/Disabled

If this parameter is specified, the command displays the FlexCache volumes matching the global file locking mode.

[-is-nfsv4-enabled {true|false}] - NFSv4 access Enabled/Disabled (privilege: advanced)

If this parameter is specified, the command displays the FlexCache volumes matching the value for NFSv4 access.

[-is-writeback-enabled {true|false}] - Is Writeback Enabled (privilege: advanced)

This parameter specifies whether writeback is enabled for the FlexCache volume. Writeback is a storage method where data is first written to the FlexCache volume and then written to the origin of a FlexCache volume.

`[-is-relative-size-enabled {true|false}] - Is Relative Sizing Enabled/Disabled (privilege: advanced)`

If this parameter is specified, the command displays information only about the FlexCache volume or volumes that match the given value for `-is-relative-size-enabled`. When relative size is enabled, the FlexCache size will be automatically updated based on the size of the origin.

`[-relative-size-percentage <integer>] - Relative FlexCache Size Percentage (privilege: advanced)`

If this parameter is specified, the command displays the information only about the FlexCache volume or volumes for which the value of `-relative-size-percentage` is specified. The relative size percentage is used to calculate the size of a FlexCache volume when the `is-relative-size-enabled` field is `true`.

`[-override-encryption {true|false}] - Override FlexCache Volume Encryption`

If this parameter is specified, the command displays information only about the FlexCache volume or volumes that match the given value for the `-override-encryption`. If this parameter is `true`, then the volume is a plaintext FlexCache volume.

`[-is-atime-scrub-enabled {true|false}] - Is atime-Based Scrubbing Enabled/Disabled (privilege: advanced)`

If this parameter is specified, the command displays information only about the FlexCache volume or volumes that match the given value for `-is-atime-scrub-enabled`. When atime scrubber is enabled, the FlexCache volume will automatically scrub inactive files based on the specified atime duration.

`[-atime-scrub-period <integer>] - Days After Which An Inactive File Is Scrubbed (privilege: advanced)`

If this parameter is specified, the command displays information only about the FlexCache volume or volumes for which the value of `-atime-scrub-period` is specified. The atime scrub period is used to remove the files from the FlexCache volume when the duration is met. The default duration to scrub files is set to *30 days*.

`[-is-cifs-change-notify-enabled {true|false}] - Is CIFS CHANGE_NOTIFY Enabled/Disabled (privilege: advanced)`

If this parameter is specified, the command displays information only about the FlexCache volume or volumes that match the given value for `-is-cifs-change-notify-enabled`. When CIFS CHANGE_NOTIFY is enabled, the FlexCache volume will notify clients about changes on a directory.

Examples

The following example displays information about all FlexCache volumes on the Vserver named "vs34":

```

cluster1::> flexcache show -vserver vs34
          (volume flexcache show)
          Vserver Volume      Size      Origin-Vserver Origin-Volume
Origin-Cluster
-----
vs34     fc1          800MB     vs34          origin1
cluster-2
vs34     fc2          800MB     vs34          origin1
cluster-2
2 entries were displayed.

```

The following example displays detailed information about a FlexCache volume named fc1 on an SVM named vs34:

```

cluster1::> flexcache show -vserver vs34 -volume fc1 -instance
          (volume flexcache show)
Vserver: vs34
          Cache Volume Name: fc1
List of Aggregates for FlexGroup Constituents: aggr34
          Volume Size: 800MB
          Cache Flexgroup MSID: 2155934574
          Origin Vserver Name: vs34
          Origin Vserver UUID: a8717aeb-2826-11e8-bf56-
00505695f37a
          Origin Volume Name: origin1
          Origin Volume MSID: 2155934545
          Origin Cluster Name: cluster-2
          Cache Junction Path: -
          FlexCache Create Time: Thu Aug 23 04:36:19 2018
          Relationship Type: inter-vserver
          Preserve Origin MSID on the Cache: false

```

volume flexcache sync-properties

Sync volume properties

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `volume flexcache sync-properties` command is used to sync a FlexCache volume's properties with its origin volume or to update FlexCache origin volume name in FlexCache relationship during a origin volume rename operation in disconnected mode.



This command should be used when there is a failure to sync properties from the origin of a FlexCache volume to the FlexCache volume or to update the FlexCache origin volume name in the FlexCache relationship in the case of a volume rename from the origin in disconnected mode.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

Specifies the SVM on which the volume is located.

-volume <volume name> - Volume Name (privilege: advanced)

Specifies the volume name whose properties need to be synced. If the volume is a FlexCache volume, its properties are synced with its origin volume. If the volume is a FlexCache origin volume, the origin volume name is updated in the FlexCache relationship where the origin volume is renamed in disconnected mode.

Examples

```
cluster1::> flexcache sync-properties -volume fc1 -vserver vs34
(volume flexcache sync-properties)
```

volume flexcache config modify

Modify cache configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume flexcache config modify` command is used to modify the configuration for a FlexCache volume.

Parameters

-vserver <Vserver Name> - Vserver Name (privilege: advanced)

This specifies the Vserver in which the FlexCache volume exists.

-volume <volume name> - Cache Volume Name (privilege: advanced)

This specifies the name of the FlexCache volume.

[-is-cifs-enabled {true|false}] - Is CIFS Access Enabled (privilege: advanced)

This parameter specifies whether CIFS access is allowed on a FlexCache volume.

[-is-nfsv4-enabled {true|false}] - Is NFSv4 Access Enabled (privilege: advanced)

This parameter specifies whether NFSv4 access is allowed on a FlexCache volume.

`[-is-writeback-enabled {true|false}] - Is Writeback Enabled (privilege: advanced)`

This parameter specifies whether writeback is enabled for the FlexCache volume. Writeback is a storage method where data is first written to the FlexCache volume and then written to the origin of a FlexCache volume.

`[-is-relative-size-enabled {true|false}] - Is Relative Sizing Enabled (privilege: advanced)`

This parameter specifies whether the relative sizing is enabled for the FlexCache volume. Relative sizing is introduced as a part of follow the origin feature. When relative sizing is enabled, it will block any modifications done manually in the absolute size of the FlexCache. The size of FlexCache will be calculated and entered automatically based on the size of the origin. The default value for create is *false*.

`[-relative-size-percentage <integer>] - Relative FlexCache Size Percentage (privilege: advanced)`

This parameter specifies the percent size FlexCache should have relative to the total size of the origin. The default value for create is 10%.

`[-is-atime-scrub-enabled {true|false}] - Is atime based Scrubbing Enabled/Disabled (privilege: advanced)`

This parameter specifies if the files can be scrubbed out automatically if they have not been accessed for a specified duration. This feature to scrub out inactive files is dependent on enabling the atime feature on the FlexCache volume. The default value for create is *disabled*.

`[-atime-scrub-period <integer>] - Days after which an inactive file is scrubbed (privilege: advanced)`

This parameter specifies the duration in days after which the inactive files can be scrubbed out from the cache volume. The default value for create is *30 days* when atime-scrub is enabled.

`[-is-cifs-change-notify-enabled {true|false}] - Is CIFS CHANGE_NOTIFY Enabled/Disabled (privilege: advanced)`

This parameter specifies whether the FlexCache volume sends change notifications to CIFS clients. The default value is *false*.

Examples

The following example triggers FlexCache volume configuration modify:

```
cluster1::> volume flexcache config modify -vserver vs34 -volume fc1 -is
-cifs-enabled true
```

```
cluster1::> volume flexcache config modify -volume originvol -is-cifs
-enabled false
(volume flexcache config modify)
Error: command failed: FlexCache config modify is only applicable for
volumes
with "flexcache-endpoint-type" cache.
```

volume flexcache config show

Display cache configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

```
The `volume flexcache config show` command is used to view the FlexCache
configuraiton for a FlexCache volume.
```

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <Vserver Name>] - Vserver Name (privilege: advanced)

If this parameter and the `-volume` parameter are specified, the command displays config information about the specified volume. If this parameter is specified by itself, the command displays information about FlexCache volumes on the specified Vserver.

[-volume <volume name>] - Cache Volume Name (privilege: advanced)

If this parameter is specified, the command displays config information about the specified FlexCache volumes matching the specified name.

[-is-cifs-enabled {true|false}] - Is CIFS Access Enabled (privilege: advanced)

If this parameter is specified, the command displays the FlexCache volumes matching the value for CIFS access.

[-is-nfsv4-enabled {true|false}] - Is NFSv4 Access Enabled (privilege: advanced)

If this parameter is specified, the command displays the FlexCache volumes matching the value for NFSv4 access.

[-is-writeback-enabled {true|false}] - Is Writeback Enabled (privilege: advanced)

If this parameter is specified, the command displays the FlexCache volumes matching the writeback mode. Writeback is a storage method where data is first written to the FlexCache volume and then written to the origin of a FlexCache volume.

[-is-relative-size-enabled {true|false}] - Is Relative Sizing Enabled (privilege: advanced)

If this parameter is specified, the command displays the FlexCache volumes matching the value for relative size.

[-relative-size-percentage <integer>] - Relative FlexCache Size Percentage (privilege: advanced)

If this parameter is specified, the command displays the FlexCache volumes matching the percentage for

relative size.

`[-is-atime-scrub-enabled {true|false}]` - Is atime based Scrubbing Enabled/Disabled (privilege: advanced)

If this parameter is specified, the command displays the FlexCache volumes matching the atime scrubbing mode.

`[-atime-scrub-period <integer>]` - Days after which an inactive file is scrubbed (privilege: advanced)

If this parameter is specified, the command displays the FlexCache volumes matching the period for atime scrubbing.

`[-is-cifs-change-notify-enabled {true|false}]` - Is CIFS CHANGE_NOTIFY Enabled/Disabled (privilege: advanced)

If this parameter is specified, the command displays the FlexCache volumes matching the value for CIFS CHANGE_NOTIFY mode.

Examples

The following example triggers FlexCache volume configuration show:

```
cluster1::> flexcache config show
(volume flexcache config show)
Vserver Volume      CIFS Access Enabled
-----
vs1      cachevol      false
vs1      cachevol_2    true
2 entries were displayed.
```

volume flexcache origin show-caches

Display all the caches connected to origin

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume flexcache origin show-caches` command displays FlexCache relationships for origin of a FlexCache volumes on the origin cluster.

Parameters

`{ [-fields <fieldname>,...]`

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

`| [-instance] }`

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-origin-vserver <Vserver Name>] - Vserver Name

If this parameter and the `origin-volume` parameter are specified, the command displays the FlexCache relationship for the specified origin of a FlexCache volume. If this parameter is specified by itself, the command displays all the FlexCache relationships for all origin of a FlexCache volumes in the specified Vserver.

[-origin-volume <volume name>] - Origin Volume

If this parameter is specified, the command displays all the FlexCache relationships for the specified origin of a FlexCache volume.

[-cache-vserver <vserver name>] - Cache Vserver

If this parameter is specified then the command displays FlexCache relationships for the specified Vserver hosting FlexCache volumes.

[-cache-volume <volume name>] - Cache Volume

If this parameter is specified then the command displays FlexCache relationships for the specified FlexCache volume.

[-cache-cluster <Cluster name>] - Cache Cluster

If this parameter is specified then the command displays FlexCache relationships for the specified cluster hosting FlexCache volumes.

[-cache-volume-msid <integer>] - Cache Volume MSID

If this parameter is specified then the command displays FlexCache relationships for the specified FlexCache volume's FlexGroup master data-set ID.

[-relationship-create-time <Date>] - Relationship Create Time

If you specify this parameter, the command displays FlexCache relationships for which the `relationship-create-time` option matches the specified input.

[-relationship-type {Intravserver|Intervserver|Intercluster}] - Relationship Type (privilege: advanced)

If this parameter is specified, the command displays the origin of a FlexCache volume matching the relationship type

[-preserve-msid {true|false}] - Preserve Origin MSID on the Cache (privilege: advanced)

If this parameter is specified, the command displays the origin of a FlexCache volume whose Msid is preserved

Examples

The following example displays information about all origin of a FlexCache volumes on the Vserver named vs34:

```

cluster1::> flexcache origin show-caches -origin-vserver vs34
(volume flexcache origin show-caches)
Origin-Vserver  Origin-Volume  Cache-Vserver  Cache-Volume  Cache-Cluster
-----
vs34            origin1        vs56           fc1_c3_origin1
                                     cluster-3
vs34            origin1        vs34           fc1            cluster-2
vs34            origin1        vs34           fc2            cluster-2
vs34            origin2_new    vs56           fc1_c3_origin2
                                     cluster-3

4 entries were displayed.

```

```

cluster1::> flexcache origin show-caches -origin-vserver vs56 -instance
(volume flexcache origin show-caches)
Vserver: vs56
    Origin Volume: origin
    Cache Vserver: vs56
    Cache Volume: fc1
    Cache Cluster: cluster-3
    Cache Volume MSID: 2156002002
Relationship Create Time: Thu Aug 23 04:36:24 2018
Relationship Type: intra-vserver

```

volume flexcache origin config modify

Modify FlexCache Origin Configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume flexcache origin config modify` command modifies volume specific options of an origin of a FlexCache volume.

Parameters

-origin-vserver <Vserver Name> - Vserver Name (privilege: advanced)

This specifies the Vserver in which the origin of a FlexCache volume is located.

-origin-volume <volume name> - Origin Volume Name (privilege: advanced)

This specifies the origin of a FlexCache volume.

[-is-bli-enabled {true|false}] - Is BLI Enabled (privilege: advanced)

This parameter specifies whether data invalidation at block level is enabled or not. By default, `-is-bli-enabled` is *false*.

[*-is-global-file-locking-enabled* {*true*|*false*}] - Is Global File Locking Enabled (privilege: advanced)

This parameter specifies the global file locking mode for the volume. By default, *-is-global-file-locking-enabled* is *false*. When global file locking mode is enabled for the volume, the "is_disconnected_mode_off_for_locks" flag is always set to *true*.

Examples

The following example modifies the *is-bli-enabled* field value to *true*:

```
cluster1::> flexcache origin config modify -origin-volume origin_flexgroup
-origin-vserver vs3_c2 -is-bli-enabled true
(volume flexcache origin config show)
```

```
cluster1::> flexcache origin config show -origin-volume origin_flexgroup
-origin-vserver vs3_c2
(volume flexcache origin config show)
Vserver: vs3_c2
Origin Volume Name: origin_flexgroup
Is BLI Enabled: true
```

volume flexcache origin config show

Show FlexCache Origin Configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume flexcache origin config show` command displays volume specific options of an origin of a FlexCache volume.

Parameters

{ [*-fields* <fieldname>,...]

If you specify the *-fields* <fieldname>, ... parameter, the command output also includes the specified field or fields. You can use '*-fields ?*' to display the fields to specify.

| [*-instance*] }

If you specify the *-instance* parameter, the command displays detailed information about all fields.

[*-origin-vserver* <Vserver Name>] - Vserver Name (privilege: advanced)

If this parameter and the *-origin-volume* parameter are specified, the command displays options for the specified origin of a FlexCache volume. If this parameter is specified by itself, the command displays options for all origin of a FlexCache volumes in the specified Vserver.

[*-origin-volume* <volume name>] - Origin Volume Name (privilege: advanced)

If this parameter is specified, the command displays options for the specified origin of a FlexCache volume.

[*-is-bli-enabled* {true|false}] - Is BLI Enabled (privilege: advanced)

If this parameter is specified, the command displays the origin of a FlexCache volume matching the specified value. Possible values are:

- *false* - Block level invalidation is not set
- *true* - Block level invalidation is set

[*-is-global-file-locking-enabled* {true|false}] - Is Global File Locking Enabled (privilege: advanced)

If this parameter is specified, the command displays the origin configurations matching the specified value. Global file locking mode is a mode where protocol read locking semantics are enforced across all FlexCaches and origins of a FlexCache volume. To completely disable/enable global file locking mode, kindly use "flexcache origin config modify" command in advanced privilege with *-is-global-file-locking-enabled* set to *true/false* depending upon the requirement. When global file locking mode is enabled for the volume, the "is_disconnected_mode_off_for_locks" flag is always set to *true*. Possible values are:

- *false* - Global file locking mode is not set
- *true* - Global file locking mode is set

Examples

The following example displays FlexCache origin options:

```

cluster1::> volume flexcache origin config show
(volume flexcache origin config show)
Origin-Vserver      Origin-Volume      Is-BLI-Enabled      Is-Global-File-Locking-
                    Origin-Volume      Is-BLI-Enabled      Enabled
-----
-----
vs34                origin1             true                 true
vs34                origin2_new         false                false
2 entries were displayed.

```

volume flexcache prepopulate start

Prepopulate flexcache volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume flexcache populate start` command is used to populate a FlexCache volume with the data of given directory path(s) from the origin volume.



The FlexCache populate job fails when the directories listed with the `-path-list` have one or more non-existing paths.

Parameters

`-cache-vserver <vserver name>` - Cache Vserver Name (privilege: advanced)

This specifies the Vserver that has the FlexCache volume to be populated.

`-cache-volume <volume name>` - Cache Volume Name (privilege: advanced)

This specifies the FlexCache volume which needs to be populated.

`-path-list <text>,...` - List of Paths (privilege: advanced)

This specifies an array of directory names that need to be populated in the FlexCache volume.

`[-isRecursion {true|false}]` - Recursively populate (privilege: advanced)

This specifies that the directories listed with the `-path-list` need to be recursively populated in the `-cache-volume` FlexCache volume.

The default value is true.

`[-exclude-path-list <text>,...]` - Exclude List of Paths (privilege: advanced)

This specifies an array of directory names that need not be populated.



The path provided should be absolute path. If same path provided is for include and exclude path list then include-path will take the preference

Examples

The following example triggers FlexCache volume populate:

```
cluster1::> flexcache populate start -cache-volume
vol_cache2_vs3_c2_vol_origin1_vs1_c1 -cache-vserver vs3_c2 -path-list
/dir1 -isRecursion false
(volume flexcache populate start)
[JobId 164]: FlexCache populate job queued.
```

```
cluster1::> flexcache prepopulate start -cache-volume
vol_cache2_vs3_c2_vol_origin1_vs1_c1 -cache-vserver vs3_c2 -path-list
/dir1
(volume flexcache prepopulate start)
[JobId 188]: FlexCache prepopulate job queued.
```

```
cluster1::> flexcache prepopulate start -cache-volume
vol_cache2_vs3_c2_vol_origin1_vs1_c1 -cache-vserver vs3_c2 -path-list
/dir1, 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".
```

volume flexgroup commands

volume flexgroup qtree-disable

Disable qtree support on a FlexGroup

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume flexgroup qtree-disable` command disables qtree support on a FlexGroup.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

This specifies the Vserver in which the FlexGroup is located.

-volume <volume name> - Volume Name (privilege: advanced)

This specifies the name of the FlexGroup.

Examples

The following example disables the qtree support on a FlexVol named "fg" in Vserver "vs0":

```
cluster::*> volume flexgroup qtree-disable -vserver vs0 -volume fg
```

volume inode-upgrade commands

volume inode-upgrade resume

Resume suspended inode upgrade

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `volume inode-upgrade resume` command resumes suspended inode upgrade process. The inode upgrade process may be suspended earlier due to performance reasons.

Parameters

-vserver <vserver name> - VServer Name (privilege: advanced)

This specifies the Vserver on which the volume is located.

-volume <volume name> - Volume Name (privilege: advanced)

This specifies the volume for which the inode upgrade process is to be resumed.

Examples

The following example resumes a volume upgrade process.

```
cluster1::> volume inode-upgrade resume -vserver vs0 -volume vol1
```

volume inode-upgrade show

Display Inode Upgrade Progress

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `volume inode-upgrade show` command displays information about volumes in the middle of the inode upgrade process. The command output depends on the parameter or parameters specified with the command. If no parameters are specified, the command displays the default fields about all volumes in the middle of the inode upgrade process. Default fields are `vserver`, `volume`, `aggregate`, `status`, `scan-percent`, `remaining-time`, `space-needed`, and `scanner-progress`.

Parameters

{ [-fields <fieldname>,...]

This specifies the fields that need to be displayed.

| [-instance] }

If this parameter is specified, the command displays information about all entries.

[-vserver <vserver name>] - Vserver (privilege: advanced)

If this parameter and the `-volume` parameter are specified, the command displays detailed information about the specified volume. If this parameter is specified by itself, the command displays information about

volumes on the specified Vserver.

[`-volume <volume name>`] - Volume (privilege: advanced)

If this parameter and the `-vserver` parameter are specified, the command displays detailed information about the specified volume. If this parameter is specified by itself, the command displays information about all volumes that match the specified name.

[`-node <nodename>`] - Node Name (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that are located on the specified storage system.

[`-vol-dsid <integer>`] - Volume DSID (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that match the specified data set ID.

[`-vol-uuid <UUID>`] - Volume UUID (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that match the specified UUID.

[`-volume-msid <integer>`] - Volume MSID (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that match the specified master data set ID.

[`-vserver-uuid <UUID>`] - Vserver UUID (privilege: advanced)

If this parameter is specified, the command displays information only about the volume on the Vserver that has the specified UUID.

[`-aggregate <aggregate name>`] - Aggregate Name (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that are located on the specified storage aggregate.

[`-aggregate-uuid <UUID>`] - Aggregate UUID (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that are located on the storage aggregate with the specified UUID.

[`-status {pending|scanning|suspended-initializing|suspended|cleanup-pending|cleanup|cleanup-done|suspended-aborting|suspended-removing|suspended-while-removing|suspended-ironing}`] - Upgrade Status (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that match the specified inode upgrade status.

[`-scan-percent <percent>`] - Upgrade Scan Percent Complete (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that match the specified inode upgrade progress percentage.

[`-space-needed {<integer>[KB|MB|GB|TB|PB]}`] - Space Needed to Complete Upgrade (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes where the space needed to complete the upgrade process match the specified size.

[`-remaining-time` <[<integer>d] [<integer>h] [<integer>m] [<integer>s]>] - Remaining Upgrade Time (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes where the remaining time to complete the inode upgrade process match the specified time.

[`-scanner-progress` <text>] - Scanner Progress (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes where the progress of the inode upgrade process match the input.

Examples

The following example displays information about all volumes in the middle of the inode upgrade process on the Vserver named vs0:

```
cluster1::> volume inode-upgrade show -vserver vs0
Vserver Volume Aggregate Status %Complete Time          Space   Inode
          Remaining Needed Progress
-----
vs0      voll1  aggr1  pending  0%      -      3.07MB Public : Inode
0 out of 3822
```

volume move commands

volume move abort

Stop a running volume move operation

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The "volume move abort" command sends an abort message to the volume move operation and returns immediately. The volume move operation might not abort immediately depending on the stage it is in. For example, if the volume move operation is in a cut-over or clean-up phase, the abort is ignored. You invoke the "volume move show" command to view the list of running volume move operations and monitor the progress of the abort operation. This command has the same behavior as the `job stop`-id` <job-id>` command where the job-id is the identifier of the volume move job.

Parameters

`-vserver` <vserver name> - Vserver Name

This specifies the Vserver on which the volume is located.

`-volume` <volume name> - Volume Name

This specifies the name of the volume being moved.

Examples

The following example aborts running volume move operation on volume `voll1`

```

cluster1::> volume move show
      Vserver   Volume   State   Move Phase Percent-Complete Time-To-
Complete
-----
-----
      vs0       vol1     alert   cutover_hard_deferred 0%   -
      vs0       vol2     failed  failed           -           -
2 entries were displayed.
cluster1::> volume move abort -vserver vs0 -volume vol1
cluster1::> volume move show -vserver vs0 -volume vol1 -fields completion-
status
      vserver volume completion-status
-----
vs0      vol1      "Volume move job stopped."

```

The following example shows command failed to abort on *vol2* as volume move operation is completed.

```

cluster1::> volume move show
      Vserver   Volume   State   Move Phase Percent-Complete Time-To-
Complete
-----
-----
      vs0       vol1     alert   cutover_hard_deferred 0%   -
      vs0       vol2     failed  failed           -           -
2 entries were displayed.
cluster1::> volume move abort -vserver vs0 -volume vol2
Error: command failed: There is no volume move operation running on the
specified volume.

```

Related Links

- [job stop](#)

volume move modify

Modify parameters for a running volume move operation

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `volume move modify` command modifies the parameters used by the volume move operation in progress. These modified values can be verified by invoking the `volume move show` command. The volume move operation will use the modified cutover parameters in its next cutover attempt. Note that the modifications to the job are not applied if the move is in the "finishing" state. This command is not supported.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This specifies the name of the volume being moved.

[-cutover-action {abort_on_failure|defer_on_failure|force|wait|retry_on_failure}] - Specified Action For Cutover (privilege: advanced)

Specifies the action to be taken for cutover. If the effective cluster version is Data ONTAP 8.3 and later, the default is *retry_on_failure*; otherwise the default is *defer_on_failure*. If the *abort_on_failure* action is specified, the job will try to cutover until cutover attempts are exhausted. If it fails to cutover, it will cleanup and end the operation. If the *defer_on_failure* action is specified, the job will try to cutover until the cutover attempts are exhausted. If it fails to cutover, it will move into the "cutover_hard_deferred" state. The volume move job waits for a [volume move trigger-cutover](#) command to restart the cutover process. If the *force* action is specified, the job will try to cutover until the cutover attempts are exhausted and force the cutover at the expense of disrupting the clients. If the *wait* action is specified, when the job hits the decision point, it will not go into cutover automatically, instead it will wait for a [volume move trigger-cutover](#) command as the signal to try the cutover.

[-cutover-window <integer>] - Specified Cutover Time Window (privilege: advanced)

This specifies the time interval in seconds to completely cutover operations from the original volume to the moved volume. The default value is 30 seconds. The range for valid input is from 30 to 300 seconds, inclusive.

[-enforce-network-throttling {true|false}] - Enforce Replication Engine Node Level Network Throttling (privilege: advanced)

This specifies whether this move should enforce replication engine node level network throttling. If this parameter is specified with *true*, the move will be throttled by the replication engine. The default is *false*.

Examples

The following example modifies the parameters for volume move operation on volume vol2.

```
cluster1::*> volume move show -vserver vs0 -volume vol2
Vserver Name: vs0
          Volume Name: vol2
Actual Completion Time: -
          Bytes Remaining: 172KB
Specified Action For Cutover: wait
Specified Cutover Time Window: 40
Time Cutover was Triggered: -
Time Cutover was last triggered: -
          Destination Aggregate: node_1_aggr1
          Destination Node: node_1
          Detailed Status: Cutover Deferred, Waiting for user
intervention(69.79MB Sent)::Volume move job preparing transfer
          Estimated Time of Completion: -
```

```

                Job ID: 105
                Managing Node: node-2
                Percentage Complete: 50%
                Move Phase: cutover_hard_deferred
Estimated Remaining Duration: -
                Replication Throughput: -
                Duration of Move: 1 days 00:04
                Source Aggregate: node_2_aggr1
                Source Node: node_2
                Start Time of Move: Sun Sep 18 16:40:37 2011
                Move State: alert

cluster1::*> volume move modify -vserver vs0 -volume vol2 -cutover-action
abort_on_failure -cutover-window 50

cluster1::*> volume move show -vserver vs0 -volume vol2
Vserver Name: vs0
                Volume Name: vol2
                Actual Completion Time: -
                Bytes Remaining: 172KB
                Specified Action For Cutover: abort_on_failure
                Specified Cutover Time Window: 50
                Time Cutover was Triggered: -
                Time Cutover was last triggered: -
                Destination Aggregate: node_1_aggr1
                Destination Node: node_1
                Detailed Status: Cutover Deferred, Waiting for user
intervention(69.79MB Sent)::Volume move job preparing transfer
                Estimated Time of Completion: -
                Job ID: 106
                Managing Node: node-2
                Percentage Complete: 50%
                Move Phase: cutover_hard_deferred
Estimated Remaining Duration: -
                Replication Throughput: -
                Duration of Move: 1 days 00:05
                Source Aggregate: node_2_aggr1
                Source Node: node_2
                Start Time of Move: Sun Sep 18 16:40:37 2011
                Move State: alert

```

The following example shows command failed to modify on vol1 as volume move operation is completed.


```

cluster1::*> volume move show -vserver vs0 -volume voll
Vserver Name: vs0
                Volume Name: voll
                Actual Completion Time: Sun Sep 18 16:34:27 2011
                Bytes Remaining: 172KB
                Specified Action For Cutover: wait
                Specified Cutover Time Window: 30
                Time Cutover was Triggered: -
                Time Cutover was last triggered: -
                Destination Aggregate: node_1_aggr1
                Destination Node: node_1
                Detailed Status: Volume move failed because of a job
restart
                Estimated Time of Completion: -
                Job ID: 108
                Managing Node: node-2
                Percentage Complete: -
                Move Phase: failed
                Estimated Remaining Duration: -
                Replication Throughput: -
                Duration of Move: 15 days 08:07
                Source Aggregate: node_2_aggr1
                Source Node: node_2
                Start Time of Move: Sat Sep 03 08:27:06 2011
                Move State: failed

cluster1::*> volume move modify -vserver vs0 -volume voll -cutover-action
abort_on_failure -cutover-window 40

Error: command failed: There is no volume move operation running on the
specified volume.

```

Related Links

- [volume move show](#)
- [volume move trigger-cutover](#)

volume move show

Show status of a volume moving from one aggregate to another aggregate

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `volume move show` command displays information about volume moves in the cluster. By default, with

no parameters, it only shows volume moves that have failed or are currently running. The command display output depends on the parameters passed. If `-vserver` and `-volume` are specified, the following information is displayed:

- **Vserver Name:** The Vserver on which the volume is located.
- **Volume Name:** The volume that is part of a completed or running volume move operation.
- **Actual Completion Time:** The date and time in the cluster time zone when the volume move completed.
- **Bytes Remaining:** The number of bytes remaining to be sent during volume move. This is an approximate number and lags the current status by a few minutes while the volume move is in operation.
- **Specified Action for Cutover:** The action to be taken for cutover or during cutover failure. This is the input given during the start of volume move or the value specified during a volume move modification.
- **Specified Cutover Time Window:** The time window in seconds given as an input for the cutover phase of volume move. This is the input given during the start of volume move or the value specified during a volume move modification.
- **Job ID:** The Job-ID of move job.
- **Destination Node:** The name of the node where the destination aggregate is present.
- **Source Node:** The name of the node where the source aggregate is present.
- **Prior Issues Encountered:** The latest issues or transient errors encountered causing the move operation to retry the data copy phase or the cutover phase.
- **Move Initiated by Auto Balance Aggregate:** The value "true" indicates the move is initiated by Auto Balance Aggregate feature.
- **Destination Aggregate:** The name of the aggregate to which the volume is moved.
- **Detailed Status:** The detail about any warnings, errors, and state of the move operation.
- **Estimated Time of Completion:** The approximate date and time in the cluster time zone when the entire volume move operation is expected to complete. Note that this time may keep increasing when the move goes into cutover-deferred mode. In those cases where the input for cutover-action is wait, during the data copy phase, the estimated time of completion will approximate the time to reach the cutover point and wait for user intervention.
- **Managing Node:** The node in the cluster on which the move job is or was running. This is usually on the node hosting the volume to be moved.
- **Percentage Complete:** The amount of work to move the volume completed thus far in terms of percentage.
- **Move Phase:** The phase of the move operation.
- **Estimated Remaining Duration:** The approximate amount of time in terms of days, hours, minutes and seconds remaining to complete the volume move.
- **Replication Throughput:** The current replication throughput of the move operation in terms of Kb/s, Mb/s or Gb/s.
- **Duration of Move:** The duration in days, hours and minutes for which the volume move was or is in progress.
- **Source Aggregate:** The name of the aggregate where the volume being moved originally resides or resided.
- **Start Time of Move:** The date and time in the cluster time zone when the volume move operation started.
- **Move State:** The state of the volume move at the time of issuing the command and the system gathering up the information about the move.
- **Original Job ID:** The job-id assigned when the job was first created. This value will only be populated when

the original job-id differs from the current job-id.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Name

This specifies the Vserver on which the volume is located. If this parameter and the `-volume` parameter are specified, the command displays detailed information about latest move performed on the specified volume. If this parameter is specified by itself, the command displays information about latest moves performed on volumes of the specified Vserver.

[-volume <volume name>] - Volume Name

This specifies the volume that is part of a completed or running volume move operation. If this parameter and the `-vserver` parameter are specified, the command displays detailed information about latest move performed on the specified volume. If this parameter is specified by itself, the command displays information about the latest move on all volumes matching the specified name.

[-actual-completion-time <Date>] - Actual Completion Time

If this parameter is specified, the command displays move operations that match the specified date and time in the cluster time zone when the volume move completed.

[-bytes-remaining {<integer>[KB|MB|GB|TB|PB]}] - Bytes Remaining

If this parameter is specified, the command displays move operations that match the specified number of bytes remaining to be sent during volume move.

**[-cutover-action {abort_on_failure|defer_on_failure|force|wait|retry_on_failure}]
- Specified Action For Cutover (privilege: advanced)**

If this parameter is specified, the command displays move operations that match the specified action to be taken for cutover or during cutover failure.

[-cutover-window <integer>] - Specified Cutover Time Window (privilege: advanced)

If this parameter is specified, the command displays move operations that match the specified time window in seconds for the cutover phase of volume move.

[-destination-aggregate <aggregate name>] - Destination Aggregate

If this parameter is specified, the command displays move operations that match the specified name of the aggregate to which the volume is being moved.

[-destination-node <nodename>] - Destination Node (privilege: advanced)

If this parameter is specified, the command displays move operations that match the specified name of the node where the destination aggregate is present.

[-details <text>] - Detailed Status

If this parameter is specified, the command displays move operations that match the specified detail about any warnings, errors and state of the move operation.

[-estimated-completion-time <Date>] - Estimated Time of Completion

If this parameter is specified, the command displays move operations that match the specified date and time in the cluster time zone when the entire volume move operation is expected to complete.

[-job-id <integer>] - Job ID (privilege: advanced)

If this parameter is specified, the command displays move operations that match the specified Job-ID of the move job.

[-managing-node <nodename>] - Managing Node

If this parameter is specified, the command displays move operations that match the specified node in the cluster on which the move job is or was running.

[-percent-complete <percent>] - Percentage Complete

If this parameter is specified, the command displays move operations that match the specified amount of work to move the volume completed thus far in terms of percentage.

[-phase

{queued|initializing|replicating|cutover|cutover_hard_deferred|cutover_soft_deferred|aborting|completed|cleaning_up|failed|restarting|finishing}] - Move Phase

If this parameter is specified, the command displays move operations that match the specified phase of the move operation.

[-prior-issues <text>] - Prior Issues Encountered (privilege: advanced)

If this parameter is specified, the command displays move operations that match the specified issues or transient errors encountered causing the move operation to retry the data copy phase or the cutover phase.

[-estimated-remaining-duration {<seconds>| [<d> days] <hh>:<mm>[:<ss>]}] - Estimated Remaining Duration

If this parameter is specified, the command displays move operations that match the specified time.

[-replication-throughput <text>] - Replication Throughput

If this parameter is specified, the command displays move operations that match the specified replication throughput of the move operation in terms of Kb/s, Mb/s or Gb/s.

[-actual-duration {<seconds>| [<d> days] <hh>:<mm>[:<ss>]}] - Duration of Move

If this parameter is specified, the command displays move operations that match the specified duration for which the volume move was or is in progress.

[-source-aggregate <aggregate name>] - Source Aggregate

If this parameter is specified, the command displays move operations that match the specified name of the aggregate where the volume being moved originally resides or resided.

[-source-node <nodename>] - Source Node (privilege: advanced)

If this parameter is specified, the command displays move operations that match the specified name of the node where the source aggregate is present.

[-start-time <Date>] - Start Time of Move

If this parameter is specified, the command displays move operations that match the specified date and time in the cluster time zone when the volume move operation started.

[-state {healthy|warning|alert|failed|done}] - Move State

If this parameter is specified, the command displays move operations that match the specified state of the volume move operation.

[-moved-by-autobalance {true|false}] - Move Initiated by Auto Balance Aggregate (privilege: advanced)

If this parameter is specified, the command displays move operations that match the specified value of this parameter.

[-original-job-id <integer>] - Original Job ID (privilege: advanced)

If this parameter is specified, the command displays move operations that match the specified value of this parameter.

[-is-source-encrypted {true|false}] - Is Source Volume Encrypted

If this parameter is specified, the command displays move operations that match the specified value of this parameter.

[-source-key-id <text>] - Encryption Key ID of Source Volume

If this parameter is specified, the command displays move operations that match the specified value of this parameter.

[-is-destination-encrypted {true|false}] - Is Destination Volume Encrypted

If this parameter is specified, the command displays move operations that match the specified value of this parameter.

[-destination-key-id <text>] - Encryption Key ID of Destination Volume

If this parameter is specified, the command displays move operations that match the specified value of this parameter.

[-enforce-network-throttling {true|false}] - Enforce Replication Engine Node Level Network Throttling (privilege: advanced)

If this parameter is specified, the command displays move operations that match the specified value of this parameter.

[-is-capacity-tier-optimized {true|false}] - Is Capacity Tier Move Optimized (privilege: advanced)

If this parameter is specified, the command displays move operations that match the specified value of this parameter.

Examples

The following example lists status of volume move operation for a volume vol2 on a Vserver vs0

```
cluster1::> volume move show -vserver vs0 -volume vol2
Vserver Name: vs0
                Volume Name: vol2
Actual Completion Time: -
                Bytes Remaining: 6.37GB
Destination Aggregate: cluster1_aggr2
                Detailed Status: Transferring data: 3.67GB sent.
Estimated Time of Completion: Sat Jul 16 20:25:50 2011
                Managing Node: node1
                Percentage Complete: 36%
                Move Phase: replicating
Estimated Remaining Duration: 00:01
                Replication Throughput: 61.08MB/s
                Duration of Move: 00:02
                Source Aggregate: cluster1_aggr1
                Start Time of Move: Sat Jul 16 20:22:01 2011
                Move State: healthy
```

The following example lists status of volume move operation for a volume vol2 on a Vserver vs0 in advanced mode

```

cluster1::*> volume move show -vserver vs0 -volume vol2
Vserver Name: vs0
                Volume Name: vol2
    Actual Completion Time: -
        Bytes Remaining: 156KB
    Specified Action For Cutover: wait
    Specified Cutover Time Window: 30
        Destination Aggregate: cluster1_aggr2
            Destination Node: node2
                Detailed Status: Cutover Deferred, Waiting for
user intervention (2.04MB Sent)::Volume move job preparing transfer.
        Estimated Time of Completion: -
            Job ID: 265
                Managing Node: node1
                    Percentage Complete: -
                        Move Phase: cutover_hard_deferred
        Prior Issues Encountered: -
    Estimated Remaining Duration: -
        Replication Throughput: -
            Duration of Move: 00:24:59
                Source Aggregate: cluster1_aggr1
                    Source Node: node1
                        Start Time of Move: Tue Mar 17 22:31:32 2011
                            Move State: alert
    Move Initiated by Auto Balance Aggregate: false
        Original Job ID: -

```

The following example lists status of running and failed volume move operations in the cluster.

```

cluster1:::> volume move show
                Vserver   Volume   State   Move Phase   Percent-
Complete Time-To-Complete
-----
vs0             s1       alert   cutover_hard_deferred
                                           98%
-
vs0             vol2     failed  failed       -
-
                2 entries were displayed.

```

The following example lists status of all the volume move operations in the cluster.

```

cluster1::> vol move show -phase *
              (volume move show)
              Vserver  Volume  State  Move Phase Percent-
Complete Time-To-Complete
-----
-----
              vs0     s1     alert  cutover_hard_deferred
              98%
-
              vs0     s2     done   completed  100%
-
              vs0     vol1   failed failed      -
-
              3 entries were displayed.

```

volume move start

Start moving a volume from one aggregate to another aggregate

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `volume move start` command moves a flexible volume from one storage aggregate to another. The destination aggregate can be located on the same node as the original aggregate or on a different node. The move occurs within the context of the same Vserver.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This specifies the volume that will be moved.

-destination-aggregate <aggregate name> - Destination Aggregate

This specifies the aggregate to which the volume will be moved.

[-cutover-window <integer>] - Cutover time window in seconds (privilege: advanced)

This specifies the time interval to completely cutover operations from the original volume to the moved volume. The default value is 30 seconds. The range for valid input is from 30 to 300 seconds, inclusive.

**[-cutover-action {abort_on_failure|defer_on_failure|force|wait|retry_on_failure}]
- Action for Cutover (privilege: advanced)**

Specifies the action to be taken for cutover. If the effective cluster version is Data ONTAP 8.3 and later, the default is `retry_on_failure`; otherwise the default is `defer_on_failure`. If the `abort_on_failure` action is specified, the job tries to cutover until cutover attempts are exhausted. If it fails to cutover, it cleans up and ends the operation. If the `defer_on_failure` action is specified, the job

tries to cutover until the cutover attempts are exhausted. If it fails to cutover, it moves into the "cutover deferred" state. The volume move job waits to issue a `volume move trigger-cutover` command to restart the cutover process. If the `force` action is specified, the job tries to cutover until the cutover attempts are exhausted and forces the cutover at the expense of disrupting the clients. If the `wait` action is specified, when the job hits the decision point, it does not go into cutover automatically, instead it waits to issue a `volume move trigger-cutover` command as the signal to try the cutover. Once cutover is manually triggered, the cutover action changes to `defer_on_failure`. If the `retry_on_failure` action is specified, the job retries to cutover indefinitely and it never enters a "hard-deferred" state. After exhausting cutover attempts, the move job waits one hour before trying to cutover again. Issue a `volume move trigger-cutover` command at any time to restart the cutover process.

`[-perform-validation-only <true>]` - Performs validation checks only

This is a boolean option allowing to perform pre-move validation checks for the intended volume. When set to `true`, the command only performs the checks without creating a move job. The default value is `false`.

`[-foreground {true|false}]` - Foreground Process

This specifies whether the volume move operation runs as a foreground process. The default setting is `false` (that is, the operation runs in the background). Note that using this parameter will not affect how long it takes for the operation to complete.

`[-encrypt-destination {true|false}]` - Encrypt Destination Volume

This specifies whether the move operation should result in creating an encrypted volume on the destination aggregate. When this option is set to `true`, the destination volume will be encrypted. When it is set to `false`, the destination volume will be a plain-text volume. When this parameter is not specified, then destination will be same as the source type.

`[-generate-destination-key {true|false}]` - Generate New Encryption Key for Destination Volume

This option is specified along with `-encrypt-destination`, a new key will be generated, and that new key will be used for encrypting the destination volume.

`[-tiering-policy <Tiering Policy>]` - Volume Tiering Policy

This optional parameter specifies the tiering policy to apply to the destination volume. Tiering policies decide whether the user data blocks of a volume in a FabricPool will be tiered to the cloud tier when they become cold. FabricPool combines Flash (performance tier) with an object store (cloud tier) into a single aggregate. The temperature of a volume block increases if it is accessed frequently and decreases when it is not.

The available tiering policies are:

- `snapshot-only` - Only the volume Snapshot copies not associated with the active file system are tiered to the cloud tier.
- `auto` - Both Snapshot copy data and active file system user data are tiered to the cloud tier.
- `none` - Volume blocks are not tiered to the cloud tier.
- `all` - Both Snapshot copy data and active file system user data are tiered to the cloud tier as soon as possible without waiting for a cooling period.

`[-allow-mixed-aggr-types {true|false}]` - Allow Mixing FabricPool and non-FabricPool

If set to `true`, moving a FlexGroup constituent from a FabricPool to a non-FabricPool, or vice versa, is allowed. The default value is `false`. This parameter is only supported for FlexGroup constituents.

[`-encrypt-with-aggr-key {true|false}`] - Encrypt Destination Volume with aggr key

This specifies whether the move operation should result in creating an encrypted volume with aggr key on the destination aggregate. When this option is set to `true`, the destination volume will be encrypted as NAE (NetApp Aggregate Encryption) volume.

[`-enforce-network-throttling {true|false}`] - Enforce Replication Engine Node Level Network Throttling (privilege: advanced)

This specifies whether this move should enforce replication engine node level network throttling. If this parameter is specified with `true`, the move will be throttled by the replication engine. The default is `false`.

[`-is-capacity-tier-optimized {true|false}`] - Capacity Tier Optimized Volume Move (privilege: advanced)

This specifies whether this move should perform capacity tier optimized volume move, in which data in the capacity tier does not need to be copied over. If this parameter is specified as `true` and optimized volume move is not supported, the volume move operation will fail. If this parameter is specified as `false`, non-optimized volume move will be performed. When this parameter is not specified, optimized volume move is attempted, and if it is not supported, non-optimized volume move will be performed automatically.

Examples

The following examples perform a validation-check for a volume named `volume_test` on a Vserver named `vs0` to determine if it can be moved to a destination-aggregate named `dest_aggr`.

```
cluster1::> volume move start -vserver vs0 -volume volume_test
-destination-aggregate dest_aggr -perform-validation-only true
      Error: command failed: There is 2.54GB of available space on the
aggregate
      dest_aggr which is not enough to accommodate a volume.
cluster1::> volume move start -vserver vs0 -volume volume_test
-destination-aggregate dest_aggr -perform-validation-only true
      Validation succeeded.
```

The following example performs a volume move start operation to move a volume named `volume_test` on a Vserver name `vs0` to a destination-aggregate named `dest_aggr`.

```
cluster1::> volume move start -vserver vs0 -volume volume_test
-destination-aggregate dest_aggr
      [Job 267] Job is queued: Move "volume_test" in Vserver
"vs0" to aggregate "dest_aggr".
      Use the "volume move show -vserver vs0 -volume
volume_test" command to view the status of this operation.
```

The following example performs a volume move start operation to move a plain-text volume named `vol1` to an encrypted volume on destination-aggregate `aggr1`.

```
cluster1::> volume move start -volume voll -destination-aggregate aggr1
-encrypt-destination true
    [Job 267] Job is queued: Move "voll" in Vserver "vs1" to aggregate
"aggr1".
        Use the "volume move show -vserver vs1 -volume voll" command to
view the status of this operation.
```

Related Links

- [volume move trigger-cutover](#)

volume move trigger-cutover

Trigger cutover of a move job

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

This command causes a replicating or deferred volume move job to attempt cutover. Unless the force option is set, cutover entry is not guaranteed.

Parameters

-vserver <vserver name> - Vserver Name

The Vserver on which the volume is located.

-volume <volume name> - Volume Name

The volume that is being moved.

[-force <true>] - Force Cutover

If this parameter is specified, the cutover is done without confirming the operation - even if the operation could cause client I/O disruptions.

Examples

```
cluster1::> volume move trigger-cutover -vserver vs0 -volume testvol_1 -force
```

volume move recommend show

Display Move Recommendations

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `volume move recommend show` command displays moves that were recommended by the Auto Balance Aggregate feature.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Name (privilege: advanced)

If this parameter is specified, the display will be limited to only those recommendations with a Vserver that matches the specified value.

[-volume <volume name>] - Volume Name (privilege: advanced)

If this parameter is specified, the display will be limited to only those recommendations with a volume that matches the specified value.

[-creation-time <MM/DD/YYYY HH:MM:SS>] - Time Stamp of Recommendation (privilege: advanced)

If this parameter is specified, the display will be limited to only those recommendations with a creation-time that matches the specified value.

[-source-aggregate <aggregate name>] - Unbalanced Aggregate Name (privilege: advanced)

If this parameter is specified, the display will be limited to only those recommendations with a source-aggregate that matches the specified value.

[-source-space-after <percent>] - Space Free After Move (%) (privilege: advanced)

If this parameter is specified, the display will be limited to only those recommendations with a source-space-after that matches the specified value.

[-destination-aggregate <aggregate name>] - Destination Aggregate Name (privilege: advanced)

If this parameter is specified, the display will be limited to only those recommendations with a destination-aggregate that matches the specified value.

[-destination-space-after <percent>] - Space Bump After Move (%) (privilege: advanced)

If this parameter is specified, the display will be limited to only those recommendations with a destination-space-after that matches the specified value.

Examples

The following example displays information about the recommendations made by the Auto Balance Aggregate feature.

```
cluster1::*> volume move recommend show -instance
                Vserver Name: vs0.example.com
                Volume Name: ro10
Time Stamp of Recommendation: 3/13/2014 16:26:39
    Unbalanced Aggregate Name: aggr_1
    Space Free After Move (%): 36%
    Destination Aggregate Name: aggr_3
    Space Bump After Move (%): 36%
```

volume move target-aggr show

List target aggregates compatible for volume move

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The volume move target-aggr show displays information about compatible target aggregates for the specified volume to be moved to.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Name (Required field)

Selects information about compatible target aggregates for volumes of the specified Vserver.

[-volume <volume name>] - Volume Name (Required field)

Selects information about compatible target aggregates that have enough space for the specified volume.

[-aggregate <aggregate name>] - Aggregate Name

Selects information about compatible target aggregates with the specified aggregate name (to which the volume might be moved).

[-tiering-policy <Tiering Policy>] - Volume Tiering Policy

Selects information about compatible target aggregates with the specified destination tiering policy.

[-availsize {<integer>[KB|MB|GB|TB|PB] }] - Available size

Selects information about compatible target aggregates that have the specified available size.

[`-storagetype <text>`] - Storage Type

Selects information about compatible target aggregates with the specified storage type. Examples of storage types are “ATA”, “BSAS”, “FCAL”, “LUN”, “SATA”, “SAS” and “SSD”.

[`-allow-mixed-aggr-types {true|false}`] - Allow Mixing FabricPool and non-FabricPool

If set to true, moving a FlexGroup constituent from a FabricPool to a non-FabricPool, or vice versa, is allowed. The default value is false. This parameter is only supported for FlexGroup constituents.

Examples

The following example lists target aggregates compatible for moving a volume vol1 on a Vserver vs1.

```
cluster1::> volume move target-aggr show -vserver vs1 -volume vol1
Aggregate Name      Available Size      Storage Type
-----
aggr1                113.5GB             FCAL
aggr2                113.5GB             FCAL
2 entries were displayed.
```

volume object-store commands

volume object-store tiering show

Display Tiering Status of FabricPool Volumes

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `volume object-store tiering show` command displays information about the tiering scanner status of FabricPool aggregate volumes. The tiering scanner will, depending on the combination of the tiering-policy and cloud-retrieval-policy settings, both retrieve from and tier to the capacity tier. This show command can be used to display the status of the scanner, whether it has aborted, what errors it has encountered and when it will be scheduled to run again.

Parameters

{ [`-fields <fieldname>`,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [`-instance]` }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[`-vserver <vserver name>`] - Vserver (privilege: advanced)

vserver name

[`-volume <volume name>`] - Volume (privilege: advanced)

volume name

[`-node <nodename>`] - Node Name (privilege: advanced)

node name

[`-vol-dsid <integer>`] - Volume DSID (privilege: advanced)

volume dsid

[`-aggregate <aggregate name>`] - Aggregate Name (privilege: advanced)

aggregate name

[`-scanner-state {ready|active|waiting}`] - State (privilege: advanced)

The state of the scanner. The values active, ready and waiting may appear in this field. Active means the scanner is currently running, waiting means the scanner is waiting to be scheduled, and ready means the scanner is ready to run. Once a scanner is ready to run, it must be scheduled by the system's scan scheduler, the amount of time it takes for it to become active is variable and depends on how busy the system is..

[`-scanner-last-status {aborted-nospace-promote|aborted-fabricpool|aborted-nospace-defrag|aborted-nospace|aborted-exception|aborted-policy|aborted|completed|prev-not-run|aborted-defrag-throttled|aborted-nospace-revert-format|aborted-revert-format-paused|aborted-dual-tier-dirty-failed}`] - Previous Run Status (privilege: advanced)

The status of the tiering scanner the last time it ran. This field will indicate if the scanner has aborted or completed. Possible status options are:

- `aborted-nospace-promote` : The tiering scanner was aborted because of an internal space check error or a lack of space while promoting data from the capacity tier to the performance tier. Please check the `promotion-space-needed` field.
- `aborted-fabricpool`: The tiering scanner was aborted because it tried running inactive data reporting on a FabricPool volume.
- `aborted-nospace-defrag` : The tiering scanner was aborted because of an internal space check error while defragmenting.
- `aborted-nospace` : The tiering scanner was aborted because of an internal space check error.
- `aborted-exception` : The tiering scanner encountered an exception while reading data from the capacity tier. Check the `scanner-abort-exception` field to see the error code that caused the exception.
- `aborted-policy` : The tiering scanner was aborted because either the `tiering-policy` or the `cloud-retrieval-policy` changed while the scan was in progress.
- `aborted` : The tiering scan was manually aborted.
- `completed` : The tiering scanner was completed successfully.
- `prev-not-run` : The tiering scanner has not run before.
- `aborted-defrag-throttled` : The tiering scanner was aborted because the defragmentation cannot continue due to its threshold limit.
- `aborted-nospace-revert-format` : The tiering scanner was aborted because of an internal space check error while reverting unified object format to native format.
- `aborted-revert-format-paused`: The tiering scanner was aborted because it failed to drain waiting objects

while reverting unified object format to native format.

- aborted-dual-tier-dirty-failed: The tiering scanner was aborted because it failed to tier blocks while reverting unified object format to native format.

[-scanner-abort-exception <integer>] - Aborted Exception Status (privilege: advanced)

If the scanner last status indicates the scanner aborted with an exception then this field is set to the exception error code that caused the abort.

[-time-last-scan <Date>] - Time Scanner Last Finished (privilege: advanced)

This field displays the date and time the last scanner status field was set.

[-scan-percent <percent>] - Scanner Percent Complete (privilege: advanced)

This field displays the completion percentage of scans in the active state.

[-time-next-scan <Date>] - Time Waiting Scan will be scheduled (privilege: advanced)

This field displays the date and time scans in the waiting state will enter the ready state.

[-tiering-policy <Tiering Policy>] - Tiering Policy (privilege: advanced)

The tiering policy. This field displays the current tiering-policy in effect on the volume.

[-promotion-space-needed {<integer>[KB|MB|GB|TB|PB]}] - Estimated Space Needed for Promotion (privilege: advanced)

If the scanner aborted due to no space for promote, then this field is set to the estimated minimum space required for promotion to take effect.

[-scan-start-time <Date>] - Time Scan Started (privilege: advanced)

This field displays the date and time the scan, if in active state, started.

[-cloud-retrieval-policy {default|on-read|never|promote}] - Cloud Retrieval Policy (privilege: advanced)

The cloud retrieval policy. This field displays the current cloud-retrieval-policy in effect on the volume.

[-scan-elapsed-time [<integer>d [<integer>h [<integer>m [<integer>s]]] - Elapsed Time Scanner Ran (privilege: advanced)

This field displays the time it took for the last scan to run to complete or abort.

[-revert-scan-state <text>] - Object Format Revert Scan State (privilege: advanced)

This field displays the state of object format revert scan. The values inactive, ready, active and complete may appear in this field. Inactive means revert scan is not running on this volume, ready means scan is waiting to start running, active mean scanner is running and complete indicates revert scan is complete on this volume.

[-scan-last-space-err <integer>] - space error code encountered in last scan (privilege: advanced)

This field displays the most recent space check error code, if any, that was encountered.

[-blocks-promoted <integer>] - Blocks Promoted (privilege: advanced)

This field displays the number of blocks that were promoted from the capacity tier to the performance tier.

Examples

```
cluster1::> volume object-store tiering show

                                     Policy
Vserver Volume Aggregate State   %Complete Tiering CloudRetrieve
LastStatus
-----
vs1      vol      aggr1      waiting      - snapshot-only
                                     default
completed
```

Shows the tiering status for all FabricPool volumes.

volume object-store tiering trigger

Trigger a tiering scan

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `volume object-store tiering trigger` command triggers a tiering and retrieve scan. Tiering and retrieval behavior is driven by the tiering and cloud retrieve policy settings on the volume. The cloud retrieve policy must be set to promote to enable scanner based retrieval.

Parameters

-vserver <vserver name> - VServer Name (privilege: advanced)

vserver name

-volume <volume name> - Volume Name (privilege: advanced)

volume name

volume qtree commands

volume qtree create

Create a new qtree

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command creates a qtree in the Vserver and volume you specify. You can create up to 4,994 qtrees per volume.

You can optionally specify the following attributes when creating a new qtree:

- Security style
- Opportunistic lock mode
- User ID
- Group ID
- UNIX permissions
- Export Policy

Parameters

-vserver <vserver name> - Vserver Name

This specifies the name of the Vserver on which the volume containing the qtree belongs.

{ -volume <volume name> - Volume Name

This specifies the name of the volume that will contain the qtree you are creating.

-qtree <qtree name> - Qtree Name

This specifies the name of the qtree you are creating.

A qtree name cannot contain a forward slash (/). The qtree name cannot be more than 64 characters long.

| -qtree-path <qtree path> - Actual (Non-Junction) Qtree Path }

The qtree path argument in the format `/vol/<volume name>/<qtree name>` can be specified instead of specifying volume and qtree as separate arguments.

[-security-style <security style>] - Security Style

This optionally specifies the security style for the qtree, which determines how access to the qtree is controlled. The supported values are *unix* (for UNIX uid, gid and mode bits), *ntfs* (for CIFS ACLs), and *mixed* (for NFS and CIFS access). If you do not specify a security style for the qtree, it inherits the security style of its containing volume.

[-oplock-mode {enable|disable}] - Oplock Mode

This optionally specifies whether oplocks are enabled for the qtree. If you do not specify a value for this parameter, it inherits the oplock mode of its containing volume.

[-user <user name>] - User ID

This optionally specifies the name or ID of the user that is set as the owner of the qtree.

[-group <group name>] - Group ID

This optionally specifies the name or ID of the group that is set as the owner of the qtree.

[-m, -unix-permissions <unix perm>] - Unix Permissions

This optionally specifies the UNIX permissions for the qtree when the `-security-style` is set to *unix* or *mixed*. You can specify UNIX permissions either as a four-digit octal value (for example, 0700) or in the style of the UNIX `ls` command (for example, `-rwxr-x---`). For information on UNIX permissions, see the UNIX or Linux documentation. If you do not specify UNIX permissions for the qtree, it inherits the UNIX permissions of its containing volume.

[`-export-policy <text>`] - Export Policy

This optional parameter specifies the name of the export policy associated with the qtree. For information on export policies, see the documentation for the [vserver export-policy create](#) command. If you do not specify a value for this parameter, it inherits the export policy of its containing volume.

[`-qos-policy-group <text>`] - QoS policy group

A policy group defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated. If you do not assign a policy group to a qtree, the system does not monitor and control the traffic to it.

Examples

The following example creates a qtree named qtree1. The Vserver name is vs0 and the volume containing the qtree is named vol1. The qtree has a mixed security style. Its other attributes are inherited from volume vol1.

```
cluster1::> volume qtree create -vserver vs0 -volume vol1 -qtree qtree1
-security-style mixed
```

The following example uses a 7G-compatible command to create the qtree.

```
cluster1::> vserver context vs0
vs0::> qtree create /vol/vol1/qtree1
```

Related Links

- [vserver export-policy create](#)

volume qtree delete

Delete a qtree

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command deletes a qtree. The length of time that it takes to delete a qtree depends on the number of directories and files it contains. You can monitor the progress of the delete operation by using the [job show](#) and [job watch-progress](#) commands, respectively.

The automatically created qtree in the volume - qtree0, listed in CLI output as "" - cannot be deleted.



Quota rules associated with this qtree in all the quota policies will be deleted when you delete this qtree. Qtree deletion will not be allowed if Storage-level Access Guard (SLAG) is configured.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the name of the Vserver on which the volume containing the qtree belongs.

{ -volume <volume name> - Volume Name

This specifies the name of the volume containing the qtree to be deleted.

-qtree <qtree name> - Qtree Name

This specifies the name of the qtree to be deleted.

| -qtree-path <qtree path> - Actual (Non-Junction) Qtree Path }

The qtree path argument in the format /vol/<volume name>/<qtree name> can be specified instead of specifying volume and qtree as separate arguments.

[-force <true>] - Force Delete (privilege: advanced)

This optionally forces the qtree delete operation to proceed when the qtree contains files. The default setting is false (that is, the qtree will not be deleted if it contains files). This parameter is available only at the advanced privilege and higher.

[-foreground <true>] - Foreground Process

This optionally specifies whether the qtree delete operation runs as a foreground process. The default setting is false (that is, the operation runs in the background).

Examples

The following example deletes a qtree named qtree4. The Vserver name is vs0 and the volume containing the qtree is named vol1.

```
cluster1::> volume qtree delete -vserver vs0 -volume vol1 -qtree qtree4
WARNING: Are you sure you want to delete qtree qtree4 in volume vol1
vserver vs0? {y|n}: y
[Job 38] Job is queued: Delete qtree qtree4 in volume vol1 vserver vs0.
```

Related Links

- [job show](#)
- [job watch-progress](#)

volume qtree modify**Modify qtree attributes**

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command allows you to modify the following attributes of an existing qtree in the given Vserver and volume:

- Security style

- Opportunistic lock mode
- User ID
- Group ID
- UNIX permissions
- Export policy

Parameters

-vserver <vserver name> - Vserver Name

This specifies the name of the Vserver on which the volume containing the qtree belongs.

{ -volume <volume name> - Volume Name

This specifies the name of the volume containing the qtree to be modified.

-qtree <qtree name> - Qtree Name

This specifies the name of the qtree to be modified. You can modify the attributes of qtree0 (represented as "" in the CLI) by omitting the `-qtree` parameter from the command or by specifying the value "" for the `-qtree` parameter.

| -qtree-path <qtree path> - Actual (Non-Junction) Qtree Path }

The qtree path argument in the format `/vol/<volume name>/<qtree name>` can be specified instead of specifying volume and qtree as separate arguments. The automatically created qtree0 can be represented as `/vol/<volume name>`.

[-security-style <security style>] - Security Style

This optionally modifies the security style for the qtree. The supported values are *unix* (for UNIX uid, gid and mode bits), *ntfs* (for CIFS ACLs), and *mixed* (for NFS and CIFS access). Modifying a qtree's security style will not affect any of the files in the other qtrees of this volume.

[-oplock-mode {enable|disable}] - Oplock Mode

This optionally modifies whether oplocks are enabled for the qtree.

Modifying qtree0's oplock mode will not affect any of the files in the other qtrees of this volume.

[-user <user name>] - User ID

This optionally modifies the name or ID of the user that is set as the owner of the qtree.

[-group <group name>] - Group ID

This optionally modifies the name or ID of the group that is set as the owner of the qtree.

[-unix-permissions <unix perm>] - Unix Permissions

This optionally modifies the UNIX permissions for the qtree. You can specify UNIX permissions either as a four-digit octal value (for example, 0700) or in the style of the UNIX `ls` command (for example, `-rwxr-x---`). For information on UNIX permissions, see the UNIX or Linux documentation.

The unix permissions can be modified only for qtrees with unix or mixed security style.

[-export-policy <text>] - Export Policy

This optional parameter modifies the export policy associated with the qtree. If you do not specify an export

policy name, the qtree inherits the export policy of the containing volume. For information on export policy, see the documentation for the [vserver export-policy create](#) command.

[`-qos-policy-group <text>`] - QoS Policy Group

This optional parameter specifies which QoS policy group to apply to the qtree. This policy group defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy group is associated. If you do not assign a policy group to a qtree, the system does not monitor and control the traffic to it. To remove this qtree from a policy group, enter the reserved keyword "none".

Examples

The following example modifies a qtree named qtree1. The Vserver name is vs0 and the volume containing the qtree is named vol1. The qtree now has a UNIX security style and oplocks are enabled.

```
cluster1::> volume qtree modify -vserver vs0 -volume vol1 -qtree qtrees1
-security-style unix -oplocks enabled
```

Related Links

- [vserver export-policy create](#)

volume qtree oplocks

Modify qtree oplock mode

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command allows you to display or modify the opportunistic lock mode of a qtree.

Parameters

`-vserver <vserver name>` - Vserver Name

This specifies the name of the Vserver on which the volume containing the qtree belongs.

{ `-volume <volume name>` - Volume Name

This specifies the name of the volume containing the qtree.

`-qtree <qtree name>` - Qtree Name

This specifies the name of the qtree for which the oplock mode is being displayed or modified.

| `-qtree-path <qtree path>` - Actual (Non-Junction) Qtree Path }

The qtree path argument in the format `/vol/<volume name>/<qtree name>` can be specified instead of specifying volume and qtree as separate arguments. The automatically created qtree0 can be represented as `/vol/<volume name>`.

[`-oplock-mode {enable|disable}`] - Oplock Mode

This specifies the new oplock mode of the qtree. If this parameter is not specified, then the current oplock mode of the qtree is displayed.

Modifying qtree0's oplock mode will not affect any of the files in the other qtrees of this volume.

Examples

The following example displays the oplock mode of a qtree called qtree1. The Vserver name is vs0 and the volume containing the qtree is named vol1.

```
cluster1::> volume qtree oplocks -vserver vs0 -volume vol1 -qtree qtree1
/vol/vol1/qtree1 has mixed security style and oplocks are disabled.
```

The following example modifies the oplock mode of a qtree called qtree2 to enabled. The Vserver name is vs0 and the volume containing the qtree is named vol1.

```
cluster1::> volume qtree oplocks -vserver vs0 -volume vol1 -qtree qtree2
-oplock-mode enable
```

The following example uses a 7G-compatible command to display and modify the oplock mode of a qtree.

```
cluster1::> vserver context vs0
vs0::> qtree oplocks /vol/vol1/qtree1
/vol/vol1/qtree1 has mixed security style and oplocks are disabled.
vs0::> qtree oplocks /vol/vol1/qtree2 enable
```

volume qtree rename

Rename an existing qtree

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command allows you to rename an existing qtree.

The automatically created qtree in the volume - qtree0, listed in CLI output as "" - cannot be renamed.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the name of the Vserver on which the volume containing the qtree belongs.

{ -volume <volume name> - Volume Name

This specifies the name of the volume containing the qtree to be renamed.

-qtree <qtree name> - Qtree Name

This specifies the name of the qtree to be renamed.

| -qtree-path <qtree path> - Actual (Non-Junction) Qtree Path }

The qtree path argument in the format `/vol/<volume name>/<qtree name>` can be specified instead of specifying volume and qtree as separate arguments.

[-newname <qtree name>] - Qtree New Name

This specifies the new name of the qtree. The new qtree name cannot contain a forward slash (/) and cannot be more than 64 characters long.

Examples

The following example renames a qtree named qtree3 to qtree4. The Vserver name is vs0 and the volume containing the qtree is named vol1.

```
cluster1::> volume qtree rename -vserver vs0 -volume vol1 -qtree qtree3
-newname qtree4
```

volume qtree security**Modify qtree security style**

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command allows you to display or modify the security style of a qtree.

Parameters**-vserver <vserver name> - Vserver Name**

This specifies the name of the Vserver on which the volume containing the qtree belongs.

{ -volume <volume name> - Volume Name

This specifies the name of the volume containing the qtree.

-qtree <qtree name> - Qtree Name

This specifies the name of the qtree for which the security style is being displayed or modified.

| -qtree-path <qtree path> - Actual (Non-Junction) Qtree Path }

The qtree path argument in the format `/vol/<volume name>/<qtree name>` can be specified instead of specifying volume and qtree as separate arguments. The automatically created qtree0 can be represented as `/vol/<volume name>`.

[-security-style <security style>] - Security Style

This specifies the new security style of the qtree. If this parameter is not specified, then the current security style of the qtree is displayed. The supported values are *unix* (for UNIX uid, gid and mode bits), *ntfs* (for CIFS ACLs), and *mixed* (for NFS and CIFS access). Modifying a qtree's security style will not affect any of

the files in the other qtrees of this volume.

Examples

The following example displays the security style of a qtree called qtree1. The Vserver name is vs0 and the volume containing the qtree is named vol1.

```
cluster1::> volume qtree security -vserver vs0 -volume vol1 -qtree qtree1
/vol/vol1/qtree1 has mixed security style and oplocks are disabled.
```

The following example modifies the security style of a qtree called qtree2 to unix. The Vserver name is vs0 and the volume containing the qtree is named vol1.

```
cluster1::> volume qtree security -vserver vs0 -volume vol1 -qtree qtree2
-security-style unix
```

The following example uses a 7G-compatible command to display and modify the security style of a qtree.

```
cluster1::> vserver context vs0
vs0::> qtree security /vol/vol1/qtree1
/vol/vol1/qtree1 has mixed security style and oplocks are disabled.
vs0::> qtree security /vol/vol1/qtree2 unix
```

volume qtree show

Display a list of qtrees

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays information about qtrees for online volumes. By default, the command displays the following information about all qtrees in the cluster:

- Vserver name
- Volume name
- Qtree name
- Security style (unix, ntfs, mixed or unified)
- Whether oplocks is enabled
- Status (normal or readonly)

The display will also include information about Qtree 0. When you create a volume, a special qtree referred to

as "qtree0", also called the default qtree is automatically created for the volume. It represents all of the data stored in a volume that is not contained in a qtree. In the CLI output, qtree0 is denoted by empty quotation marks ("") and has the ID zero (0). The qtree called qtree0 cannot be manually created or deleted.

The qtree status indicates readonly for data protection and load sharing volumes.

To display detailed information about a single qtree, run the command with the `-instance` and `-qtree` parameters. The detailed view adds the following information:

- User ID
- Group ID
- UNIX permissions
- Qtree ID
- Export policy
- Is Export Policy Inherited

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-exports]

Displays the following information about qtree exports:

- Vserver - The name of the Vserver the qtree belongs to
- Volume - The name of the volume the qtree resides on
- Qtree name - The name of the qtree
- Policy Name - The name of the export policy assigned to the qtree
- Is Export Policy Inherited - Whether the export policy assigned to the qtree is inherited

| [-id]

Displays qtree IDs in addition to the default output.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Name

Selects information about the qtrees in the specified Vserver.

{ [-volume <volume name>] - Volume Name

Selects information about the qtrees in the specified volume.

[-qtree <qtree name>] - Qtree Name

Selects information about the qtrees that have the specified name.

[`-qtree-path <qtree path>`] - Actual (Non-Junction) Qtree Path }

Selects information about the qtrees that have the specified path.

[`-security-style <security style>`] - Security Style

Selects information about the qtrees that have the specified security style.

[`-oplock-mode {enable|disable}`] - Oplock Mode

Selects information about the qtrees that have the specified oplock mode.

[`-user <user name>`] - User ID

Selects information about the qtrees that have the specified user.

[`-group <group name>`] - Group ID

Selects information about the qtrees that have the specified group.

[`-m, -unix-permissions <unix perm>`] - Unix Permissions

Selects information about the qtrees that have the specified UNIX permissions.

[`-qtree-id <integer>`] - Qtree Id

Selects information about the qtrees that have the specified ID. A valid qtree ID is an integer from 0 to 4994. All qtree0 (automatically created) qtrees have an ID of zero (0).

[`-status {normal|readonly}`] - Qtree Status

Selects information about the qtrees that have the specified status.

[`-export-policy <text>`] - Export Policy

Selects information about the qtrees that use the specified export policy.

[`-is-export-policy-inherited {true|false}`] - Is Export Policy Inherited

Selects information about the qtrees that inherit (true) or not inherit (false) the export policy of containing volume.

[`-qos-policy-group <text>`] - QoS policy group

Selects information about the qtrees that use the specified QoS policy.

Examples

The following example displays default information about all qtrees along with each qtree ID. Note that on vs0, no qtrees have been manually created, so only the automatically created qtrees referred to as qtree 0 are shown. On vs1, the volume named vs1_vol1 contains qtree 0 and two manually created qtrees, qtree1 and qtree2.

```

cluster1::> volume qtree show -id
Vserver      Volume          Qtree           Style           Oplocks         Status         Id
-----
vs0          vs0_vol1        ""              unix            enable          readonly      0
vs0          vs0_vol2        ""              unix            enable          normal        0
vs0          vs0_vol3        ""              unix            enable          readonly      0
vs0          vs0_vol4        ""              unix            enable          readonly      0
vs0          root_vs_vs0     ""              unix            enable          normal        0
vs1          vs1_vol1        ""              unix            enable          normal        0
vs1          vs1_vol1        qtree1          unix            disable         normal        1
vs1          vs1_vol1        qtree2          unix            enable          normal        2
vs1          root_vs_vs1     ""              unix            enable          normal        0
9 entries were displayed.

```

volume quota commands

volume quota modify

Modify quota state for volumes

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command allows you to modify the following quota attributes for one or more volumes:

- Quota state
- Whether quota exceeded messages are logged or not
- Frequency with which quota exceeded messages are logged

Modifications to the quota state for a volume creates a job to perform the quota state changes for that volume. You can monitor the progress of the job by using the [job show](#) and [job watch-progress](#) commands.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the name of the Vserver on which the volume whose quota attributes you are modifying is located.

-volume <volume name> - Volume Name

This specifies the name of the volume whose quota attributes you are modifying.

[-state <quota_state>] - Quota State

This parameter optionally modifies the quota state to one of the following:

- `off` - This indicates that quotas be deactivated for the specified volume.
- `on` - This indicates that quotas be activated for the specified volume.

- `resize` - This indicates that the quota limits be resized according to the values specified in the quota policy assigned to the Vserver. Note that quotas must be activated first for a volume before a resize operation can be performed.

Both quota activation and quota resize operations apply the quota rules configured for the volume within the quota policy that is currently assigned to the Vserver. These quota rules are managed by using the commands in the `volume quota policy rule` menu. Quotas, when activated for a volume, go through an initialization process. As part of the quota initialization all the quota rules are applied to the volume. In addition, a filesystem scanner is started to scan the entire filesystem within the volume to bring the quota accounting and reporting up to date. The quota job finishes after the filesystem scanner is started on the volume. The quota state for the volume is `initializing` until the filesystem scanner finishes scanning the entire filesystem. After the scanning is complete, the quota state will be `on`.

When quotas are resized, the quota state is `resizing` until the resizing operation finishes. As part of this operation, the quota limits for quotas currently in effect are resized to the limits currently configured for the volume. After the quota resize operation finishes, the quota state will be `on`.

Quota state changes can also be performed using the commands `volume quota on`, `volume quota off` and `volume quota resize`.

`[-logging {on|off}] - Logging Messages`

This parameter optionally specifies whether quota exceeded syslog/EMS messages are logged in the system log messages. When it is set to `on`, quota exceeded messages are generated when the user exceeds the quota's disk limit or the file limit through a NFS/CIFS operation or any operation within the Data ONTAP software. When set to `off` no quota exceeded messages are generated. This parameter can be changed only after quotas are activated for a volume.

`[-logging-interval <text>] - Logging Interval`

This parameter optionally specifies a logging interval, which indicates the frequency with which quota exceeded messages are generated. You can specify a logging interval in the `<integer><suffix>` format, where suffix can be minutes (`m`), hours (`h`), or days (`d`), but not combinations thereof (in other words, `90m` is a valid logging interval, but `1h30m` is not a valid logging interval). You can modify the logging interval only when the logging is `on`. When quotas are first activated, the logging parameter is automatically set to `on`, and the logging interval set to `1h`. If continuous logging is required, an interval of `0m` should be specified. This parameter can be changed only after quotas are activated for a volume.



quota message logging may not occur at exactly the same interval rate as specified by the user, especially for very small intervals. This is due to the behavior of the logging system that buffers messages instead of outputting them immediately. Setting the logging interval to `0m` can cause lots of quota exceeded messages to be logged in the system log messages.

`[-foreground <>true>] - Foreground Process`

This parameter optionally specifies whether the job created by quota state modify operation runs as a foreground process. The default setting is `false` (that is, the quota state modify operation runs in the background). When set to `true`, the command will not return until the job completes.

Examples

The following example activates quotas on the volume named `vol1`, which exists on Vserver `vs0`.

```
cluster1::> volume quota modify -vserver vs0 -volume voll -state on
[Job 24] Job is queued: Quota ON Operation on vserver vs0 volume voll.
```

The following example turns on quota message logging and sets the logging interval to 4 hours.

```
cluster1::> volume quota modify -vserver vs0 -volume voll -logging on
-logging-interval 4h
```

The following example resizes quota limits on a volume.

```
cluster1::> volume quota modify -vserver vs0 -volume voll -state resize
-foreground true
[Job 80] Job succeeded: Successful
```

Related Links

- [job show](#)
- [job watch-progress](#)
- [volume quota on](#)
- [volume quota off](#)
- [volume quota resize](#)

volume quota off

Turn off quotas for volumes

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command creates a job to deactivate quotas for the specified volume. You can monitor the progress of the job by using the [job show](#) and [job watch-progress](#) commands.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the name of the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This specifies the name of the volume on which you are deactivating quotas.

[-foreground <true>] - Foreground Process

This optionally specifies whether the job created for deactivating quotas runs as a foreground process. The default setting is *false* (that is, the operation runs in the background). When set to *true*, the command will not return until the job completes.

Examples

The following example deactivates quotas on the volume named `vol1` , which exists on Vserver `vs0` .

```
cluster1::> volume quota off -vserver vs0 -volume vol1
[Job 23] Job is queued: Quota OFF Operation on vserver vs0 volume vol1.
```

The following example uses a 7G-compatible command to deactivate quotas on the volume named `vol1` which exists on Vserver `vs0` .

```
cluster1::> vserver context vs0
vs0::> quota off vol1
[Job 25] Job is queued: Quota OFF Operation on vserver vs0 volume vol1.
```

Related Links

- [job show](#)
- [job watch-progress](#)

volume quota on

Turn on quotas for volumes

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command creates a job to activate quotas for the specified volume. You can monitor the progress of the job by using the [job show](#) and [job watch-progress](#) commands.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the name of the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This specifies the name of the volume on which you are activating quotas.

[-w, -foreground <true>] - Foreground Process

This optionally specifies whether the job created for activating quotas runs as a foreground process. The default setting is *false* (that is, the operation runs in the background). When set to *true* , the command will not return until the job completes. The quota job finishes after the filesystem scanner is started. The quota state for the volume is *initializing* until the filesystem scanner finishes scanning the entire filesystem. After the scanning is complete, the quota state will be *on* .

Examples

The following example activates quotas on the volume named `vol1` , which exists on Vserver `vs0` .

```
cluster1::> volume quota on -vserver vs0 -volume vol1
[Job 23] Job is queued: Quota ON Operation on vserver vs0 volume vol1.
```

The following example uses a 7G-compatible command to activate quotas on the volume named `vol1` which exists on Vserver `vs0` .

```
cluster1::> vserver context vs0
vs0::> quota on -w vol1
[Job 25] Job is queued: Quota ON Operation on vserver vs0 volume vol1.

[Job 25] Job succeeded: Successful
```

Related Links

- [job show](#)
- [job watch-progress](#)

volume quota report

Display the quota report for volumes

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays the quota report for all volumes in each Vserver that are online and for which quotas are activated. Quota report includes the quota rules (default, explicit, and derived) in effect and the associated resource usage (disk space and files). If quotas are still initializing for a specific volume, that volume is not included.

This command displays the following information:

- Vserver name
- Volume name
- Index - This is a unique number within a volume assigned to each quota rule displayed in the quota report.
- Tree name - This field gives the name of the qtree if the quota rule is at the qtree level. It is empty if the quota rule is at the volume level.
- Quota type - Type of quota rule (`tree` or `user` or `group`).
- Quota target - This field gives the name of the target of the quota rule. For tree quota rules, it will be the qtree ID of the qtree. For user quota rules, it will be the UNIX user name or the Windows user name. For group quota rules, it will be the UNIX group name. For default rules (`tree` or `user` or `group`), this will display "*" . If the UNIX user identifier, UNIX group identifier, or Windows security identifier no longer exists or if the identifier-to-name conversion fails, the target appears in numeric form.
- Quota target ID - This field gives the target of the quota rule in numeric form. For tree quota rules, it will be the qtree ID of the qtree. For group quota rules, it will be the UNIX group identifier. For UNIX user quota rules, it will be the UNIX user identifier. For Windows user quota rules, it will be the Windows security

identifier in its native format. For default rules (tree or user or group), "*" will be displayed.

- Disk space used - For a default quota, the value is 0.
- Disk space limit
- Soft disk space limit
- Threshold for disk space limit
- Files used - For a default quota, the value is 0.
- File limit
- Soft file limit
- Quota specifier - For an explicit quota, this field shows how the quota target was configured by the administrator using the volume quota policy rule command. For a default quota, the field shows "" . For a derived tree quota, this field shows the qtree path. For a derived user and group quota, the field is either blank or "" .

The following parameters: `-soft` , `-soft-limit-thresholds` , `-target-id` , `-thresholds` , `-fields` and `-instance` display different set of fields listed above. For example, `-soft` will display the soft disk space limit and soft file limit apart from other information. Similarly `-target-id` will display the target in the numeric form.

A quota report is a resource intensive operation. If you run it on many 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 a Vserver.

Depending upon the quota rules configured for a volume, the quota report for a single volume can be large. If you want to monitor the quota report entry for a particular tree/user/group repeatedly, find the index of that quota report entry and use the `-index` field to view only that quota report entry. See the examples section for an illustration.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-soft]

If this parameter is specified, the command display will include the soft disk space limit and the soft file limit.

| [-soft-limit-thresholds]

If this parameter is specified, the command display will include the soft disk space limit, threshold for disk space limit and soft file limit.

| [-target-id]

If this parameter is specified, the command will display the target of a user or group quota rule in numeric form.

| [-thresholds]

If this parameter is specified, the command display will include the threshold for disk space limit.

[`-instance`] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[`-vserver <vserver name>`] - Vserver Name

If this parameter is specified, the command displays the quota report for volumes in the specified Vserver.

[`-volume <volume name>`] - Volume Name

If this parameter is specified, the command displays the quota report for the specified volume.

[`-index <integer>`] - Index

If this parameter is specified, the command displays the quota report for the quota rules that have the specified index.

[`-tree <qtree name>`] - Qtree Name

If this parameter is specified, the command displays the quota report for the quota rules that have the specified qtree name.

[`-quota-type {tree|user|group}`] - Quota Type

If this parameter is specified, the command displays the quota report for the quota rules of the given type.

[`-quota-target <text>,...`] - Quota Target

If this parameter is specified, the command displays the quota report for the quota rules that have the specified quota target.

[`-quota-target-id <text>,...`] - Quota Target ID

If this parameter is specified, the command displays the quota report for the quota rules that have the specified quota target identifier.

[`-disk-used {<integer>[KB|MB|GB|TB|PB]}`] - Disk Space Used

If this parameter is specified, the command displays the quota report for the quota rules that have the specified disk space used value.

[`-disk-limit {<integer>[KB|MB|GB|TB|PB]}`] - Disk Space Limit

If this parameter is specified, the command displays the quota report for the quota rules that have the specified disk space limit.

[`-files-used <integer>`] - Files Used

If this parameter is specified, the command displays the quota report for the quota rules that have the specified files used value.

[`-file-limit <integer>`] - Files Limit

If this parameter is specified, the command displays the quota report for the quota rules that have the specified file limit.

[`-threshold {<integer>[KB|MB|GB|TB|PB]}`] - Disk Space Threshold

If this parameter is specified, the command displays the quota report for the quota rules that have the specified threshold for disk space limit.

[-soft-disk-limit {<integer>[KB|MB|GB|TB|PB] }] - Soft Disk Space Limit

If this parameter is specified, the command displays the quota report for the quota rules that have the specified soft disk space limit.

[-soft-file-limit <integer>] - Soft Files Limit

If this parameter is specified, the command displays the quota report for the quota rules that have the specified soft file limit.

[-quota-specifier <text>] - Quota Specifier

If this parameter is specified, the command displays the quota report for the quota rules that have the specified quota specifier.

[-path <text>] - Path

If this parameter is specified, the command will display the quota report for the quota rules that are applicable for the file in the specified path. The format of the path to the file should begin with /vol/<volume name>/. The quota rules that are applicable typically consists of the tree quota rule corresponding to the qtree in which the file resides within the volume, user quota rule at the volume and qtree level corresponding to the UNIX user identifier or the Windows security identifier associated with the file and the group quota rule at the volume and qtree level corresponding to the UNIX group identifier associated with the file.

[-disk-used-pct-soft-disk-limit <percent_no_limit>] - Disk Space Used % Soft Disk Space Limit

If this parameter is specified, the command displays the quota report for entries that have the specified percent utilization. The attribute value is computed from `disk-used` and `soft-disk-limit`.

[-disk-used-pct-threshold <percent_no_limit>] - Disk Space Used % Disk Space Threshold

If this parameter is specified, the command displays the quota report for entries that have the specified percent utilization. The attribute value is computed from `disk-used` and `threshold`.

[-disk-used-pct-disk-limit <percent_no_limit>] - Disk Space Used % Disk Space Limit

If this parameter is specified, the command displays the quota report for entries that have the specified percent utilization. The attribute value is computed from `disk-used` and `disk-limit`.

[-files-used-pct-soft-file-limit <percent_no_limit>] - Files Used % Soft File Limit

If this parameter is specified, the command displays the quota report for entries that have the specified percent utilization. The attribute value is computed from `files-used` and `soft-file-limit`.

[-files-used-pct-file-limit <percent_no_limit>] - Files Used % File Limit

If this parameter is specified, the command displays the quota report for entries that have the specified percent utilization. The attribute value is computed from `files-used` and `file-limit`.

Examples

The following example displays the quota report for all the volumes.

cluster1::> volume quota report

Vserver: vs0

Volume	Tree	Type	ID	Used	Limit	Used	Limit	
vol2		tree	*	0.00B	100MB	0	10000	*
vol2	vxw02	tree	3	0.00B	200MB	1	20000	vxw02
vol2		user	*	0.00B	50MB	0	-	*
vol2	vxw02	user	sam,Engr\Sammy	0.00B	100MB	0	-	sam
vol2		group	*	0.00B	500MB	0	-	*
vol2	q1	tree	1	1MB	100MB	2	10000	q1
vol2	q1	user	*	0.00B	50MB	0	-	
vol2	q1	group	*	0.00B	500MB	0	-	
vol2	q1	group	root	1MB	-	2	-	
vol2	vxw01	tree	2	0.00B	100MB	1	10000	vxw01
vol2	vxw01	user	*	0.00B	50MB	0	-	
vol2	vxw01	group	*	0.00B	500MB	0	-	
vol2	vxw01	group	root	0.00B	-	1	-	
vol2	vxw02	user	*	0.00B	50MB	0	-	
vol2	vxw02	group	*	0.00B	500MB	0	-	
vol2	vxw02	group	root	0.00B	-	1	-	
vol2	vxw03	tree	4	0.00B	100MB	1	10000	vxw03
vol2	vxw03	user	*	0.00B	50MB	0	-	
vol2	vxw03	group	*	0.00B	500MB	0	-	
vol2	vxw03	group	root	0.00B	-	1	-	
vol2		group	root	1MB	-	6	-	
vol2	q1	user	root,Engr\root	0.00B	-	1	-	
vol2	vxw01	user	root,Engr\root	0.00B	-	1	-	
vol2	vxw02	user	root,Engr\root	0.00B	-	1	-	
vol2	vxw03	user	root,Engr\root	0.00B	-	1	-	
vol2		user	root,Engr\root	0.00B	-	5	-	
vol2		user	john,Engr\John	1MB	50MB	1	-	*
vol2	q1	user	john,Engr\John	1MB	50MB	1	-	

28 entries were displayed.

The following example displays the quota report for the quota rules that are applicable for the given path to a file.

```
cluster1::> volume quota report -path /vol/vol2/q1/file1
Vserver: vs0
----Disk----  ----Files-----  Quota
Volume  Tree      Type   ID          Used  Limit   Used  Limit
Specifier
-----  -
vol2    q1         tree   1           1MB  100MB   2    10000  q1
vol2    q1         group  root        1MB   -       2     -
vol2    q1         group  root        1MB   -       6     -
vol2    q1         user   john,Engr\John
                               1MB   50MB   1     -      *
vol2    q1         user   john,Engr\John
                               1MB   50MB   1     -

5 entries were displayed.
```

The following example displays the quota report with the target in the numeric form for the given path to a file.

```
cluster1::> volume quota report -path /vol/vol2/q1/file1 -target-id
Vserver: vs0
----Disk----  ----Files-----  Quota
Volume  Tree      Type   ID          Used  Limit   Used  Limit
Specifier
-----  -
vol2    q1         tree   1           1MB  100MB   2    10000  q1
vol2    q1         group  0           1MB   -       2     -
vol2    q1         group  0           1MB   -       6     -
vol2    q1         user   8017,S-1-5-21-3567637-1906459281-1427260136-
60871
                               1MB   50MB   1     -      *
vol2    q1         user   8017,S-1-5-21-3567637-1906459281-1427260136-
60871
                               1MB   50MB   1     -

5 entries were displayed.
```

The following example shows how to monitor the quota report for a particular user/tree/group. First, the quota report command is issued with `-instance` to see the index field for the quota report entry we are interested in. Next, the quota report is issued with the `-index` field specified to fetch only that particular quota report entry repeatedly to view the usage over time.

```
cluster1::> volume quota report -vserver vs0 -volume vol1 -quota-type user
```

```
-quota-target john -tree q1 -instance
```

```
Vserver Name: vs0
```

```
Volume Name: voll1
Index: 10
Qtree Name: q1
Quota Type: user
Quota Target: john
Quota Target ID: 5433
Disk Space Used: 50.5MB
Disk Space Limit: 100MB
Files Used: 205
Files Limit: -
Disk Space Threshold: 95MB
Soft Disk Space Limit: 80MB
Soft Files Limit: -
Quota Specifier: john
Disk Space Used % Soft Disk Space Limit: 63%
Disk Space Used % Disk Space Threshold: 53%
Disk Space Used % Disk Space Limit: 51%
Files Used % Soft File Limit: -
Files Used % File Limit: -
```

```
cluster1::> volume quota report -vserver vs0 -volume voll1 -index 10
```

```
Vserver Name: vs0
```

```
Volume Name: voll1
Index: 10
Qtree Name: q1
Quota Type: user
Quota Target: john
Quota Target ID: 5433
Disk Space Used: 55MB
Disk Space Limit: 100MB
Files Used: 410
Files Limit: -
Disk Space Threshold: 95MB
Soft Disk Space Limit: 80MB
Soft Files Limit: -
Quota Specifier: john
Disk Space Used % Soft Disk Space Limit: 69%
Disk Space Used % Disk Space Threshold: 58%
Disk Space Used % Disk Space Limit: 55%
Files Used % Soft File Limit: -
Files Used % File Limit: -
```

```
cluster1::> volume quota report -vserver vs0 -volume voll1 -index 10
```

```
Vserver Name: vs0
```

```
Volume Name: voll
Index: 10
Qtree Name: q1
Quota Type: user
Quota Target: john
Quota Target ID: 5433
Disk Space Used: 60.7MB
Disk Space Limit: 100MB
Files Used: 500
Files Limit: -
Disk Space Threshold: 95MB
Soft Disk Space Limit: 80MB
Soft Files Limit: -
Quota Specifier: john
Disk Space Used % Soft Disk Space Limit: 76%
Disk Space Used % Disk Space Threshold: 64%
Disk Space Used % Disk Space Limit: 61%
Files Used % Soft File Limit: -
Files Used % File Limit: -
```

volume quota resize

Resize quotas for volumes

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command resizes the quota limits for the quotas currently in effect for the specified volume. It creates a job to resize quotas. You can monitor the progress of the job by using the [job show](#) and [job watch-progress](#) commands.



Quotas must be activated before quota limits can be resized.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the name of the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This specifies the name of the volume on which you are resizing the quota limits and threshold.

[-foreground <true>] - Foreground Process

This optionally specifies whether the job created for resizing quotas runs as a foreground process. The default setting is *false* (that is, the operation runs in the background). When set to *true*, the command will not return until the job completes.

Examples

The following example resizes quotas on the volume named `voll` , which exists on Vserver `vs0` .

```
cluster1::> volume quota resize -vserver vs0 -volume voll
[Job 34] Job is queued: Quota RESIZE Operation on vserver vs0 volume voll.
```

The following example uses a 7G-compatible command to resize quotas on the volume named `voll` which exists on Vserver `vs0` .

```
cluster1::> vserver context vs0
vs0::> quota resize voll
[Job 35] Job is queued: Quota RESIZE Operation on vserver vs0 volume voll.
```

Related Links

- [job show](#)
- [job watch-progress](#)

volume quota show

Display quota state for volumes

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays information about quotas for online volumes. The command output depends on the parameters specified with the command. Quotas can only be administered on FlexVol volumes. If no parameters are specified, the command displays the following information for all volumes:

- Vserver name
- Volume name
- Quota state - quota state for this volume. The possible values are as follows:
 - `off` - this state indicates that quotas are deactivated for the volume.
 - `on` - this state indicates that quotas are activated for the volume.
 - `initializing` - this state indicates that quotas are being initialized for the volume.
 - `resizing` - this state indicates that quota limits are being resized for the volume.
 - `corrupt` - this state indicates that quotas are corrupt for this volume.
 - `mixed` - this state may only occur for FlexGroups, and indicates that the constituent volumes are not all in the same state.
- Scan status - percentage of the files in the volume scanned by the quota scanner that runs as part of activating quotas.
- Last error - most recently generated error message as part of the last quota operation (`on` or `resize`).

Present only if an error has been generated.

To display detailed information about all volumes, run the command with the `-instance` parameter. The detailed view provides all of the information in the previous list and the following additional information:

- Logging messages - whether quota exceeded syslog/EMS messages are logged or not. For volumes where the quota logging parameter is set to `on`, quota exceeded messages are generated when a NFS/CIFS operation or any internal Data ONTAP operation is being prevented because the quota disk usage is exceeding the disk limit or the quota file usage is exceeding the file limit. For quotas where the logging parameter is set to `off`, no quota exceeded messages are generated.
- Logging interval - frequency with which quota exceeded messages are logged. This parameter only applies to volumes that have the logging parameter set to `on`.
- Sub status - additional status about quotas for this volume. Following are the possible values reported:
 - `upgrading` - this indicates that the quota metadata format is being upgraded from an older version to a newer version for the volume.
 - `setup` - this indicates that the quotas are being setup on the volume.
 - `transferring rules` - this indicates that the quota rules are being transferred to the volume.
 - `scanning` - this indicates that the quota filesystem scanner is currently running on the volume.
 - `finishing` - this indicates that the quota `on` or `resize` operation is in the final stage of the operation.
 - `done` - this indicates that the quota operation is finished.
 - `none` - this indicates that there is no additional status.
- All errors - collection of all the error messages generated as part of the last quota operation (`on` or `resize`) since the volume was online.
- User quota enforced (advanced privilege only) - indicates whether there are user quota rules being enforced.
- Group quota enforced (advanced privilege only)- indicates whether there are group quota rules being enforced.
- Tree quota enforced (advanced privilege only) - indicates whether there are tree quota rules being enforced.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-logmsg]

If this parameter is specified, the command displays whether quota exceeded messages are logged and the logging interval for the quota messages.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Name

If this parameter is specified, the command displays information for the volumes in the specified Vserver.

[-volume <volume name>] - Volume Name

If this parameter is specified, the command displays information for the specified volume.

[-state <quota_state>] - Quota State

If this parameter is specified, the command displays information for the volumes that have the specified quota state.

[-scan-status <percent>] - Scan Status

If this parameter is specified, the command displays information about the volumes whose scan-status matches the specified percentage value. The scan status is displayed in the format [0-100]%.

[-logging {on|off}] - Logging Messages

If this parameter is specified, the command displays information about the volumes that have the specified logging setting.

[-logging-interval <text>] - Logging Interval

If this parameter is specified, the command displays information about the volumes that have the specified quota logging interval.

[-sub-status <text>] - Sub Quota Status

If this parameter is specified, the command displays information about the volumes that have the specified quota sub-status.

[-last-error <text>] - Last Quota Error Message

If this parameter is specified, the command displays information about the volumes whose last error matches the specified error message.

[-errors <text>] - Collection of Quota Errors

If this parameter is specified, the command displays information about the volumes whose collection of errors match the specified error message.

[-is-user-quota-enforced {true|false}] - User Quota enforced (privilege: advanced)

If this parameter is specified, the command displays information about the volumes that have the specified value for this field.

[-is-group-quota-enforced {true|false}] - Group Quota enforced (privilege: advanced)

If this parameter is specified, the command displays information about the volumes that have the specified value for this field.

[-is-tree-quota-enforced {true|false}] - Tree Quota enforced (privilege: advanced)

If this parameter is specified, the command displays information about the volumes that have the specified value for this field.

Examples

The following example displays information about all volumes on the Vserver named `vs0`.

```

cluster1::> volume quota show -vserver vs0

```

Vserver	Volume	State	Scan Status
vs0	root_vs0	off	-
vs0	vol1	off	-
Last Error: Volume vol1 has no valid quota rules			
vs0	vol2	on	-
vs0	vol3	initializing	30%

4 entries were displayed.

The following example displays the logging information for all the volumes.

```

cluster1::> volume quota show -logmsg

```

Vserver	Volume	State	Logging Message	Logging Interval
vs0	root_vs0	off	-	-
vs0	vol1	off	-	-
vs0	vol2	on	on	1h
vs0	vol3	on	on	1h

4 entries were displayed.

The following example displays detailed information in advanced privilege for a volume `vol1`, which exists on Vserver `vs0`

```
cluster1::> set advanced
```

```
Warning: These advanced commands are potentially dangerous; use them only  
when
```

```
directed to do so by NetApp personnel.
```

```
Do you want to continue? {y|n}: y
```

```
cluster1::*> volume quota show -instance -vserver vs0 -volume vol1
```

```
Vserver Name: vs0  
    Volume Name: vol1  
    Quota State: on  
    Scan Status: -  
    Logging Messages: on  
    Logging Interval: 1h  
    Sub Quota Status: none  
Last Quota Error Message: -  
Collection of Quota Errors: -  
    User Quota enforced: true  
    Group Quota enforced: false  
    Tree Quota enforced: false
```

The following example displays the detailed information when quotas are upgrading for volume `vol1` , which exists on Vserver `vs0` .

```
cluster1::> volume quota show -instance -vserver vs0 -volume vol1
```

```
Vserver Name: vs0  
    Volume Name: vol1  
    Quota State: initializing  
    Scan Status: 3%  
    Logging Messages: -  
    Logging Interval: -  
    Sub Quota Status: upgrading  
Last Quota Error Message: -  
Collection of Quota Errors: -
```

The following example displays the "Last Quota Error Message" and the "Collection of Quota Errors" for volume `vol1` , which exists on Vserver `vs0`

```
cluster1::> volume quota show -instance -vserver vs0 -volume voll
    Vserver Name: vs0
    Volume Name: voll
    Quota State: on
    Scan Status: -
    Logging Messages: on
    Logging Interval: 1h
    Sub Quota Status: none
    Last Quota Error Message: second definition for user2 (type:user
target:user2,user4 qtree:"").
    Collection of Quota Errors: second definition for user1 (type:user
target:user1,user3 qtree:"").
                                second definition for user2 (type:user
target:user2,user4 qtree:"").
```

volume quota policy copy

Copy a quota policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command copies a quota policy and the rules it contains. You must enter the following information to copy a quota policy:

- Vserver name
- Policy name
- New policy name

Parameters

-vserver <vserver name> - Vserver

This parameter specifies the Vserver from which you are copying the quota policy.

-policy-name <text> - Policy Name

This parameter specifies the name of the quota policy you are copying.

-new-policy-name <text> - New Policy Name

This parameter specifies the name of the new quota policy you are copying to. The new name cannot have more than 32 characters.

Examples

The following example copies a quota policy named `quota_policy_0` on Vserver `vs0`. It is copied to `quota_policy_1`.

```
cluster1::> volume quota policy copy -vserver vs0 -policy-name
quota_policy_0 -new-policy-name quota_policy_1
```

volume quota policy create

Create a quota policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

A quota policy is collection of quota rules for all the volumes in a specific Vserver. This command creates a quota policy for a specific Vserver. Multiple quota policies can be created for a Vserver, but only one of them can be assigned to the Vserver. A Vserver can have a maximum of five quota policies. If five quota policies already exist, the command fails and a quota policy must be deleted before another quota policy can be created.

When you turn on quotas for a volume, the quota rules to be enforced on that volume will be picked from the quota policy that is assigned to the Vserver containing that volume. The quota policy for clustered volumes is equivalent to the `/etc/quotas` file in 7G.

You must enter the following information to create a quota policy:

- Vserver name
- Policy name

Parameters

-vserver <vserver name> - Vserver

This parameter specifies the Vserver for which you are creating the quota policy. You can create a quota policy only for a data Vserver. Quota policies cannot be created for a node or admin Vserver.

-policy-name <text> - Policy Name

This parameter specifies the name of the quota policy you are creating. The quota policy name cannot be more than 32 characters long and must be unique within the Vserver.

Examples

The following example creates a quota policy named `quota_policy_0` on Vserver `vs0`.

```
cluster1::> volume quota policy create -vserver vs0 -policy-name
quota_policy_0
```

volume quota policy delete

Delete a quota policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command deletes a quota policy and all the rules it contains. The policy can be deleted only when it is not assigned to the Vserver. You must enter the following information to delete a quota policy:

- Vserver name
- Policy name

Parameters

-vserver <vserver name> - Vserver

This parameter specifies the Vserver containing the quota policy you want to delete.

-policy-name <text> - Policy Name

This parameter specifies the name of the quota policy you want to delete.

Examples

The following example deletes a quota policy named `quota_policy_5` on Vserver `vs0`.

```
cluster1::> volume quota policy delete -vserver vs0 -policy-name
quota_policy_5
```

volume quota policy rename

Rename a quota policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command renames a quota policy. You must enter the following information to rename a quota policy:

- Vserver name
- Policy name
- New policy name

Parameters

-vserver <vserver name> - Vserver

This parameter specifies the Vserver containing the quota policy you want to rename.

-policy-name <text> - Policy Name

This parameter specifies the name of the quota policy you are renaming.

-new-policy-name <text> - New Policy Name

This parameter specifies the new name of the quota policy. The new name cannot be more than 32 characters long.

Examples

The following example renames a quota policy named `quota_policy_0` on Vserver `vs0`. The policy's new name is `quota_policy_1`.

```
cluster1::> volume quota policy rename -vserver vs0 -policy-name
quota_policy_0 -new-policy-name quota_policy_1
```

volume quota policy show

Display the quota policies

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays information about quota policies. The command displays the following information about all quota policies:

- Vserver name
- Policy name
- When the policy was last modified

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver

If this parameter is specified, the command displays information about the quota policies for the specified Vserver.

[-policy-name <text>] - Policy Name

If this optional parameter is specified, the command displays information about quota policies that match the specified name.

[-last-modified <MM/DD/YYYY HH:MM:SS>] - Last Modified

If this optional parameter is specified, the command displays information about quota policies with the last modified time that match the given time.

Examples

The following example displays information about all quota policies.


```

cluster1::> volume quota policy show
Vserver          Policy Name          Last Modified
-----
vs0              quota_policy_vs0    10/16/2008 17:40:05
vs1              quota_policy_vs1    10/16/2008 17:47:45
vs2              quota_policy_vs2    10/16/2008 17:44:13
vs3              quota_policy_vs3    10/16/2008 17:44:13
4 entries were displayed.

```

volume quota policy rule create

Create a new quota rule

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command creates a quota policy rule. You must enter the following information to create a quota policy rule:

- Vserver name
- Quota policy name
- Volume name
- Quota target type
- Target to which the rule applies
- Qtree to which the rule applies

You can optionally specify the following additional attributes for the quota policy rule:

- User mapping
- Hard disk limit
- Hard file limit
- Threshold for disk limit
- Soft disk limit
- Soft file limit



For a new quota policy rule to get enforced on the volume, you should create the rule in the quota policy assigned to the Vserver. Additionally, a quota off and on or a quota resize operation must be done using the "[volume quota modify](#)" command.

Parameters

-vserver <vserver name> - Vserver

This parameter specifies the Vserver containing the quota policy for which you are creating a rule.

-policy-name <text> - Policy Name

This parameter specifies the name of the quota policy in which you are creating a rule.

-volume <volume name> - Volume Name

This parameter specifies the name of the volume for which you are creating a rule.

-type {tree|user|group} - Type

This parameter specifies the quota target type of the rule you are creating.

-target <text> - Target

This parameter specifies the target to which the quota policy rule applies. For default quota rules, this parameter should be specified as "". For explicit tree quotas rules, this parameter should indicate the qtree name. For explicit user quota rules, this parameter can contain UNIX user name, UNIX user identifier, Windows user name, Windows Security Identifier or a path to an existing object within the volume. If a name contains a space, enclose the entire value in quotes. A UNIX user name cannot include a backslash (\) or an @ sign; user names with these characters are treated as Windows names. For multi-user quotas, this parameter can contain multiple user targets separated by a comma. For explicit group quota rules, this parameter can contain UNIX group name or UNIX group identifier or a path to an existing object within the volume. When a path is specified as the target, it should be of the format /vol/<vol-name>/<path to file from volume root> where the volume matches that of the -volume parameter.

-qtree <qtree name> - Qtree Name

This 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 "".

[-user-mapping {on|off}] - User Mapping

This parameter optionally specifies if user mapping needs to be performed for a user quota rule. If this parameter is "on", the UNIX user name specified as the quota target will be mapped to the corresponding Windows user name or vice-versa and quota accounting will be performed for the users together. The mapping will be obtained as configured in "vserver name-mapping".

Note that this parameter can be specified only for quota policy rules of type user. A value of "on" can be specified for this parameter only if the quota target is a UNIX user name or a Windows user name and cannot be specified for multi-user quota targets.

[-disk-limit {<size>|-}] - Disk Limit

This parameter optionally specifies a hard limit for the disk space that can be consumed by the quota target. The default unit for the disk limit is assumed to be Kilobytes if no units are specified. When the hard disk space limit is reached, no additional disk space can be consumed by the specified target. The value that you specify for this parameter should be greater than or equal to the threshold and soft disk limit. A disk limit of unlimited can be specified with a "-" for this parameter or by not specifying this parameter and will be indicated by a "-". The maximum value is 1,125,899,906,842,620 KB, which is approximately 1,023 PB. The value should be a multiple of 4 KB. If it is not, this field can appear incorrect in quota reports. This happens because the field is always rounded up to the nearest multiple of 4 KB to match disk space limits, which are translated into 4-KB disk blocks. The value can be larger than the amount of disk space available in the volume.

[-file-limit {<integer>|-}] - Files Limit

This parameter optionally specifies a hard limit for the number of files permitted on the quota target. When the hard number of files limit is reached, no additional files can be created by the specified target. The value that you specify for this parameter should be greater than or equal to the soft file limit. A file limit of

unlimited can be specified with a "-" for this parameter or by not specifying this parameter and will be indicated by a "-". The maximum value is 4,294,967,295.

[`-threshold {<size>|-}`] - Threshold for Disk Limit

This parameter optionally specifies the disk limit threshold for the quota target. The default unit for the disk limit threshold is assumed to be Kilobytes if no units are specified. When the disk limit threshold is exceeded, a console message, EMS events, and SNMP traps are generated. The value that you specify for this parameter should be greater than or equal to the soft disk limit and equal to or less than the disk limit. A threshold of unlimited can be specified with a "-" for this parameter or by not specifying this parameter and will be indicated by a "-". The maximum value is 1,125,899,906,842,620 KB, which is approximately 1,023 PB. The value should be a multiple of 4 KB. If it is not, this field can appear incorrect in quota reports. This happens because the field is always rounded up to the nearest multiple of 4 KB to match disk space limits, which are translated into 4-KB disk blocks.

[`-soft-disk-limit {<size>|-}`] - Soft Disk Limit

This parameter optionally specifies a soft limit for the disk space that can be consumed by the quota target. The soft disk limit indicates that the hard limit for the disk space will soon be exceeded. The default unit for the soft disk limit is assumed to be Kilobytes if no units are specified. When the soft limit for the disk space is exceeded, a console message, EMS events and SNMP traps are generated. The same happens when the disk space used goes below the specified limit. The value that you specify for this parameter should be equal to or less than the threshold and the disk limit. A soft disk limit of unlimited can be specified with a "-" for this parameter or by not specifying this parameter and will be indicated by a "-". The maximum value is 1,125,899,906,842,620 KB, which is approximately 1,023 PB. The value should be a multiple of 4 KB. If it is not, this field can appear incorrect in quota reports. This happens because the field is always rounded up to the nearest multiple of 4 KB to match disk space limits, which are translated into 4-KB disk blocks.

[`-soft-file-limit {<integer>|-}`] - Soft Files Limit

This parameter optionally specifies a soft limit for the number of files permitted on the quota target. The soft file limit indicates that the hard limit for the number of files will soon be exceeded. When the soft limit for the number of files is exceeded, a console message, EMS events and SNMP traps are generated. The same happens when the files used goes below the specified limit. The value that you specify for this parameter should be equal to or less than the file limit. A soft file limit of unlimited can be specified with a "-" for this parameter or by not specifying this parameter and will be indicated by a "-". The maximum value is 4,294,967,295.

Examples

The following example creates a default tree quota rule for volume `vol0` in Vserver `vs0` and in the quota policy named `quota_policy_0`. This quota policy applies to all qtrees on volume `vol0`.

```
cluster1::> volume quota policy rule create -vserver vs0
-policy-name quota_policy_0 -volume vol0 -type tree -target ""
```

The following example creates a quota policy rule for volume `vol0` in Vserver `vs0` and in the quota policy named `quota_policy_0`. This quota policy applies to the UNIX user `myuser` for a qtree named `qtrees1` on volume `vol0` with a disk limit of 20 Gigabytes, soft disk limit of 15.4 Gigabytes and threshold limit of 15.4 Gigabytes. User mapping is turned on for this rule.

```
cluster1::> volume quota policy rule create -vserver vs0
-policy-name quota_policy_0 -volume vol0 -type user -target myuser
-qtrees qtrees1 -user-mapping on -disk-limit 20GB -soft-disk-limit 15.4GB
-threshold 15.4GB
```

The following example creates a quota policy rule for volume *vol0* in Vserver *vs0* and in the quota policy named *quota_policy_0*. This quota policy applies to the Windows user *DOMXYZ\myuser* for a qtrees named *qtrees1* on volume *vol0* with a file limit of *40000* and a soft file limit of *26500*. User mapping is turned on for this rule.

```
cluster1::> volume quota policy rule create -vserver vs0
-policy-name quota_policy_0 -volume vol0 -type user -target DOMXYZ\myuser
-qtrees qtrees1 -user-mapping on -file-limit 40000 -soft-file-limit 26500
```

The following example creates a quota policy rule for volume *vol0* in Vserver *vs0* and in the quota policy named *quota_policy_0*. This quota policy applies to the UNIX user identifier *12345* for a qtrees named *qtrees1* on volume *vol0*.

```
cluster1::> volume quota policy rule create -vserver vs0
-policy-name quota_policy_0 -volume vol0 -type user -target 12345
-qtrees qtrees1
```

The following example creates a quota policy rule for volume *vol0* in Vserver *vs0* and in the quota policy named *quota_policy_0*. This quota policy applies to the Windows Security Identifier *S-123-456-789* for a qtrees named *qtrees1* on volume *vol0*.

```
cluster1::> volume quota policy rule create -vserver vs0
-policy-name quota_policy_0 -volume vol0 -type user
-target S-123-456-789 -qtrees qtrees1
```

The following example creates a quota policy rule for volume *vol0* in Vserver *vs0* and in the quota policy named *quota_policy_0*. This quota policy applies to the UNIX group *enr* for a qtrees named *qtrees1* on volume *vol0*.

```
cluster1::> volume quota policy rule create -vserver vs0
-policy-name quota_policy_0 -volume vol0 -type group -target enr
-qtrees qtrees1
```

The following example creates a quota policy rule for volume *vol0* in Vserver *vs0* and in the quota policy named *quota_policy_0*. This quota policy applies to the user who is the owner of the file */vol/vol0/qtrees1/file1.txt* for qtrees *qtrees1* on volume *vol0*.

```
cluster1::> volume quota policy rule create -vserver vs0 -policy-name
quota_policy_0 -volume vol0 -type user -target /vol/vol0/qtreen1/file1.txt
-qtreen qtreen1
```

The following example creates a quota policy rule for volume *vol0* in Vserver *vs0* and in the quota policy named *quota_policy_0*. This quota policy applies to the users specified in the target for qtree *qtreen1* on volume *vol0*.

```
cluster1::> volume quota policy rule create -vserver vs0
-policy-name quota_policy_0 -volume vol0 -type user
-target user1,DOMXYZ\user2,23457,S-126-098-567,/vol/vol0/qtreen1/file2.txt
-qtreen qtreen1
```

Related Links

- [volume quota modify](#)

volume quota policy rule delete

Delete an existing quota rule

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume quota policy rule delete` command deletes a quota policy rule. You must enter the following information to delete a quota policy rule:

- Vserver name
- Quota policy name
- Volume name
- Quota target type
- Target to which the rule applies
- Qtree to which the rule applies



If the rule being deleted belongs to the quota policy that is currently assigned to the Vserver, enforcement of the rule on the volume must be terminated by performing a quota off and on or a quota resize operation using the "[volume quota modify](#)" command.

Parameters

-vserver <vserver name> - Vserver

This parameter specifies the Vserver containing the quota policy for which you are deleting a rule.

-policy-name <text> - Policy Name

This parameter specifies the name of the quota policy in which you are deleting a rule.

-volume <volume name> - Volume Name

This parameter specifies the name of the volume for which you are deleting a rule.

-type {tree|user|group} - Type

This parameter specifies the quota target type for the rule.

-target <text> - Target

This parameter specifies the target to which the quota policy rule applies.

-qtree <qtree name> - Qtree Name

This parameter specifies the name of the qtree for which you are deleting a rule.

Examples

The following example deletes a quota policy rule on Vserver vs1 for the quota policy named quota_policy_1. This quota policy applies to the group named engr for the qtree named qtree1 on volume vol1.

```
cluster1::> volume quota policy rule delete -vserver vs1
-policy-name quota_policy_1 -volume vol1 -type group -target engr
-qtree qtree1
```

Related Links

- [volume quota modify](#)

volume quota policy rule modify

Modify an existing quota rule

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command can be used to modify the following attributes of a quota policy rule:

- User mapping
- Hard disk limit
- Hard file limit
- Threshold for disk limit
- Soft disk limit
- Soft file limit



If the rule being modified belongs to the quota policy that is currently assigned to the Vserver, rule enforcement on the volume must be enabled by performing a quota off and on or a quota resize operation using the "[volume quota modify](#)" command.

Parameters

-vserver <vserver name> - Vserver

This parameter specifies the Vserver containing the quota policy for which you are modifying a rule.

-policy-name <text> - Policy Name

This parameter specifies the name of the quota policy in which you are modifying a rule.

-volume <volume name> - Volume Name

This parameter specifies the name of the volume for which you are modifying a rule.

-type {tree|user|group} - Type

This parameter specifies the quota target type for the rule you are modifying.

-target <text> - Target

This parameter specifies the target to which the quota policy rule applies. If the target is a user, the user ID or username must be the same one that was used to create the quota. The same restriction is there for both group ID or groupname and Windows SID or Windows account name.

-qtree <qtree name> - Qtree Name

This parameter specifies the name of the qtree to which the quota policy rule applies.

[-user-mapping {on|off}] - User Mapping

This parameter optionally modifies the user mapping for a user quota rule. The value for this parameter can be modified only for quota policy rules of type user. A value of "on" can be specified for this parameter only if the quota target is a unix user name or a Windows user name and cannot be specified for multi-user quota targets. If this parameter is "on", the unix user name specified as the quota target will be mapped to the corresponding Windows user name or vice-versa and quota accounting will be performed for the users together.

[-disk-limit {<size>|-}] - Disk Limit

This parameter optionally modifies the hard limit for the disk space that can be consumed by the quota target. The default unit for the disk limit is assumed to be Kilobytes if no units are specified. The value that you specify for this parameter should be greater than or equal to the threshold and soft disk limit. A disk limit of unlimited can be specified with a "-" for this parameter.

[-file-limit {<integer>|-}] - Files Limit

This parameter optionally modifies the hard limit for the number of files permitted on the quota target. The value that you specify for this parameter should be greater than or equal to the soft file limit. A file limit of unlimited can be specified with a "-" for this parameter.

[-threshold {<size>|-}] - Threshold for Disk Limit

This parameter optionally modifies the disk limit threshold for the quota target. The default unit for the disk limit threshold is assumed to be Kilobytes if no units are specified. The value that you specify for this parameter should be greater than or equal to the soft disk limit and equal to or less than the disk limit. A threshold limit of unlimited can be specified with a "-" for this parameter.

[`-soft-disk-limit` {<size>|-}] - Soft Disk Limit

This parameter optionally modifies the soft limit for the disk space that can be consumed by the quota target. The default unit for the soft disk limit is assumed to be Kilobytes if no units are specified. The value that you specify for this parameter should be equal to or less than the threshold and the disk limit. A soft disk limit of unlimited can be specified with a "-" for this parameter.

[`-soft-file-limit` {<integer>|-}] - Soft Files Limit

This parameter optionally modifies the soft limit for the number of files permitted on the quota target. The value that you specify for this parameter should be equal to or less than the file limit. A soft file limit of unlimited can be specified with a "-" for this parameter.

Examples

The following example modifies a quota policy rule for the quota policy named `quota_policy_0`. This quota policy exists on Vserver `vs0` and applies to the user named `myuser` for qtree named `qtree1` on volume `vol0`. The user mapping is turned on, the hard disk limit is set to 20 GB and the hard file limit is set to 100,000 files.

```
cluster1::> volume quota policy rule modify -vserver vs0
-policy-name quota_policy_0 -volume vol0 -type user -target myuser
-qtree qtree1 -user-mapping on -disk-limit 20GB -file-limit 100000
```

Related Links

- [volume quota modify](#)

volume quota policy rule show

Display the quota rules

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays the following information about quota policy rules by default.

- Vserver name
- Quota policy name
- Volume name
- Type of quota policy rule
- Target of the quota policy rule
- Qtree name
- User mapping
- Hard disk limit
- Soft disk limit
- Hard file limit
- Soft file limit

- Threshold for disk limit

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver

If this parameter is specified, the command displays information about quota rules for the quotas contained on volumes on the specified Vserver.

[-policy-name <text>] - Policy Name

If this parameter is specified, the command displays information about quota rules for the specified quota policy.

[-volume <volume name>] - Volume Name

If this parameter is specified, the command displays information about quota rules for the quota policy associated with the specified volume.

[-type {tree|user|group}] - Type

If this parameter is specified, the command displays information about quota rules for the specified quota type.

[-target <text>] - Target

If this parameter is specified, the command displays information about quota rules for the specified target.

[-qtree <qtree name>] - Qtree Name

If this parameter is specified, the command displays information about quota rules for the specified qtree.

[-user-mapping {on|off}] - User Mapping

If this parameter is specified, the command displays information about quota rules having the specified user-mapping value.

[-disk-limit {<size>|-}] - Disk Limit

If this parameter is specified, the command displays information about quota rules having the specified hard disk limit.

[-file-limit {<integer>|-}] - Files Limit

If this parameter is specified, the command displays information about quota rules having the specified hard file limit.

[-threshold {<size>|-}] - Threshold for Disk Limit

If this parameter is specified, the command displays information about quota rules having the specified disk limit threshold.

[`-soft-disk-limit` {<size>|-}] - Soft Disk Limit

If this parameter is specified, the command displays information about quota rules having the specified soft disk limit.

[`-soft-file-limit` {<integer>|-}] - Soft Files Limit

If this parameter is specified, the command displays information about quota rules having the specified soft file limit.

Examples

The following example displays information about all the quota policy rules in a cluster. There is one user rule that exists on Vserver vs0 for the quota policy named `quota_policy_0`. This quota policy applies to the user named `myuser` for qtree named `qtree0` on volume `vol0`.

```
cluster1::> volume quota policy rule show
Vserver: vs0      Policy: quota_policy_0      Volume: vol0
Soft      Soft
Type      Target  Qtree  User      Disk      Disk      Files      Files
-----  -
tree      myuser  qtree0 on      20GB      18GB      100000     80000     16GB
```

volume quota policy rule count show

Display count of quota rules

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

This command displays various counts of quota policy rules defined within a quota policy. By default, the subtotal for each volume is displayed. Optionally, the command can provide the total rule count across the entire quota policy or detailed subtotals organized by qtree and quota rule type.

Parameters

{ [`-fields` <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [`-detail`]

Displays rule count subtotals for each quota rule type. The subtotals for each type are computed for a specific volume and qtree.

| [-hierarchy]

Displays rule count subtotals in hierarchical format with subtotals at the quota policy, volume, qtree, and quota rule type levels.

| [-total]

Displays the total rule count for each Vserver and quota policy.

| [-instance] }

Displays detailed information about all fields.

[-vserver <vserver name>] - Vserver

Displays quota rule counts for the specified Vserver.

[-policy-name <text>] - Policy Name

Displays quota rule counts for the specified quota policy.

[-volume <volume name>] - Volume Name

Displays quota rule counts for the specified volume.

[-qtree <qtree name>] - Qtree Name

Displays quota rule counts for the specified qtree.

[-type {tree|user|group}] - Type

Displays quota rule counts for the specified quota rule type.

[-count-where-policy-volume-qtree-type <integer>] - Qtree/Type Subtotal

Subtotal of rules matching the given Vserver, quota policy, volume, qtree, and quota rule type. If specified as input, only matching totals are displayed.

[-count-where-policy-volume-qtree <integer>] - Qtree Subtotal

Subtotal of rules matching the given Vserver, quota policy, volume, and qtree. All quota rule types are included. If specified as input, only matching totals are displayed.

[-count-where-policy-volume-type <integer>] - Volume/Type Subtotal

Subtotal of rules matching the given Vserver, quota policy, volume, and quota rule type. All qtrees are included. If specified as input, only matching totals are displayed.

[-count-where-policy-volume <integer>] - Volume Subtotal

Subtotal of rules matching the given Vserver, quota policy, and volume. All qtrees and quota rule types are included. If specified as input, only matching totals are displayed.

[-count-where-policy-type <integer>] - Policy/Type Subtotal

Subtotal of rules matching the given Vserver, quota policy, and quota rule type. All volumes and qtrees are included. If specified as input, only matching totals are displayed.

[-count-where-policy <integer>] - Policy Total

Total rule count matching the given Vserver and quota policy. All volumes, qtrees, and quota rule types are included. If specified as input, only matching totals are displayed.

Examples

The following example shows quota rule counts for Vserver *vs0*, quota policy *default*. The total number of rules in quota policy *default* is 7500. There are two volumes with quota rules configured. Volume *volume0* has a total of 1000 rules, and *volume1* has a total of 6500 rules.

```
cluster1::> volume quota policy rule count show -vserver vs0 -policy-name
default

Vserver: vs0                Policy: default
Rule
Volume                      Count
-----
volume0                     1000
volume1                     6500
2 entries were displayed.
```

volume reallocation commands

volume reallocation measure

Start reallocate measure job

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

Performs a measure-only reallocation check on a LUN, NVMe namespace, file, or volume. At the end of each check, the system logs the optimization results in the Event Message System (EMS). If you use the `logfile`, the system records detailed information about the LUN, NVMe namespace, file, or volume layout in the log file. To view previous measure-only reallocation checks, use the [volume reallocation show](#) command.



This command is not supported for FlexGroups or FlexGroup constituents.

Parameters

-vserver <vserver name> - Vserver

Specifies the Vserver.

-path <text> - Path

Specifies the path of the reallocation for a LUN, NVMe namespace, file, or volume.

{ [-i, -interval <text>] - Interval Schedule

Specifies the reallocation scan interval in

- m for minutes
- h for hours

- d for days

For example, 30m is a 30 minute interval. The countdown to the next scan begins after the first scan is complete.

The default interval is 24 hours.

[`-o`, `-once <true>`] - Once }

Specifies that the job runs once and then is automatically removed from the system when set to true. If you use this command without specifying this parameter, its effective value is false and the reallocation scan runs as scheduled. If you enter this parameter without a value, it is set to true and a reallocation scan runs once.

[`-l`, `-logpath <text>`] - Log Path

Specifies the path for reallocation logs.

[`-t`, `-threshold <integer>`] - Threshold

Specifies the threshold when a LUN, NVMe namespace, file, or volume is considered unoptimized and a reallocation should be performed. Once the threshold is reached, the system creates a diagnostic message that indicates that a reallocation might improve performance.

The threshold range is from 3 (the layout is moderately optimized) to 10 (the layout is not optimal). The threshold default is 4.

Examples

```
cluster1::> volume reallocation measure -path /vol/vol2 -once
[Job 167] Job is queued: Reallocate Job.
```

Performs a one-time, measure-only reallocation scan on volume vol2.

Related Links

- [volume reallocation show](#)

volume reallocation off

Disable reallocate jobs

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

Disables all reallocation jobs globally in a cluster. After you use this command, you cannot start or restart any reallocation jobs. All jobs that are executing when you use this command are stopped. You must use the `reallocate on` command to enable or restart reallocation jobs globally in a cluster.



This command is not supported for FlexGroups or FlexGroup constituents.

Examples

```
cluster1::> volume reallocation off
```

Disables all reallocation jobs globally in a cluster.

volume reallocation on

Enable reallocate jobs

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

Globally enables all reallocation jobs in a cluster. You must globally enable reallocation scans in the cluster before you can run a scan or schedule regular scans. Reallocation scans are disabled by default.



This command is not supported for FlexGroups or FlexGroup constituents.

Examples

```
cluster1::> volume reallocation on
```

Globally enables all reallocation jobs on a cluster.

volume reallocation quiesce

Quiesce reallocate job

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

Temporarily stops any reallocation jobs that are in progress. When you use this command, the persistent state is saved. You can use the [volume reallocation restart](#) command to restart a job that is quiesced.

There is no limit to how long a job can remain in the quiesced state.



This command is not supported for FlexGroups or FlexGroup constituents.

Parameters

-vserver <vserver name> - Vserver

Specifies the Vserver.

-path <text> - Path

Specifies the file path of the LUN, NVMe namespace, file, or volume that you want to stop temporarily.

Examples

```
cluster1::> volume reallocation quiesce /vol/vol2
2 entries were acted on.
```

Temporarily stops all reallocation jobs on volume vol2.

Related Links

- [volume reallocation restart](#)

volume reallocation restart

Restart reallocate job

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

Starts a reallocation job. Use this command to start a quiesced (temporarily stopped) job or a scheduled scan that is idle.



This command is not supported for FlexGroups or FlexGroup constituents.

Parameters

-vserver <vserver name> - Vserver

Specifies the Vserver.

-path <text> - Path

Specifies the file path of the LUN, NVMe namespace, file, or volume on which you want to restart reallocation scans.

[-i, -ignore-checkpoint <true>] - Ignore Checkpoint

Restarts the job at the beginning when set to true. If you use this command without specifying this parameter, its effective value is false and the job starts the scan at the point where it stopped. If you specify this parameter without a value, it is set to true and the scan restarts at the beginning.

Examples

```
cluster1::> volume reallocation restart /vol/vol2
2 entries were acted on.
```

Restarts two reallocation jobs on volume vol2.

volume reallocation schedule

Modify schedule of reallocate job

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

Schedules a reallocation scan for an existing reallocation job. If the reallocation job does not exist, use the [volume reallocation start](#) command to define a reallocation job.

You can delete an existing reallocation scan schedule. However, if you do this, the job's scan interval reverts to the schedule that was defined for it when the job was created with the [volume reallocation start](#) command.



This command is not supported for FlexGroups or FlexGroup constituents.

Parameters

-vserver <vserver name> - Vserver

Specifies the Vserver.

-path <text> - Path

Specifies the path of the reallocation for a LUN, NVMe namespace, file, or volume.

[-d, -del <>true>] - Delete

Deletes an existing reallocation schedule when set to true. If you use this command without specifying this parameter, its effective value is false and the reallocation schedule is not deleted. If you specify this parameter without a value, it is set to true and the reallocation schedule is deleted.

[-s, -cron <text>] - Cron Schedule

Specifies the schedule with the following four fields in sequence. Use a space between field values. Enclose the values in double quotes.

- minute is a value from 0 to 59.
- hour is a value from 0 (midnight) to 23 (11:00 p.m.).
- day of week is a value from 0 (Sunday) to 6 (Saturday).
- day of month is a value from 1 to 31.



If you specify 31 as the value for the day of month, reallocation scans will not run in any months with fewer than 31 days.

Use an asterisk "*" as a wildcard to indicate every value for that field. For example, an * in the day of month field means every day of the month. You cannot use the wildcard in the minute field.

You can enter a number, a range, or a comma-separated list of values for a field.

Examples

```
cluster1::> volume reallocation schedule -s "0 23 6 *" /vol/db/lun1
```

Schedules a reallocation job to run at 11 pm every Saturday on lun1.

Related Links

- [volume reallocation start](#)

volume reallocation show

Show reallocate job status

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

Displays the status of a reallocation scan, including the state, schedule, interval, optimization, and log files. If you do not specify the `path` for a particular reallocation scan, then the command displays all the reallocation scans.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-v]

Specify this parameter to display the output in a verbose format.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver

Specify this parameter to display reallocation scans that match the Vserver that you specify.

[-path <text>] - Path

Specify this parameter to display reallocation scans that match the path that you specify.

[-t, -threshold <integer>] - Threshold

Specify this parameter to display reallocation scans that match the threshold that you specify.

[-id <integer>] - Job ID

Specify this parameter to display reallocation scans that match the reallocation job ID that you specify.

[-description <text>] - Job Description

Specify this parameter to display reallocation scans that match the text description that you specify.

[-state

{Initial|Queued|Running|Waiting|Pausing|Paused|Quitting|Success|Failure|Reschedule|Error|Quit|Dead|Unknown|Restart|Dormant}] - Job State

Specify this parameter to display reallocation jobs that match the state that you specify.

[-progress <text>] - Execution Progress

Specify this parameter to list the running reallocation jobs whose progress indicator matches the text that you provide. For example, if you specify "Starting ..." as the text string for the progress option, then the

system lists all of the jobs that are starting.

[`-schedule <job_schedule>`] - Schedule Name

Specify this parameter to display reallocation scans that match the schedule name that you specify. If you want a list of all job schedules, use the [job schedule show](#) command.

[`-global-status <text>`] - Global State of Scans

Specify this parameter to indicate if reallocation scans are on or off globally. You must type either of the following text strings:

- "Reallocation scans are on"
- "Reallocation scans are off"

Examples

```
cluster1::> volume reallocation show
Vserver      Description                               Schedule      State
-----      -
Reallocation scans are on
vs0          /vol/vol2,space-optimized               reallocate_1d  Queued
```

Displays the Vserver, description, schedule, and state for the reallocation scans on the local node.

Related Links

- [job schedule show](#)

volume reallocation start

Start reallocate job

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

Begins a reallocation scan on a LUN, NVMe namespace, file, or volume when you specify the path. If a volume has several small files that would benefit from periodic optimization, specify the `/vol/volname`.

Before performing a reallocation scan, the reallocation job normally performs a check of the current layout optimization. If the current layout optimization is less than the threshold, then the system does not perform a reallocation on the LUN, NVMe namespace, file, or volume.

You can define the reallocation scan job so that it runs at a specific interval, or you can use the [volume reallocation schedule](#) command to schedule reallocation jobs.



This command is not supported for FlexGroups or FlexGroup constituents.

Parameters

-vserver <vserver name> - Vserver

Specifies the Vserver.

-path <text> - Path

Specifies the path of the reallocation for a LUN, NVMe namespace, file, or volume.

{ [-i, -interval <text>] - Interval Schedule

Specifies the reallocation scan interval in

- m for minutes
- h for hours
- d for days

For example, 30m is a 30 minute interval. The countdown to the next scan begins after the first scan is complete.

The default interval is 24 hours.

| [-o, -once <>true>] - Once

Specifies that the job runs once and then is automatically removed from the system when set to true. If you use this command without specifying this parameter, its effective value is false and the reallocation scan runs as scheduled. If you enter this parameter without a value, it is set to true and a reallocation scan runs once.

| [-f, -force <>true>] - Force }

Performs a one-time full reallocation on a LUN, file, or volume when set to true. A forced reallocation rewrites blocks on a LUN, file, or volume unless the reallocation would result in worse performance. If you use this command without specifying this parameter, its effective value is false and a forced reallocation is not performed. If you specify this parameter without a value, it is set to true, and a forced reallocation is performed.

{ [-p, -space-optimized <>true>] - Space Optimized

Specifies that snapshot blocks are not copied to save space when set to true. If you use this command without specifying this parameter, its effective value is false and snapshot blocks are copied. However, reads from snapshots might have a slightly higher latency. If you specify this parameter without a value, it is set to true and snapshot blocks are not copied. You cannot use the `space-optimized` option with the `unshare` option.

| [-u, -unshare <>true>] - Unshare Deduplicated Blocks }

Specifies that blocks that are shared by deduplication will be unshared. This option can help remove fragmentation caused on dense volumes. This may result in increased disk usage, especially for full reallocation. You cannot use the `unshare` option with the `space-optimized` option.

{ [-t, -threshold <integer>] - Threshold

Specifies the threshold when a LUN, NVMe namespace, file, or volume is considered unoptimized and a reallocation should be performed. Once the threshold is reached, the system creates a diagnostic message that indicates that a reallocation might improve performance.

The threshold range is from 3 (the layout is moderately optimized) to 10 (the layout is not optimal). The threshold default is 4.

| [-n, -no-check <true>] - No Threshold Check }

Does not check the current layout to determine if a reallocation is needed when set to true. If you use this command without specifying this parameter, its effective value is false and the system does check the current layout to determine if a reallocation is needed. If you specify this parameter without a value, it is set to true and the system does not check the current layout to determine if a reallocation is needed.

Examples

```
cluster1::> volume reallocation start -path /vol/vol2 -interval 30m
[Job 165] Job is queued: Reallocate Job.
```

Starts a reallocation job on volume vol2 every 30 minutes.

Related Links

- [volume reallocation schedule](#)

volume reallocation stop

Stop reallocate job

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

Stops and deletes any reallocation scans on a LUN, NVMe namespace, file, or volume. This command stops and deletes in-progress, scheduled, and quiesced scans.



This command is not supported for FlexGroups or FlexGroup constituents.

Parameters

-vserver <vserver name> - Vserver

Specifies the Vserver.

-path <text> - Path

Specifies the path of the reallocation for a LUN, NVMe namespace, file, or volume.

Examples

```
cluster1::> volume reallocation stop /vol/vol2
1 entry was deleted.
```

Stops and deletes one reallocation scan on volume vol2.

volume rebalance commands

volume rebalance modify

Modify the configuration for volume capacity rebalancing operations.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume rebalance modify` command modifies the configuration for volume capacity rebalancing operations on a FlexGroup volume in the specified Vserver. When a new capacity rebalancing operation is started on a FlexGroup volume, it will use the current configuration values. Once the operation is started, any changes to the configuration will not affect the currently running capacity rebalancing operation. Only future capacity rebalancing operations will use the new configuration values.

The `volume rebalance modify` command is only supported on FlexGroup volumes.

Parameters

-vserver <vserver name> - Vserver Name

Identifies the Vserver that contains the FlexGroup volume whose capacity rebalancing configuration will be modified.

-volume <volume name> - Volume Name

Identifies the FlexGroup volume whose capacity rebalancing configuration will be modified.

[-max-runtime <time_interval>] - Maximum Runtime

Specifies the maximum time a capacity rebalancing operation will run for.

[-max-threshold <percent>] - Maximum Imbalance Threshold per Constituent

Specifies the maximum imbalance percentage for constituents. When a constituent's imbalance percentage is larger than this value, files will be moved from the constituent. The default value is 20%.

[-min-threshold <percent>] - Minimum Imbalance Threshold per Constituent

Specifies the minimum imbalance percentage for constituents. When a constituent's imbalance percentage is smaller than this value, files will not be moved from the constituent. The default value is 5%.

[-max-file-moves <integer>] - Maximum Concurrent File Moves per Constituent

Specifies the maximum number of concurrent file moves in a volume capacity rebalancing operation on a constituent of the FlexGroup volume. The default value is 25.

[-min-file-size {<integer>[KB|MB|GB|TB|PB]}] - Minimum File Size

Specifies the minimum file size to consider for a volume capacity rebalancing operation. The value must be a multiple of 4KB. The default value is 100MB, and the minimum configurable value is 20MB. Setting the minimum file size to less than the default value leads to more files being moved. Moved files use granular data, which might impact read/write I/O performance.

[-exclude-snapshots {true|false}] - Exclude Files Stuck in Snapshot Copies

Specifies whether files stuck in snapshots should be excluded in a volume capacity rebalancing operation. The default value is true.

[~~-start-time~~ <Date>] - Scheduled start time

Specifies the scheduled start time of the volume capacity rebalancing operation. The value must be a time later than 'now'.

Examples

In this example, the volume capacity rebalancing configuration is modified for the FlexGroup volume "fg1" in vservers "vs0" with the maximum threshold being 30%, the minimum threshold being 10%, the maximum file moves being 35, and the minimum file size being 200MB.

```
cluster::volume rebalance*> modify -vservers vs0 -volume fg1 -max-threshold 30 -min-threshold 10 -max-file-moves 35 -min-file-size 200MB
```

In this example, a previously scheduled volume capacity rebalancing is rescheduled for a different time.

```
cluster::volume rebalance*> modify -vservers vs0 -volume fg1 -start-time 12/30/2022 22:59:59
```

volume rebalance show

Display a list of volume capacity rebalancing operations on FlexGroup volumes

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume rebalance show` command displays information about volume capacity rebalancing operations on FlexGroup volumes.

The `volume rebalance show` command is only supported on FlexGroup volumes and constituents. The command displays rebalancing information for volumes. The command output depends on the parameter or parameters specified with the command. If no parameters are specified, the command displays the following information about all volumes:

- Vserver name
- Volume name
- Rebalancing state
- Total size
- Used size for calculating imbalance
- Size used by Snapshot copies
- Target used size
- Imbalance size
- Imbalance percent

To display detailed rebalancing information about each FlexGroup volume or constituent, run the command with the `-instance` parameter. The detailed view provides the following information for FlexGroup volumes:

- Vserver name
- Volume name
- Is constituent
- Rebalance state
- Total size
- Used size for calculating imbalance
- Size used by Snapshot copies
- Imbalance size
- Imbalance percentage
- Maximum constituent imbalance percentage
- Rebalance start time
- Rebalance stop time
- Rebalance runtime
- Rebalance maximum runtime
- Maximum imbalance threshold per constituent
- Minimum imbalance threshold per constituent
- Maximum files moved per constituent
- Minimum file size
- Exclude files stuck in snapshot copies
- Rebalance notice messages

The detailed view provides the following information for FlexGroup constituents:

- Vserver name
- Volume name
- Is constituent
- Rebalance state
- Total size
- Used size for calculating imbalance
- Size used by Snapshot copies
- Constituent target used size
- Imbalance size
- Imbalance percentage
- Moved data size
- Rebalance start time
- Rebalance stop time
- Rebalance runtime
- Rebalance maximum runtime

- Rebalance maximum runtime
- Maximum imbalance threshold per constituent
- Minimum imbalance threshold per constituent
- Maximum files moved per constituent
- Minimum file size
- Exclude files stuck in snapshot copies

You can specify parameters from the above list to display information. For example, to display information only about currently running rebalancing operations, run the command with the `-state rebalancing` parameter

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-config]

If this parameter is specified, the command displays the following information:

- max-runtime
- max-threshold
- min-threshold
- max-file-moves
- min-file-size
- exclude-snapshots

| [-time]

If this parameter is specified, the command displays the following information:

- start-time
- runtime
- max-runtime

| [-instance] }

If this parameter is specified, the command displays the additional detailed information described above.

[-vserver <vserver name>] - Vserver Name

Selects information about volume capacity rebalancing operations on FlexGroup volumes in the specified Vserver.

[-volume <volume name>] - Volume Name

Selects information about volume capacity rebalancing operations on the specified FlexGroup volume.

[-is-constituent {true|false}] - Is Constituent

If specified, the command displays information about volume capacity rebalancing operations on constituents.

[-state {idle|starting|rebalancing|rebalancing-source|rebalancing-dest|scanning|stopping|paused|not-running|unknown|scheduled}] - Rebalance State

Selects information about volume capacity rebalancing operations with the specified state. While a FlexGroup volume is rebalancing, every constituent will have a rebalancing engine that can either be scanning the filesystem for space usage and files to move, actively moving files or temporarily doing neither. If one or more constituents has a state of *rebalancing-source* or *rebalancing-dest*, then files are being moved to rebalance the FlexGroup. If no files are being moved, more information about what the rebalancing engine is doing for each constituent is available using the (privilege: advanced) `volume rebalance show -instance` command. The following values apply to FlexGroup volumes.

- *not-running* - Capacity rebalancing is not running on the volume. This will only change when the user explicitly starts rebalancing.
- *idle* - Capacity rebalancing is running, however, no active scanning or file movement is currently happening.
- *starting* - The capacity rebalancing operation is starting on the volume.
- *rebalancing* - Volume capacity rebalancing is running on the volume.
- *paused* - The capacity rebalancing operation is paused on the volume.
- *stopping* - The capacity rebalancing operation is stopping on the volume.
- *unknown* - The system was unable to determine the rebalancing state for the volume.

[-notices <text>,...] - Rebalance Notice Messages

Selects information about volume capacity rebalancing operations with the specified notice messages.

[-total {<integer>[KB|MB|GB|TB|PB]}] - Total Size

Selects information about volume capacity rebalancing operations with the specified total size.

[-used-for-imbalance {<integer>[KB|MB|GB|TB|PB]}] - Used Size for Calculating Imbalance

Selects information about volume capacity rebalancing operations with the specified used size for calculating imbalance. This is calculated by subtracting the size used by Snapshot copies, the size of files pending deletion and the size used by filesystem metadata from the volume's used size.

[-size-used-by-snapshots {<integer>[KB|MB|GB|TB|PB]}] - Volume Size Used by Snapshot Copies

Selects information about volume capacity rebalancing operations with the specified volume size used by Snapshot copies.

[-target-used {<integer>[KB|MB|GB|TB|PB]}] - Constituent Target Used Size

Selects information about volume capacity rebalancing operations with the specified target used size on a constituent. Represents the average used size among all constituents.

[-imbalance-size {<integer>[KB|MB|GB|TB|PB]}] - Imbalance Size

Selects information about volume capacity rebalancing operations with the specified imbalance size.

[-imbalance-percent <percent_no_limit_signed>] - Imbalance Percentage

Selects information about volume capacity rebalancing operations with the specified imbalance percentage.

[-data-moved {<integer>[KB|MB|GB|TB|PB]}] - Moved Data Size

Selects information about volume capacity rebalancing operations with the specified moved data size.

[-max-constituent-imbalance-percent <percent_no_limit>] - Maximum Constituent Imbalance Percentage

Selects information about volume capacity rebalancing operations with the specified maximum constituent imbalance percentage. Represents the imbalance percentage of the constituent most out of balance.

[-start-time <Date>] - Rebalance Start Time

Selects information about volume capacity rebalancing operations with the specified start time.

[-stop-time <Date>] - Rebalance Stop Time

Selects information about volume capacity rebalancing operations with the specified stop time.

[-runtime <time_interval>] - Rebalance Runtime

Selects information about volume capacity rebalancing operations with the specified run time.

[-max-runtime <time_interval>] - Rebalance Maximum Runtime

Selects information about volume capacity rebalancing operations with the specified maximum run time. Represents the maximum amount of time a rebalancing operation can run for before stopping.

[-max-threshold <percent>] - Maximum Imbalance Threshold per Constituent

Selects information about volume capacity rebalancing operations with the specified maximum threshold. Represents the acceptable imbalance threshold beyond which rebalancing should be started on that constituent in the FlexGroup volume.

[-min-threshold <percent>] - Minimum Imbalance Threshold per Constituent

Selects information about volume capacity rebalancing operations with the specified minimum threshold. Represents the minimum imbalance threshold at which rebalancing is stopped for the constituent in the FlexGroup volume.

[-max-file-moves <integer>] - Maximum Concurrent File Moves per Constituent

Selects information about volume capacity rebalancing operations with the specified maximum concurrent file moves. Represents the maximum number of concurrent file moves allowed per constituent.

[-min-file-size {<integer>[KB|MB|GB|TB|PB]}] - Minimum File Size

Selects information about volume capacity rebalancing operations with the specified minimum file size. Represents the minimum size of file to be considered for rebalancing a constituent in a FlexGroup volume.

[-exclude-snapshots {true|false}] - Exclude Files Stuck in Snapshot Copies

Selects information about volume capacity rebalancing operations with the specified excluding snapshot flag. Represents whether or not files that are in Snapshot copies as candidates for rebalancing are ignored.

Examples

The following example displays rebalancing information about all FlexGroup volumes on Vserver "vs0":

```

cluster::> volume rebalance show -vserver vs0
Vserver: vs0

```

Imbalance		Volume	State	Total	Projected	Target
Size	%				Used	Used
176KB	0%	fg1	rebalancing	400MB	117.7MB	-
8KB	0%	fg2	not-running	800MB	115.0MB	-

2 entries were displayed.

The following example displays rebalancing information about all FlexGroup constituents on Vserver "vs0":

```

cluster::> volume rebalance show -vserver vs0 -is-constituent True
Vserver: vs0

```

Imbalance		Volume	State	Total	Projected	Target
Size	%				Used	Used
176KB	0%	fg1__0001	idle	200MB	59.03MB	58.86MB
-172KB	0%	fg1__0002	idle	200MB	58.69MB	58.86MB
4KB	0%	fg2__0001	not-running	200MB	28.77MB	28.76MB
-4KB	0%	fg2__0002	not-running	200MB	28.76MB	28.76MB
4KB	0%	fg2__0003	not-running	200MB	28.77MB	28.76MB
-4KB	0%	fg2__0004	not-running	200MB	28.76MB	28.76MB

6 entries were displayed.

The following example displays detailed rebalancing information about a specific FlexGroup volume "fg1" on Vserver "vs0":

```

cluster::> volume rebalance show -vserver vs0 -volume fg1 -instance
Vserver Name: vs0
                                Volume Name: fg1
                                Is Constituent: false
                                Rebalance State: not-running
                                Rebalance Notice Messages: -
                                    Total Size: 400MB
                                Projected AFS Used Size: 117.7MB
                                Volume Size Used by Snapshot Copies: 10.57MB
                                Constituent Target Used Size: -
                                    Imbalance Size: 176KB
                                    Imbalance Percentage: 0%
                                    Moved Data Size: -
                                Maximum Constituent Imbalance Percentage: 0%
                                Rebalance Start Time: Fri Mar 25 17:26:09
2022
                                    Rebalance Stop Time: Fri Mar 25 17:29:29
2022
                                    Rebalance Runtime: 0h3m20s
                                    Rebalance Maximum Runtime: 0h3m20s
                                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 Snapshot Copies: true

```

The following example displays detailed rebalancing information about a specific FlexGroup constituent "fg1__0001" on Vserver "vs0":

```

cluster::> volume rebalance show -vserver vs0 -volume fg1__0001 -instance
-is-constituent True
Vserver Name: vs0
                                Volume Name: fg1__0001
                                Is Constituent: true
                                Rebalance State: not-running
                                Rebalance Notice Messages: -
                                Total Size: 200MB
                                Projected AFS Used Size: 59.04MB
Volume Size Used by Snapshot Copies: 5.46MB
                                Constituent Target Used Size: 58.86MB
                                Imbalance Size: 176KB
                                Imbalance Percentage: 0%
                                Moved Data Size: 0B
                                Maximum Constituent Imbalance Percentage: -
                                Rebalance Start Time: Fri Mar 25 17:26:09
2022
                                Rebalance Stop Time: Fri Mar 25 17:29:29
2022
                                Rebalance Runtime: 0h3m20s
                                Rebalance Maximum Runtime: 0h3m20s
                                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 Snapshot Copies: true

```

volume rebalance start

Start a volume capacity rebalancing operation on a FlexGroup volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume rebalance start` command starts a volume capacity rebalancing operation on a FlexGroup volume in the specified Vserver. Only one volume capacity rebalancing operation is supported per FlexGroup volume at a time. Running capacity rebalancing operations on multiple FlexGroup volumes at the same time is supported.

The `volume rebalance start` command is only supported on FlexGroup volumes. If granular data is not already enabled on the volume, it will automatically be set to *basic* when this command is run.

Parameters

-vserver <vserver name> - Vserver Name

Identifies the Vserver that contains the FlexGroup volume whose capacity will be rebalanced.

-volume <volume name> - Volume Name

Identifies the FlexGroup volume whose capacity will be rebalanced.

[-max-runtime <time_interval>] - Maximum Runtime

Specifies the maximum time a capacity rebalancing operation will run for. The default value is 6 hours.

[-max-threshold <percent>] - Maximum Imbalance Threshold per Constituent

Specifies the maximum imbalance percentage for constituents. When a constituent's imbalance percentage is larger than this value, files will be moved from the constituent. The default value is 20%.

[-min-threshold <percent>] - Minimum Imbalance Threshold per Constituent

Specifies the minimum imbalance percentage for constituents. When a constituent's imbalance percentage is smaller than this value, files will not be moved from the constituent. The default value is 5%.

[-max-file-moves <integer>] - Maximum Concurrent File Moves per Constituent

Specifies the maximum number of concurrent file moves in a volume capacity rebalancing operation on a constituent of the FlexGroup volume. The default value is 25.

[-min-file-size {<integer>[KB|MB|GB|TB|PB]}] - Minimum File Size

Specifies the minimum file size to consider for a volume capacity rebalancing operation. The value must be a multiple of 4KB. The default value is 100MB, and the minimum configurable value is 20MB. Setting the minimum file size to less than the default value leads to more files being moved. Moved files use granular data, which might impact read/write I/O performance.

[-exclude-snapshots {true|false}] - Exclude Files Stuck in Snapshot Copies

Specifies whether files stuck in snapshots should be excluded in a volume capacity rebalancing operation. The default value is true.

[-start-time <Date>] - Scheduled start time, e.g. 12/31/2022 23:59:59

Specifies the scheduled start time of the volume capacity rebalancing operation. The value must be a time later than 'now'.

Examples

In this example, a volume capacity rebalancing operation is started for the FlexGroup volume "fg1" in vserver "vs0" with the maximum run time being 72000 seconds or 20 hours starting at 23:59:59 on 12/31/2022.

```
cluster::volume rebalance*> start -vserver vs0 -volume fg1 -max-runtime
72000 -start-time 12/31/2022 23:59:59
```

volume rebalance stop

Stop a volume capacity rebalancing operation on a FlexGroup volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume rebalance stop` command stops a volume capacity rebalancing operation on a FlexGroup volume in the specified Vserver.

The `volume rebalance stop` command is only supported on FlexGroup volumes.

Parameters

-vserver <vserver name> - Vserver Name

Identifies the Vserver that contains the FlexGroup volume whose capacity rebalancing operation will be stopped.

-volume <volume name> - Volume Name

Identifies the FlexGroup volume whose capacity rebalancing operation will be stopped.

Examples

In this example, a volume capacity rebalancing operation is stopped for the FlexGroup volume "fg1" in vserver "vs0".

```
cluster::volume rebalance rebalance*> stop -vserver vs0 -volume fg1
```

volume rebalance file-move abort

Abort a file-move operation that is in progress

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume rebalance file-move abort` command aborts a file-move operation that has been started and is in progress.

A file-move operation goes through different stages as part of the move operation. During the cutover stage of the file-move operation, the actual file is created at the destination FlexGroup constituent, after which the data blocks are transferred. Once the file is created at the destination FlexGroup constituent, the file is visible in the namespace of the FlexGroup volume. If the file-move operation has entered the cutover stage, it cannot be aborted.

The `volume rebalance file-move abort` command is only supported on FlexGroup volumes. For a given FlexGroup volume, there can be many file-move operations that are currently in progress for different sets of constituents. The parameters for this command provide a way to identify a particular file-move operation within a FlexGroup volume.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

Identifies the Vserver that contains the FlexGroup volume involved in the file-move operation.

-volume <volume name> - Volume Name (privilege: advanced)

Identifies the FlexGroup volume that is hosting the file being moved in the file-move operation.

-path <text> - Path (privilege: advanced)

Identifies the path to the file being moved within the FlexGroup volume.

-destination-constituent <volume name> - Destination Constituent Volume (privilege: advanced)

Identifies the destination constituent name of the FlexGroup volume where the file is being moved.

[-operation-uuid <UUID>] - Operation UUID (privilege: advanced)

If specified, this parameter identifies the UUID of the file-move operation. Use the [volume rebalance file-move show](#) command to obtain the UUID for a file-move operation.

Examples

The following example aborts a file-move operation for the file `"/system/build/image.tgz"` within the FlexGroup volume `"bld_volume"`:

```
cluster1::*> volume rebalance file-move abort -vserver vs0 -volume
bld_volume -path /system/build/image.tgz -destination-constituent
bld_volume__0005

File-move operation cancelled.
```

The following example aborts the file-move operation in progress for the file `"/system/build/image.tgz"` within the FlexGroup `"bld_volume"`. The UUID of the file-move operation is also provided:

```
cluster1::*> volume rebalance file-move abort -vserver vs0 -volume
bld_volume -path /system/build/image.tgz -destination-constituent
bld_volume__0005 -operation-uuid f8d780b4-32fc-4053-be5e-2f6edc0a652b

File-move operation cancelled.
```

The following example attempts to abort the file-move operation for the file `"/system/build/image.tgz"` within the FlexGroup volume `"bld_volume"`. However, the file-move operation could not be aborted as it is past the cutover stage of the operation. A file-move operation can only be aborted if it has not entered the cutover stage:


```
cluster1::*> volume rebalance file-move abort -vserver vs0 -volume
bld_volume -path /system/build/image.tgz -destination-constituent
bld_volume__0005
```

```
Error: command failed: The file move abort operation failed. Reason:
Cannot
    destroy file operation. File move operation has completed the
cutover
    phase.
```

Related Links

- [volume rebalance file-move show](#)

volume rebalance file-move modify

Modify a file-move operation

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume rebalance file-move modify` command modifies certain attributes of a file-move operation that is currently in progress. Specifically, the maximum throughput allowed on a particular file-move operation can be modified, and the data transfer of a file-move operation can be paused or resumed.

The `volume rebalance file-move modify` command is only supported on FlexGroup volumes. For a given FlexGroup volume, there can be many file-move operations that are currently in progress for different sets of constituents. The parameters for this command provide a way to identify a particular file-move operation within a FlexGroup volume.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

Identifies the Vserver that contains the FlexGroup volume involved in the file-move operation.

-volume <volume name> - Volume Name (privilege: advanced)

Identifies the FlexGroup volume that is hosting the file being moved in the file-move operation.

-path <text> - Path (privilege: advanced)

Identifies the path to the file being moved within the FlexGroup volume. This parameter can be specified with the `-destination-constituent` parameter to identify a particular file-move operation within a FlexGroup volume.

-destination-constituent <volume name> - Destination Constituent Volume (privilege: advanced)

Identifies the destination constituent name of the FlexGroup volume where the file is being moved. This parameter can be specified with the `-path` parameter to identify a particular file-move operation within a FlexGroup volume.

[`-operation-uuid <UUID>`] - Operation UUID (privilege: advanced)

Identifies the UUID of the file-move operation. Use the [volume rebalance file-move show](#) command to obtain the UUID for a file-move operation. This parameter can be specified in addition to the `-path` and `-destination-constituent` parameters to identify a particular file-move operation if there are multiple ongoing operations involving the same path and destination constituent.

[`-max-throughput {<integer>[KB|MB|GB|TB|PB]}`] - Maximum Scanner Speed (privilege: advanced)

If specified, this parameter will modify the rate limit for the data transfer in bytes per second. A non-zero value less than 1 MB/s will be set to 1 MB/s. A non-zero value greater than 1 MB/s will be truncated to the nearest integral megabyte value. A value of "0" specifies that no range will be set for the data transfer. This value does not represent a throughput the system will guarantee, but a limit that cannot be exceeded.

[`-scanner-paused {true|false}`] - Scanner Paused (privilege: advanced)

If specified, the data transfer of the file-move operation will be paused or resumed. Using a value of `false` will pause the data transfer while using a value of `true` will resume the data transfer.

Examples

The following example modifies the maximum throughput of the ongoing file-move operation for the file `/system/build/image.tgz` within the FlexGroup volume `"bld_volume"`. `-path` and `-destination-constituent` are used to identify the particular file-move operation:

```
cluster1::*> volume rebalance file-move modify -vserver vs0 -volume
bld_volume -path /system/build/image.tgz -destination-constituent
bld_volume__0005 -max-throughput 2MB
```

The following example pauses the data transfer of the ongoing file-move operation for the file `/system/build/image.tgz` within the FlexGroup volume `"bld_volume"`. This stops any further data from being transferred from the source volume to the destination volume:

```
cluster1::*> volume rebalance file-move modify -vserver vs0 -volume
bld_volume -path /system/build/image.tgz -destination-constituent
bld_volume__0005 -scanner-paused true
```

The following example resumes the data transfer of a paused file-move operation for the file `/system/build/image.tgz` within the FlexGroup volume `"bld_volume"`. `-path` and `-destination-constituent` are used to identify the particular file-move operation:

```
cluster1::*> volume rebalance file-move modify -vserver vs0 -volume
bld_volume -path /system/build/image.tgz -destination-constituent
bld_volume__0005 -scanner-paused false
```

Related Links

- [volume rebalance file-move show](#)

volume rebalance file-move show

Display a list of files being moved

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume rebalance file-move show` command displays a list of files being moved between the FlexGroup Constituents.

The `volume rebalance file-move show` command is only supported on FlexGroup volumes. The command displays information about each ongoing file-move operation, as well as any file-move operations that have completed, failed or have been cancelled within the last two minutes. For all file-move operations matching the specified parameters, the command displays the following information:

- Vserver name
- Volume name
- File Path
- Source Constituent Volume
- Destination Constituent Volume
- Current status of the file-move operation ("Preparing", "Data-Transfer", "Paused", "Complete", "Failed" or "Terminated-Manual")
- Percentage of total data bytes transferred to the destination

To display detailed information about each file-move operation, run the command with the `-instance` parameter. The detailed view provides the following information:

- Vserver Name
- Volume Name
- Destination Constituent Volume
- Source Constituent Volume
- Destination Constituent Node
- Source Constituent Node
- Operation Status: Current status of the file-move operation ("Preparing", "Data-Transfer", "Paused", "Complete", "Failed" or "Terminated-Manual")
- Scanner Progress Percentage
- Operation UUID: UUID of the overall file-move operation
- Source Operation UUID: UUID used to identify the file-move operation on the source constituent
- File Path: Path of the file being moved
- Source Constituent Master Data Set ID
- Source Constituent Data Set ID
- Source File ID: ID of the inode on the source constituent
- Source Generation: Generation number of the inode on the source constituent

- Destination File ID: ID of the inode on the destination constituent
- Destination Data Set ID
- Destination Master Data Set ID
- Destination Generation: Generation number of the inode on the destination constituent
- Maximum Scanner Speed (per sec): Maximum throughput of the data scanner
- Scanner Paused: Displays information if the underlying data transfer scanner is paused
- Scanner Status: The current stage of the underlying data transfer scanner ("Preparing", "Allocation-Map", "Data", "Destroying", "Paused-Manual", "Paused-Error", "Complete", "Terminated-Manual" or "Destroyed")
- Scanner Progress: Number of bytes transferred
- Scanner Total: Total number of bytes to be transferred
- Data Scanner Priority: Priority of the data scanner assigned to the file-move operation (high or low)
- Elapsed Time
- Cutover Time: Time elapsed during the cutover stage of the file-move operation, in seconds
- Is Snapshot Fenced: Whether snapshots are fenced for the file-move operation
- Is Destination Ready: Whether the destination volume is ready for the file-move operation
- Last Failure Time
- Last Failure Reason

You can specify parameters from the above list to display information. For example, to display information only about paused operations, run the command with the `-status Paused` parameter.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If this parameter is specified, the command displays the additional detailed information described above.

[-vserver <vserver name>] - Vserver (privilege: advanced)

If this parameter and the `-volume` parameter are specified, the command displays the file-move operations on the specified volume. If this parameter is specified by itself, the command displays the file-move operations on all volumes on the specified Vserver.

[-volume <volume name>] - Volume Name (privilege: advanced)

If this parameter and the `-vserver` parameter are specified, the command displays the file-move operations on the specified volume. If this parameter is specified by itself, the command displays the file-move operations on all volumes matching the specified name.

[-destination-constituent <volume name>] - Destination Constituent Volume (privilege: advanced)

If this parameter is specified, the command displays only the file-move operations with the specified destination FlexGroup constituent name.

[-operation-uuid <UUID>] - Operation UUID (privilege: advanced)

If this parameter is specified, the command displays only the file-move operation with the specified operation UUID.

[-destination-node <nodename>] - Destination Constituent Node (privilege: advanced)

If this parameter is specified, the command displays only file-move operations for files being moved to the specified destination node.

[-destination-dsid <integer>] - Destination Constituent DSID (privilege: advanced)

If this parameter is specified, the command displays only file-move operations with the specified destination volume data set ID.

[-source-constituent <volume name>] - Source Constituent Volume (privilege: advanced)

If this parameter is specified, the command displays only file-move operations with the specified source FlexGroup constituent name.

[-source-operation-uuid <UUID>] - Source Operation UUID (privilege: advanced)

If this parameter is specified, the command displays only the file-move operation with the specified internal source job operation UUID.

[-source-node <nodename>] - Source Constituent Node (privilege: advanced)

If this parameter is specified, the command displays only file-move operations for files being moved from the specified source node.

[-source-dsid <integer>] - Source Constituent DSID (privilege: advanced)

If this parameter is specified, the command displays only file-move operations with the specified source volume data set ID.

[-status {Preparing|Data-Transfer|Paused|Complete|Failed|Manually-Terminated}] - Operation Status (privilege: advanced)

If this parameter is specified, the command displays only file-move operations with the specified status ("Preparing", "Data-Transfer", "Paused", "Complete", "Failed" or "Terminated-Manual").

[-scanner-percent <percent>] - Scanner Progress Percentage (privilege: advanced)

If this parameter is specified, the command displays only file-move operations that have transferred the specified percentage of total bytes.

[-last-failure-reason <text>] - Last Failure Reason (privilege: advanced)

If this parameter is specified, the command displays only file-move operations whose last failure reason matches the specified string.

[-path <text>] - File Path (privilege: advanced)

If this parameter is specified, the command displays only file-move operations for files that match the specified path.

[-destination-msid <integer>] - Destination Constituent MSID (privilege: advanced)

If this parameter is specified, the command displays only file-move operations with the specified destination volume master data set ID.

[-destination-fileid <integer>] - Destination Fileid (privilege: advanced)

If this parameter is specified, the command displays only file-move operations pertaining to files with the specified file ID on the destination constituent.

[-destination-generation <integer>] - Destination Generation (privilege: advanced)

If this parameter is specified, the command displays only file-move operations pertaining to files with the specified generation number on the destination constituent.

[-source-msid <integer>] - Source Constituent MSID (privilege: advanced)

If this parameter is specified, the command displays only file-move operations with the specified source volume master data set ID.

[-source-fileid <integer>] - Source Fileid (privilege: advanced)

If this parameter is specified, the command displays only file-move operations pertaining to files with the specified file ID on the source constituent.

[-source-generation <integer>] - Source Generation (privilege: advanced)

If this parameter is specified, the command displays only file-move operations pertaining to files with the specified generation number on the source constituent.

[-max-throughput {<integer>[KB|MB|GB|TB|PB] }] - Maximum Scanner Speed (per sec) (privilege: advanced)

If this parameter is specified, the command displays only file-move operations that use the specified value as the maximum scanner speed (in MB/s). This value does not represent a throughput the system will guarantee, but a limit that cannot be exceeded.

[-scanner-paused {true|false}] - Scanner Paused (privilege: advanced)

If this parameter is specified, the command displays only file-move operations that are paused (true) or unpaused (false).

[-scanner-status {Preparing|Allocation-Map|Data|Destroying|Paused-Manual|Paused-Error|Complete|Destroyed|Terminated-Manual}] - Scanner Status (privilege: advanced)

If this parameter is specified, the command displays only file-move operations in the specified scanner stage ("Preparing", "Allocation-Map", "Data", "Destroying", "Paused-Manual", "Paused-Error", "Complete", "Terminated-Manual" or "Destroyed")

[-scanner-progress {<integer>[KB|MB|GB|TB|PB] }] - Scanner Progress (privilege: advanced)

If this parameter is specified, the command displays only file-move operations that have transferred the specified number of bytes.

[-scanner-total {<integer>[KB|MB|GB|TB|PB] }] - Scanner Total (privilege: advanced)

If this parameter is specified, the command displays only file-move operations with the specified total transfer size in bytes.

[-data-scanner-priority {high|low}] - Data Scanner Priority (privilege: advanced)

If this parameter is specified, the command displays only the file-move operations that give the data scanner high priority (high) or low priority (low).

[-elapsed-time <time_interval>] - Elapsed Time (privilege: advanced)

If this parameter is specified, the command displays only file-move operations that have a matching

elapsed-time interval. Specify an elapsed-time range by using the ".." operator between two values. For example, the following command displays file-move operations with elapsed-time between 2 hours and 3 hours 30 minutes:

```
volume rebalance file-move show -elapsed-time 2h..3h30m
```

+

Specify a comparative elapsed-time value using the ">" and "<" operators. For example, the following command displays the file-move operations with an elapsed-time greater than 5 hours and 30 minutes:

```
volume rebalance file-move show -elapsed-time >5h30m
```

[`-cutover-time` <[<integer>d] [<integer>h] [<integer>m] [<integer>s]>] - Cutover Time (privilege: advanced)

If this parameter is specified, the command displays only the file-move operations with the specified cutover time.

[`-is-snapshot-fenced` {true|false}] - Is Snapshot Fenced (privilege: advanced)

If this parameter is specified, the command displays only the file-move operations with snapshots fenced (true) or unfenced (false).

[`-is-destination-ready` {true|false}] - Is Destination Ready (privilege: advanced)

If this parameter is specified, the command displays only the file-move operations for which the destination volume is ready (true) or not ready (false).

[`-last-failure-time` <MM/DD/YYYY HH:MM:SS>] - Last Failure Time (privilege: advanced)

If this parameter is specified, the command displays file-move operations that have a matching last-failure-time value. Use the format: "MM/DD/YYYY HH:MM:SS". Specify a time range by using the ".." operator between two time statements. For example, the following command displays file-move operations with a last-failure-time value between "08/13/2014 05:55:00" and "08/13/2014 06:10:00":

```
volume rebalance file-move show -last-failure-time "08/13/2021
05:55:00".."08/13/2021 06:10:00"
```

+

Specify a comparative last-failure-time value using the ">" and "<" operators. For example, the following command displays file-move operations with a last-failure-time value after "8/27/2021 14:58:00":

```
volume rebalance file-move show -last-failure-time >"8/27/2021 14:58:00"
```

Examples

The following example displays information about all file-move operations on the Vserver named "vs1":

```

cluster1::volume rebalance file-move*> show -vserver vs1
Vserver: vs1
Volume: fg1
Path          Source          Destination    Status
Progress
-----
-----
/d1/d2/foo    fg1__0002         fg1__0003     Paused
0%
/d1/d2/bar    fg1__0002         fg1__0004     Completed
100%

Vserver: vs1
Volume: bld_volume
Path          Source          Destination    Status
Progress
-----
-----
/system/build/image.tgz  bld_volume__0001
                               bld_volume__0005
                               Data-Transfer
50%
3 entries were displayed.

```

The following example displays detailed information about a specific file-move operation on Vserver vs0:


```

cluster1::volume rebalance file-move*> show -instance -vserver vs0 -volume
fg -path /test
  (volume rebalance file-move show)
Vserver Name: vs0
      Volume Name: fg
Destination Constituent Volume: fg__0001
      Operation UUID: 283a983e-c06a-42bc-8bd7-8830211cb2a9
Destination Constituent Node: sti96-vsimsim-ucs539i
Destination Constituent DSID: 1635
      Source Constituent Volume: fg__0004
      Source Operation UUID: 283a983e-c06a-42bc-8bd7-8830211cb2a9
      Source Constituent Node: sti96-vsimsim-ucs539j
      Source Constituent DSID: 1638
      Operation Status: Paused
Scanner Progress Percentage: 0%
      Last Failure Reason: -
      File Path: /test
Destination Constituent MSID: 2154447497
      Destination Fileid: 97
      Destination Generation: 247243720
      Source Constituent MSID: 2154447500
      Source Fileid: 504
      Source Generation: 214935427
Maximum Scanner Speed (per sec): 1MB
      Scanner Paused: true
      Scanner Status: Paused-Manual
      Scanner Progress: 0B
      Scanner Total: 10MB
Data Scanner Priority: low
      Elapsed Time: 0h2m27s
      Cutover Time: 0s
      Is Snapshot Fenced: false
      Is Destination Ready: true
      Last Failure Time: -

```

volume rebalance file-move start

Moves a file from one FlexGroup constituent to another

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume rebalance file-move start` command initiates a retroactive file movement for a file residing in a FlexGroup volume in the given Vserver. You must provide the path to the file to be moved. Additional parameters may be used to control certain properties of the move operation. The parent directory's

mtime may be modified by this command which may impact analytics mtime histogram. This command only takes NFS filenames and only UTF-8 encoding is supported.

The `volume rebalance file-move start` command is only supported on FlexGroup volumes.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

Identifies the Vserver that contains the FlexGroup volume involved in this operation.

-volume <volume name> - Volume Name (privilege: advanced)

Identifies the FlexGroup volume that is hosting the file being moved in this operation.

-path <text> - Path (privilege: advanced)

Path to the file being moved within this FlexGroup volume.

[-destination-constituent <volume name>] - Destination Constituent Volume (privilege: advanced)

Identifies the destination constituent name of the FlexGroup volume where the file is being moved to. The constituent must be a member of the same FlexGroup volume as the source constituent.

[-max-throughput {<integer>[KB|MB|GB|TB|PB]}] - Maximum Scanner Speed (privilege: advanced)

Specifies the range of the data transfer in bytes per second. A non-zero value less than 1 MB/s will be set to 1 MB/s. A non-zero value greater than 1 MB/s will be truncated to the nearest integral megabyte value. If unspecified, the default value is "0" which means no range will be set for the data transfer. This value does not represent a throughput the system will guarantee, but a limit that cannot be exceeded.

[-scanner-paused {true|false}] - Scanner Paused (privilege: advanced)

Specifies that the data scanner should not be started immediately when the move operations starts. If unspecified, the default value is "false". Use the `volume rebalance file-move modify` command to start the data scanner.

[-force <true>] - Force (privilege: advanced)

If this parameter is specified, the file move operation will break the existing lock state on the file and proceed. This may cause a disruption for some client applications.

[-is-disruptive {true|false}] - Is Disruptive (privilege: advanced)

Specifies whether the file move operation should be disruptive. The default value is "false".

Examples

In this example, a file-move operation is started for the file `"/d1/d2/foo"` within the FlexGroup volume `"fg1"` in vserver `"vs0"`. The path is relative to the root directory of this FlexGroup volume, not the global namespace path of the file. In this example, the file is being moved to the destination constituent identified by the name `"fg1__0003"`.

```
cluster::volume rebalance file-move*> start -vserver vs0 -volume fg1 -path
/d1/d2/foo -destination-constituent fg1__0003
File move started with operation-uuid f8d780b4-32fc-4053-be5e-2f6edc0a652b
```

Related Links

- [volume rebalance file-move modify](#)

volume rebalance file-move statistics reset

Reset statistics for file-move operations between FlexGroup constituents

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume rebalance file-move statistics reset` command resets completed statistics for the file-move operations.

The `volume rebalance file-move statistics reset` command is only supported on FlexGroup volumes.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

If this parameter and the `-volume` parameter are specified, the command resets statistics for the FlexGroup constituents on the specified volume. If this parameter is specified by itself, the command resets statistics for the FlexGroup constituents on all volumes on the specified Vserver.

-volume <volume name> - Volume Name (privilege: advanced)

If this parameter and the `-vserver` parameter are specified, the command resets statistics for the FlexGroup constituents on the specified volume. If this parameter is specified by itself, the command resets statistics for the FlexGroup constituents on all volumes matching the specified name.

Examples

The following example resets completed statistics for all FlexGroup constituents on the FlexGroup volume "bld_volume" in Vserver "vs0":

```
cluster1::*> volume rebalance file-move statistics show -vserver vs0
-volume bld_volume

Vserver: vs0
Volume: bld_volume

                Source
Constituent      Files      Bytes      Destination
Bytes
-----
bld_volume__0001
    In Progress:      1      10GB      0
0B
    Completed:       5     100GB      0
0B
```

```

bld_volume__0002
    In Progress:           0           0B           1
10GB
    Completed:            0           0B           10
150MB
bld_volume__0003
    In Progress:           1          10GB           0
0B
    Completed:            15          100GB          0
0B
bld_volume__0004
    In Progress:           0           0B           1
10GB
    Completed:            0           0B           10
150MB
bld_volume__0005
    In Progress:           0           0B           1
10GB
    Completed:            0           0B           10
150MB
bld_volume__0006
    In Progress:           0           0B           1
10GB
    Completed:            0           0B           10
150MB
bld_volume__0007
    In Progress:           0           0B           1
10GB
    Completed:            0           0B           10
150MB
bld_volume__0008
    In Progress:           0           0B           1
10GB
    Completed:            0           0B           10
150MB

cluster1::*> volume rebalance file-move statistics reset -vserver vs0
-volume bld_volume
( volume rebalance file-move statistics reset)

cluster1::*> volume rebalance file-move statistics show -vserver vs0
-volume bld_volume

Vserver: vs0
Volume: bld_volume

Source                                     Destination

```

Constituent	Files	Bytes	Files
Bytes			
-----	-----	-----	-----

bld_volume__0001			
In Progress:	1	10GB	0
0B			
Completed:	0	0B	0
0B			
bld_volume__0002			
In Progress:	0	0B	1
10GB			
Completed:	0	0B	0
0B			
bld_volume__0003			
In Progress:	1	10GB	0
0B			
Completed:	0	0B	0
0B			
bld_volume__0004			
In Progress:	0	0B	1
10GB			
Completed:	0	0B	0
0B			
bld_volume__0005			
In Progress:	0	0B	1
10GB			
Completed:	0	0B	0
0B			
bld_volume__0006			
In Progress:	0	0B	1
10GB			
Completed:	0	0B	0
0B			
bld_volume__0007			
In Progress:	0	0B	1
10GB			
Completed:	0	0B	0
0B			
bld_volume__0008			
In Progress:	0	0B	1
10GB			
Completed:	0	0B	0
0B			

volume rebalance file-move statistics show

Display statistics for file-move operations across FlexGroup Constituents

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume rebalance file-move statistics show` command displays common statistics for the file-move operations at a FlexGroup constituent level.

The `volume rebalance file-move statistics show` command is only supported on FlexGroup volumes. For all FlexGroup constituents matching the specified parameters, the command displays the following information:

- Vserver name
- Volume name
- Constituent name
- Source Files In-Progress: Number of files currently being moved from the FlexGroup constituent
- Source Files Completed: Number of files already moved from the FlexGroup constituent
- Source Bytes In-Progress: Total number of bytes of all the files currently being moved from the FlexGroup constituent
- Source Bytes Completed: Total number of bytes of all the files already moved from the FlexGroup constituent
- Destination Files In-Progress: Number of files currently being moved to the FlexGroup constituent
- Destination Files Completed: Number of files already moved to the FlexGroup constituent
- Destination Bytes In-Progress: Total number of bytes of all the files currently being moved to the FlexGroup constituent
- Destination Bytes Completed: Total number of bytes of all the files already moved to the FlexGroup constituent

You can specify parameters from the above list to display information. For example, to display statistics only for FlexGroup constituents with more than one file currently being moved from them, run the command with the `-source-files-inprogress >1` parameter.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver (privilege: advanced)

If this parameter and the `-volume` parameter are specified, the command displays statistics for the FlexGroup constituents on the specified volume. If this parameter is specified by itself, the command displays statistics for the FlexGroup constituents on all volumes on the specified Vserver.

[`-volume <volume name>`] - Volume Name (privilege: advanced)

If this parameter and the `-vserver` parameter are specified, the command displays statistics for the FlexGroup constituents on the specified volume. If this parameter is specified by itself, the command displays statistics for the FlexGroup constituents on all volumes matching the specified name.

[`-constituent <volume name>`] - Constituent Name (privilege: advanced)

If this parameter is specified, the command displays statistics only for FlexGroup constituents matching the specified name.

[`-source-files-inprogress <integer>`] - Source Files In Progress (privilege: advanced)

If this parameter is specified, the command displays statistics only for FlexGroup constituents which have the specified number of files currently being moved from them.

[`-source-files-completed <integer>`] - Source Files Completed (privilege: advanced)

If this parameter is specified, the command displays statistics only for FlexGroup constituents which have the specified number of files already moved from them.

[`-source-bytes-inprogress {<integer>[KB|MB|GB|TB|PB]}`] - Source Bytes in Progress (privilege: advanced)

If this parameter is specified, the command displays statistics only for FlexGroup constituents which have the specified number of bytes currently being moved from them (across all files currently being moved from them).

[`-source-bytes-completed {<integer>[KB|MB|GB|TB|PB]}`] - Source Bytes Completed (privilege: advanced)

If this parameter is specified, the command displays statistics only for FlexGroup constituents which have the specified number of bytes already moved from them (across all files already moved from them).

[`-destination-files-inprogress <integer>`] - Destination Files In Progress (privilege: advanced)

If this parameter is specified, the command displays statistics only for FlexGroup constituents which have the specified number of files currently being moved to them.

[`-destination-files-completed <integer>`] - Destination Files Completed (privilege: advanced)

If this parameter is specified, the command displays statistics only for FlexGroup constituents which have the specified number of files already moved to them.

[`-destination-bytes-inprogress {<integer>[KB|MB|GB|TB|PB]}`] - Destination Bytes In Progress (privilege: advanced)

If this parameter is specified, the command displays statistics only for FlexGroup constituents which have the specified number of bytes currently being moved to them (across all files currently being moved to them).

[`-destination-bytes-completed {<integer>[KB|MB|GB|TB|PB]}`] - Destination Bytes Completed (privilege: advanced)

If this parameter is specified, the command displays statistics only for FlexGroup constituents which have the specified number of bytes already moved to them (across all files already moved to them).

Examples

The following example displays statistics for all FlexGroup constituents on the FlexGroup volume "bld_volume"

in Vserver "vs0":

```
cluster1::*> volume rebalance file-move statistics show -vserver vs0  
-volume bld_volume
```

Vserver: vs0

Volume: bld_volume

Constituent Bytes	Source		Destination
	Files	Bytes	Files
----- ----- bld_volume__0001			
In Progress:	1	10GB	0
0B			
Completed:	5	20GB	0
0B			
bld_volume__0002			
In Progress:	0	0B	1
10GB			
Completed:	0	0B	10
40GB			
bld_volume__0003			
In Progress:	1	5GB	0
0B			
Completed:	15	100GB	0
0B			
bld_volume__0004			
In Progress:	0	0B	1
5GB			
Completed:	0	0B	10
80GB			
bld_volume__0005			
In Progress:	0	0B	0
0B			
Completed:	0	0B	0
0B			
bld_volume__0006			
In Progress:	0	0B	0
0B			
Completed:	0	0B	0
0B			
bld_volume__0007			
In Progress:	0	0B	0
0B			
Completed:	0	0B	0


```

0B
bld_volume__0008
    In Progress:           0           0B           0
0B
    Completed:            0           0B           0
0B

```

The following example displays statistics for all FlexGroup constituents on the FlexGroup volume "bld_volume" in Vserver "vs0" with more than 4 files already sent from them:

```

cluster1::*> volume rebalance file-move statistics show -vserver vs0
-volume bld_volume -source-files-completed >4

Vserver: vs0
Volume: bld_volume

                Source                Destination
Constituent      Files      Bytes      Files
Bytes
-----
bld_volume__0001
    In Progress:           1      10GB           0
0B
    Completed:            5      20GB           0
0B
bld_volume__0003
    In Progress:           1       5GB           0
0B
    Completed:           15     100GB           0
0B

```

volume recovery-queue commands

volume recovery-queue modify

Modify attributes of volumes in the recovery queue

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume recovery-queue modify` command modifies the retention period of a volume in the recovery queue.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

This specifies the Vserver on which the volume is located.

-volume <volume name> - Volume Name (privilege: advanced)

This specifies the volume that is to be modified.

[-retention-hours <integer>] - Volume Delete Retention (privilege: advanced)

This specifies the time in hours that the volume will be available in the recovery queue for recovery. For example, a value of "10" means that the volume will be retained in the recovery queue for at least 10 hours. The maximum allowed Volume retention-hours is 4294967295 hours.

Examples

The following example modifies a volume named vol1_1234 on a Vserver named vs1 in the recovery queue.

```
cluster1::*> volume recovery-queue modify -vserver vs1 -volume vol1_1234
-retention-hours 10
```

volume recovery-queue purge-all

Purge all volumes from the recovery queue belonging to a Vserver

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume recovery-queue purge-all` command purges all volumes from the recovery queue belonging to a Vserver.

Parameters

-vserver <vserver name> - Vserver name (privilege: advanced)

This specifies the Vserver whose volumes in the recovery queue need to be purged.

Examples

The following example purges all volumes on a Vserver named vs1 from the recovery queue.

```
cluster1::*> volume recovery-queue purge-all -vserver vs1
```

volume recovery-queue purge

Purge volumes from the recovery queue belonging to a Vserver

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume recovery-queue purge` command purges a volume from the recovery queue belonging to a Vserver.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

This specifies the Vserver on which the volume is located.

-volume <volume name> - Volume Name (privilege: advanced)

This specifies the volume that is to be purged from the recovery queue.

Examples

The following example purges a volume named `vol1_1234` on a Vserver named `vs1` from the recovery queue.

```
cluster1::*> volume recovery-queue purge -vserver vs1 -volume vol1_1234
```

volume recovery-queue recover-all

Recover all volumes from the recovery queue belonging to a Vserver

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume recovery-queue recover-all` command recovers all the volumes from the recovery-queue belonging to a Vserver.



This command will restore volumes with their previous access types and volumes will remain offline. This command will not restore the volumes with their previous name, junction, LUN mappings, NVMe namespace mappings, quota policy rules and snapshot schedules.

Parameters

-vserver <vserver name> - Vserver name (privilege: advanced)

This specifies the Vserver for which the volumes need to be recovered.

Examples

The following example recovers all volumes on a Vserver named `vs1` from the recovery queue.

```
cluster1::*> volume recovery-queue recover-all -vserver vs1
```

volume recovery-queue recover

Recover volumes from the recovery queue belonging to a Vserver

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume recovery-queue recover` command recovers a volume from the recovery queue belonging to a Vserver.



This command will restore the volume with its previous access type and the volume will remain offline. This command will not restore the volume with its previous name, junction, LUN mappings, NVMe namespace mappings, quota policy rules and snapshot schedules.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

This specifies the Vserver on which the volume is located.

-volume <volume name> - Volume Name (privilege: advanced)

This specifies the volume that is to be recovered.

Examples

The following example recovers a volume named `vol1_1234` on a Vserver named `vs1` from the recovery queue.

```
cluster1::*> volume recovery-queue recover -vserver vs1 -volume vol1_1234
```

volume recovery-queue show

Show volumes in the recovery queue

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* privilege level.

Description

The `volume recovery-queue show` command displays information about volumes in the recovery queue. The command output depends on the parameter(s) specified with the command. If no parameters are specified, the command displays the following information about all volumes:

- Vserver name
- Volume name
- Deletion Request Time
- Retention Hours

To display information about a single volume, run the command with the `-vserver` and `-volume` parameters.

To display detailed information about all volumes, run the command with the `-instance` parameter.

You can specify additional parameters to display information that matches only those parameters. For example, to display information only about volumes with retention hours 10, run the command with the

`-retention-hours 10` parameter.

Parameters

{ [-fields <fieldname>,...]

This specifies the fields that need to be displayed. The fields `Vserver` and `Volume` are the default fields.

[[-instance]]

If this parameter is specified, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Name (privilege: advanced)

If this parameter and the `-volume` parameter are specified, the command displays detailed information about the specified volume. If this parameter is specified by itself, the command displays information about volumes on the specified Vserver.

[-volume <volume name>] - Volume Name (privilege: advanced)

If this parameter and the `-vserver` parameter are specified, the command displays detailed information about the specified volume. If this parameter is specified by itself, the command displays information about all volumes matching the specified name.

[-delete-time <Date>] - Deletion Request Time (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that match the specified delete time.

[-retention-hours <integer>] - Volume Delete Retention (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that match the specified retention hours.

[-instance-uuid <UUID>] - Volume Instance UUID (privilege: advanced)

If this parameter is specified, the command displays information only about the volume or volumes that match the specified instance uuid.

Examples

The following example displays information about all volumes on a Vserver named `vs1` in the recovery queue:

```
cluster1::*> volume recovery-queue show -vserver vs1
Vserver   Volume           Deletion Request Time  Retention Hours
-----
vs1       vol1_1234        Wed Aug 22 06:50:19 2012      10
vs1       vol2_1235        Wed Aug 22 06:50:26 2012      12
2 entries were displayed.
```

The following example displays detailed information about a volume named `vol1_1234` on a Vserver named `vs1` in the recovery queue:

```
cluster1::*> volume recovery-queue show -vserver vs1 -volume vol1_1234
Vserver Name: vs1
    Volume Name: vol1_1234
    Deletion Request Time: Wed Aug 22 06:50:19 2012
Volume Delete Retention: 10
```

volume schedule-style commands

volume schedule-style prepare-to-downgrade

Disables volume schedule style feature and sets schedule style to default (create-time)

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command will disable the volume schedule style feature and set schedule style to default (create-time).

Examples

The following example prepares the schedule-style on all volumes for revert/downgrade.

```
cluster1::*> volume schedule-style prepare-to-downgrade
```

volume snaplock commands

volume snaplock modify

Modify SnapLock attributes of a SnapLock volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

```
The ` volume snaplock modify ` command modifies one or more SnapLock
attributes of a SnapLock volume.
```

Parameters

-vserver <vserver name> - Vserver

This specifies the vserver which owns the required SnapLock volume.

-volume <volume name> - Volume

This specifies the SnapLock volume whose attribute needs to be modified.

[`-minimum-retention-period` {<integer> `seconds`|`minutes`|`hours`|`days`|`months`|`years`} | `infinite`}] - Minimum Retention Period

Specifies the minimum allowed retention period for files committed to WORM state on the volume. Any files committed with a retention period shorter than this minimum value, is assigned this minimum value.

If this option value is *infinite*, then every file committed to the volume will have a retention period that never expires.

Otherwise, the retention period is specified as a number followed by a suffix. The valid suffixes are seconds, minutes, hours, days, months, and years. For example, a value of 6months represents a retention period of 6 months. The maximum allowed retention period is 70 years. This option is not applicable while extending retention period of an already committed WORM file.

[`-default-retention-period` {<integer> `seconds`|`minutes`|`hours`|`days`|`months`|`years`} | `min` | `max` | `infinite` | `unspecified`}] - Default Retention Period

Specifies the default retention period that is applied to files while committing to WORM state without an associated retention period.

If this option value is *min*, then `minimum-retention-period` is used as the default retention period. If this option value is *max*, then `maximum-retention-period` is used as the default retention period. If this option value is *infinite*, then a retention period that never expires will be used as the default retention period. If this option value is *unspecified*, then the file will be retained forever; however, the retention time can be changed to an absolute value.

The retention period can also be explicitly specified as a number followed by a suffix. The valid suffixes are seconds, minutes, hours, days, months, and years. For example, a value of 6months represents a retention period of 6 months. The maximum valid retention period is 70 years.

[`-maximum-retention-period` {<integer> `seconds`|`minutes`|`hours`|`days`|`months`|`years`} | `infinite`}] - Maximum Retention Period

Specifies the maximum allowed retention period for files committed to WORM state on the volume. Any files committed with a retention period longer than this maximum value, is assigned this maximum value.

If this option value is *infinite*, then files that have retention period that never expires might be committed to the volume.

Otherwise, the retention period is specified as a number followed by a suffix. The valid suffixes are seconds, minutes, hours, days, months, and years. For example, a value of 6months represents a retention period of 6 months. The maximum allowed retention period is 70 years. This option is not applicable while extending retention period of an already committed WORM file.

[`-autocommit-period` {<integer> `minutes`|`hours`|`days`|`months`|`years`} | `none`}] - Autocommit Period

Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater the autocommit period of the volume are committed to WORM state.

The autocommit period option is specified as a number followed by a suffix. The valid suffixes for autocommit period are hours, minutes, days, months and years. For example, a value of 2hours represents an autocommit period of 2 hours. The minimum allowed autocommit period is 5 minutes and the maximum allowed autocommit period is 10 years.

If this option value is *none*, then autocommit is disabled on the SnapLock volume.

[`-is-volume-append-mode-enabled {true|false}`] - Is Volume Append Mode Enabled

Specifies if the volume append mode is enabled or disabled.

It can be modified only when the volume is not mounted and does not have any data or Snapshot copies.

The volume append mode is not supported on SnapLock audit log volumes.

When it is enabled, all the files created with write permissions on the volume are WORM appendable files by default. All the WORM appendable files that are not modified for a period greater than the autocommit period of the volume are also committed to the WORM read-only state.

If it is set to `true`, then the volume append mode is enabled.

If it is set to `false`, then the volume append mode is disabled.

Volume append mode is disabled by default when the volume is created.

[`-privileged-delete {disabled|enabled|permanently-disabled}`] - Privileged Delete

Specifies the privileged-delete attribute of a SnapLock volume. This parameter must be specified alone.

If it is set to `enabled` then the privileged-delete operation can be performed using the [volume file privileged-delete](#) command.

If it is set to `disabled`, then the privileged-delete operation is not supported.

Once it is set to `permanently-disabled`, then neither the privileged-delete operation nor any change in the volume privileged-delete attribute is permitted.

Examples

The following command sets `-default-retention-period` of a given SnapLock volume:

```
cluster1::> volume snaplock modify -volume vol_slc -default-retention
-period 2years

cluster1::>
```

The following command sets `-maximum-retention-period` of a given SnapLock volume to `infinite`:

```
cluster1::> volume snaplock modify -volume vol_slc -maximum-retention
-period infinite

cluster1::>
```

The following command enables the privileged-delete operation on a SnapLock volume.


```

cluster1::> volume snaplock modify -vserver vs1 -volume vol_sle
-privileged-delete enabled
[Job 38] Job succeeded: Privileged-delete Attribute Change for Volume
"vs1:vol_sle" Completed.
cluster1::>
cluster1::>volume snaplock show -vserver vs1 -volume vol_sle -fields
privileged-delete
vserver volume    privileged-delete
-----
vs1      vol_sle  enabled

cluster1::>

```

Related Links

- [volume file privileged-delete](#)

volume snaplock prepare-to-downgrade

Prepares the system for downgrade

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `volume snaplock prepare-to-downgrade` command prepares nodes to downgrade to a release without SnapLock volume append mode feature. Prior to disabling the feature, the command disables volume append mode on all SnapLock volumes in the cluster.

Examples

The following example disables the SnapLock volume append mode feature in the local cluster:

```
cluster1::> volume snaplock prepare-to-downgrade
```

volume snaplock show

Display SnapLock attributes of a SnapLock volume

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

```
The ` volume snaplock show ` command displays following information :
```

- Vserver name

- Volume name
- SnapLock Type of the volume
- Minimum retention period applicable of the volume
- Default retention period applicable of the volume
- Maximum retention period applicable of the volume
- Autocommit period of the volume
- Volume Append Mode attribute of the volume
- Privileged Delete attribute of the volume
- Litigation count on the volume
- Volume expiry time of the volume
- Volume ComplianceClock
- SnapLock audit log volume
- Unspecified retention file count on the volume

This command is applicable only for SnapLock volumes.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver

If this parameter is specified, the command displays information for all the SnapLock volumes that match the specified `-vserver` value.

[-volume <volume name>] - Volume

If this parameter is specified, the command displays information for the specified `-volume` value.

[-type {non-snaplock|compliance|enterprise}] - SnapLock Type

If this parameter is specified, the command displays all the volumes that match the specified `-type` value.

[-minimum-retention-period {{<integer> seconds|minutes|hours|days|months|years} | infinite}] - Minimum Retention Period

If this parameter is specified, the command displays all the volumes that match the specified `-minimum-retention-period` value.

[-default-retention-period {{<integer> seconds|minutes|hours|days|months|years} | min | max | infinite | unspecified}] - Default Retention Period

If this parameter is specified, the command displays all the volumes that match the specified `-default-retention-period` value.

[`-maximum-retention-period` {<integer> seconds|minutes|hours|days|months|years} | infinite}] - Maximum Retention Period

If this parameter is specified, the command displays all the volumes that match the specified `-maximum-retention-period` value.

[`-autocommit-period` {<integer> minutes|hours|days|months|years} | none}] - Autocommit Period

If this parameter is specified, the command displays all the volumes that match the specified `-autocommit-period` value.

[`-is-volume-append-mode-enabled` {true|false}] - Is Volume Append Mode Enabled

If this parameter is specified, the command displays all the volumes that match the specified `-is-volume-append-mode-enabled` value.

[`-privileged-delete` {disabled|enabled|permanently-disabled}] - Privileged Delete

If this parameter is specified, the command displays all the volumes that match the specified `-privileged-delete` value.

[`-expiry-time` <text>] - Expiry Time

If this parameter is specified, the command displays all the volumes that match the specified `-expiry-time` value.

[`-compliance-clock-time` <text>] - ComplianceClock Time

If this parameter is specified, the command displays all the volumes that match the specified `-compliance-clock-time` value.

[`-litigation-count` <integer>] - Litigation Count

If this parameter is specified, the command displays all the volumes that match the specified `-litigation-count` value.

[`-is-audit-log-volume` {true|false}] - Is SnapLock Audit Log Volume

If this parameter is specified, the command displays all the volumes that match the specified `-is-audit-log-volume` value.

[`-unspecified-retention-file-count` <integer>] - Unspecified Retention File Count

If this parameter is specified, the command displays all the volumes that match the specified `-unspecified-retention-file-count` value.

Examples

The following command shows summary of SnapLock volumes on a vservers:

```
cluster1::> volume snaplock show
Vserver          Volume          SnapLock Type ComplianceClock Time
-----
-----
vs1              vol_slc         compliance      Mon Jan 19 14:12:34 IST 2015
+05:30
vs1              vol_sle         enterprise      Mon Jan 19 14:12:34 IST 2015
+05:30
2 entries were displayed.

cluster1::>
```

The following commands lists the complete SnapLock attributes of two given SnapLock volumes:

```

cluster1::> volume snaplock show -vserver vs1 -volume vol_slc
Vserver Name: vs1
                Volume Name: vol_slc
                SnapLock Type: compliance
                Minimum Retention Period: 1 years
                Default Retention Period: max
                Maximum Retention Period: 30 years
                Autocommit Period: 12 hours
Is Volume Append Mode Enabled: false
                Privileged Delete: permanently-disabled
                Expiry Time: Thu May 11 14:37:21 GMT 2017
                ComplianceClock Time: Wed May 11 20:08:41 IST 2016 +05:30
                Litigation Count: 0
Is SnapLock Audit Log Volume: false
Unspecified Retention File Count: 0

```

```
cluster1::>
```

```

cluster1::> volume snaplock show -vserver vs1 -volume vol_sle
Vserver Name: vs1
                Volume Name: vol_sle
                SnapLock Type: enterprise
                Minimum Retention Period: 6 months
                Default Retention Period: min
                Maximum Retention Period: infinite
                Autocommit Period: none
Is Volume Append Mode Enabled: false
                Privileged Delete: enabled
                Expiry Time: infinite
                ComplianceClock Time: Wed May 11 20:08:44 IST 2016 +05:30
                Litigation Count: 0
Is SnapLock Audit Log Volume: false
Unspecified Retention File Count: 0

```

volume snapshot commands

volume snapshot compute-reclaimable

Calculate the reclaimable space if specified snapshots are deleted

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume snapshot compute-reclaimable` command calculates the volume space that can be reclaimed if one or more specified Snapshot copies are deleted.

The command heavily uses system's computational resources so it can reduce the performance for client requests and other system processes. Therefore, the queries that use query operators (* , | , etc.), are disabled for this command. You should not specify more than three Snapshot copies per query. Snapshot copies must be specified as a comma-separated list with no spaces after the commas.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This specifies the volume for which reclaimable space is to be calculated.

-snapshots <snapshot name>,... - List of Snapshots

This specifies one or more than one Snapshot copies that are to be considered for deletion. If you list more than one Snapshot copy, specify a comma-separated list with no spaces after the commas.

Examples

The following example calculates the space that can be reclaimed if the Snapshot copy named hourly.2008-01-10_1505 is deleted on a volume named vol3, which is a part of the Vserver named vs0:

```
cluster1::> volume snapshot compute-reclaimable -vserver vs0
-volume vol3 -snapshots hourly.2008-01-10_1505
```

volume snapshot create

Create a snapshot

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume snapshot create` command creates a Snapshot copy of a specified volume.

Parameters

-vserver <vserver name> - Vserver

This specifies the Vserver that contains the volume on which the snapshot is to be created.

-volume <volume name> - Volume

This specifies the volume where a Snapshot copy is to be created.

-snapshot <snapshot name> - Snapshot

This specifies the name of the Snapshot copy that is to be created.

[-comment <text>] - Comment

This optionally specifies a comment for the Snapshot copy.

[`-foreground {true|false}`] - Foreground Process

If you use this option and select `false`, the Snapshot copy creation process runs in the background. If you use this option and select `true`, the Snapshot copy creation process runs in the foreground. The default is `true`.

[`-snapmirror-label <text>`] - Label for SnapMirror Operations

If you specify this option, the Snapshot copy is created with the SnapMirror Label that you specify. If this option is not specified, the Snapshot copy is created with no SnapMirror Label. The SnapMirror Label is used by the Vaulting subsystem when you back up Snapshot copies to the Vault Destination.

[`-expiry-time <MM/DD/YYYY HH:MM:SS>`] - Expiry Time

If you specify this option, the Snapshot copy is created with the expiry time that you specify. The expiry time indicates the time at which the Snapshot copy becomes eligible for deletion.

[`-snaplock-expiry-time {MM/DD/YYYY HH:MM:SS [(+|-)hh:mm] | infinite}`] - SnapLock Expiry Time

This optionally specifies a snaplock-expiry-time for the Snapshot copy. A Snapshot copy with an expiry time cannot be deleted until the expiry time has elapsed. This option takes effect only on volumes with snapshot-locking-enabled set to true. If this option is not specified, the Snapshot copy will not be under retention.

Examples

The following example creates a Snapshot copy named `vol3_snap` on a volume named `vol3` on a Vserver named `vs0`. The Snapshot copy is given the comment "Single snapshot" and the operation runs in the background.

```
cluster1::> volume snapshot create -vserver vs0 -volume vol3 -snapshot
vol3_snapshot -comment "Single snapshot" -foreground false
```

volume snapshot delete

Delete a snapshot

Availability: This command is available to `cluster` and `Vserver` administrators at the `admin` privilege level.

Description

The `volume snapshot delete` command deletes a Snapshot copy from a specified volume.

Parameters

`-vserver <vserver name>` - Vserver

This specifies the Vserver that contains the volume on which the specified Snapshot copy is saved.

`-volume <volume name>` - Volume

This specifies the volume from which a Snapshot copy is to be deleted.

`-snapshot <snapshot name>` - Snapshot

This specifies the Snapshot copy that is to be deleted.

[`-foreground {true|false}`] - Foreground Process

If you use this option and set it to `false`, the delete operation runs as a background process. If you specify this option and set it to `true`, the operation runs as a foreground process. The default is `true`.

[`-force <true>`] - Force Delete (privilege: advanced)

If you use this switch, the Snapshot copy is immediately deleted without generating any confirmation messages. If you do not use this option the operation generates confirmation messages and the operation is disallowed on application tagged volumes. Passing in a value of `true` is supported, but not required. The `force` switch is typically used for scripting applications where users cannot directly confirm the delete operation.

[`-ignore-owners <true>`] - Ignore Snapshot Owners (privilege: advanced)

If you use this switch, the command ignores other processes that might be accessing the Snapshot copy. If you do not use this option the operation exhibits default behavior and checks the owners tags before allowing the deletion to occur. Passing in a value of `true` is supported, but not required.

Examples

The following example deletes a Snapshot copy named `vol3_daily` from a volume named `vol3` on a Vserver named `vs0`:

```
cluster1::> volume snapshot delete -vserver vs0 -volume vol3 -snapshot
vol3_daily
```

volume snapshot modify-snaplock-expiry-time

Modify expiry time of a SnapLock Snapshot copy

Availability: This command is available to `cluster` and `Vserver` administrators at the `admin` privilege level.

Description

The `volume snapshot modify-snapshot-expiry-time` extends snaplock expiry time of an existing Snapshot copy.

Parameters

`-vserver <vserver name>` - Vserver

This specifies the Vserver that contains the volume on which the Snapshot copy is located.

`-volume <volume name>` - Volume

This specifies the volume where a Snapshot copy is to be located.

`-snapshot <text>` - Snapshot

This specifies the name of the Snapshot copy locked by SnapLock whose snaplock expiry time needs to be modified.

[*-expiry-time* {MM/DD/YYYY HH:MM:SS [{+|-}hh:mm] | *infinite*}] - SnapLock Expiry Time

Specifies the new snaplock expiry that is applied to Snapshot copy locked by SnapLock.

If this option value is *infinite*, then a retention period that never expires is applied to the Snapshot copy.

Examples

The following example extends the retention period of a Snapshot copy *snap1* to *"03/03/2020 00:00:00"*:

```
cluster1::> volume snapshot modify-snaplock-expiry-time -vserver vs1
-volume voll1 -snapshot snap1 -expiry-time "03/03/2020 00:00:00"
cluster1::>
```

The following example extends the retention period of a Snapshot copy *snap2* to *infinite*:

```
cluster1::> volume snapshot modify-snaplock-expiry-time -vserver vs1
-volume voll1 -snapshot snap2 -expiry-time infinite
cluster1::>
```

```
cluster1::> volume snapshot show -vserver vs1 -fields snaplock-expiry-time
vserver volume snapshot snaplock-expiry-time
-----
vs1      voll1  snap1  3/3/2020 00:00:00 +05:30
vs1      voll1  snap2  infinite
```

volume snapshot modify

Modify snapshot attributes

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume snapshot modify` command enables you to change the text comment associated with a Snapshot copy.

Parameters

`-vserver <vserver name>` - Vserver

This specifies the Vserver that contains the volume on which the specified Snapshot copy is saved.

`-volume <volume name>` - Volume

This specifies the volume whose Snapshot copy is to be modified.

-snapshot <snapshot name> - Snapshot

This specifies the Snapshot copy whose text comment is to be modified.

[-comment <text>] - Comment

This specifies the new comment for the Snapshot copy.

[-snapmirror-label <text>] - Label for SnapMirror Operations

This specifies the SnapMirror Label for the Snapshot copy. The SnapMirror Label is used by the Vaulting subsystem when you back up Snapshot copies to the Vault Destination. If an empty label ("") is specified, the existing label will be deleted.

[-expiry-time <MM/DD/YYYY HH:MM:SS>] - Expiry Time

This specifies the expiry time for the Snapshot copy. The expiry time indicates the time at which the Snapshot copy becomes eligible for deletion. If an expiry time of ("0") is specified, the existing expiry time will be deleted.

Examples

The following example modifies the comment of a Snapshot copy named vol3_snapshot of a volume named vol3 on a Vserver named vs0. The comment is changed to "Pre-upgrade snapshot".

```
cluster1::> volume snapshot modify -vserver vs0 -volume vol3
-snapshot vol3_snapshot -comment "Pre-upgrade snapshot"
```

volume snapshot partial-restore-file

Restore part of a file from a snapshot

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume snapshot partial-restore-file` command enables you to restore a range of bytes in a file from the version of the file saved in the Snapshot copy. This command is intended to be used to restore particular pieces of LUNs, NVMe namespaces, and NFS or CIFS container files that are used by a host to store multiple sources of data. For example, a host might be storing multiple user databases in the same LUN. A partial file restore can be used to restore one of those databases in the LUN without touching other databases stored in the LUN. This command is not intended for restoring parts of normal user-level files that are stored in the volume. You should use [volume snapshot restore-file](#) command to restore normal user-level files. The volume for the partial-restore should be online during this operation.

For LUNs and NVMe namespaces, this command is supported across all LUN and NVMe namespace source and destination objects with equal logical block sizes.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver which contains the volume.

[`-volume <volume name>`] - Volume Name

This specifies the volume in which the Snapshot copy is saved.

`-s, -snapshot <snapshot name>` - Snapshot Name

This specifies the Snapshot copy which contains the version of file from which a range of bytes is restored. The source file, LUN, or NVMe namespace must be present in the Snapshot copy.

`-path <text>` - Filepath

This specifies the relative path to the file, LUN, or NVMe namespace which is partially restored from the Snapshot copy. You should specify the `-volume` option so that the file, LUN, or NVMe namespace is searched and restored from the Snapshot copy of the specified volume. If you do not specify the `-volume` then the file, LUN, or NVMe namespace is searched and restored from the Snapshot copy of the root volume. The destination file, LUN, or NVMe namespace must be present in the active file system.

`-start-byte <integer>` - Starting Byte Offset (Multiple of 4096)

This specifies the starting byte offset in the file to partially restore. The first byte of the file is byte zero. The start byte must be a multiple of 4096. In addition, the start byte must not exceed the size of the source or destination file.

`-byte-count <integer>` - Number of Bytes to Restore (Multiple of 4096)

This specifies the total number of bytes to restore, beginning at the `-start-byte` value. The `-byte-count` option must be a multiple of 4096. The maximum number of bytes that can be restored is 16 MB. The byte count must not exceed the range of the source or destination file.

Examples

The following example restores first 4096 bytes in the file `foo.txt` inside the volume `vol3` from the Snapshot copy `vol3_snap`:

```
cluster1::> volume snapshot partial-restore-file -vserver vs0 -volume vol3
-snapshot vol3_snap -path /foo.txt -start-byte 0 -byte-count 4096
```

Related Links

- [volume snapshot restore-file](#)

volume snapshot prepare-for-revert

Deletes multiple Snapshot copies of the current File System version.

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

This command will delete all Snapshot copies that have the format used by the current version of ONTAP. It will fail if Snapshot copy polices are enabled, or if Snapshot copies have an owner.



Snapshot policies must be disabled prior to running this command.

Parameters

-node <nodename> - Node (privilege: advanced)

The name of the node.

Examples

The following example prepares the Snapshot copies for revert.

```
cluster1::*> volume snapshot prepare-for-revert -node node1
```

volume snapshot rename

Rename a snapshot

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume snapshot rename` command renames a Snapshot copy.



You cannot rename a Snapshot copy that is created as a reference copy during the execution of the `volume move` command.

Parameters

-vserver <vserver name> - Vserver

This specifies the Vserver that contains the volume on which the specified Snapshot copy is to be renamed

-volume <volume name> - Volume

This specifies the volume that contains the Snapshot copy to be renamed.

-snapshot <snapshot name> - Snapshot

This specifies the Snapshot copy that is to be renamed.

-new-name <snapshot name> - Snapshot New Name

This specifies the new name for the Snapshot copy.

[-force <true>] - Force Rename (privilege: advanced)

If this parameter is specified, the Snapshot copy rename operation is allowed on application tagged volumes. Otherwise, the operation is disallowed on application tagged volumes.

Examples

The following example renames a Snapshot copy named `vol3_snap` on a volume named `vol3` and a Vserver named `vs0`. The Snapshot copy is renamed to `vol3_snap_archive`.

```
cluster1::> volume snapshot rename -vserver vs0 -volume vol3
-snapshot vol3_snap -new-name vol3_snap_archive
```

volume snapshot restore-file

Restore a file from a snapshot

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume snapshot restore-file` command enables you to restore a single file to a version saved in the Snapshot copy. You can restore a file over an existing copy of the file in the parent read-write volume or to a different location within the same parent read-write volume. If the destination file for the restore operation does not exist, a new file is created with the same version as the one saved in the Snapshot copy. If the destination file for the restore operation exists, then it is overwritten by the version from the Snapshot copy. This operation is used to restore normal user-level files, LUNs and NVMe namespaces. The command also supports restoring normal user-level files with streams. The command fails if you try to restore directories (and their contents). During the restore operation the parent read-write volume should remain online. The command fails if the destination path for the restore operation is in a different volume than the source volume.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver which contains the volume.

[-volume <volume name>] - Volume Name

This specifies the volume which contains the specified Snapshot copy.

-s, -snapshot <snapshot name> - Snapshot Name

This specifies the Snapshot copy from which the file is restored.

-path <text> - Filepath

This specifies the relative path to the file which is restored from the Snapshot copy. You should specify the `-volume` option so that the file is searched and restored from the Snapshot copy of the specified volume. If you do not specify the `-volume` then the file is searched and restored from the Snapshot copy of the root volume.

[-r, -restore-path <text>] - Restore Filepath

This option specifies the destination location inside the volume where the file is restored. If you do not specify this option, the file is restored at the same location referred by `-path` option. If you specify `-restore-path` option, then it should refer to a relative path location within the same volume which contains the source file. If you do not specify `-volume` along with the relative path, the file is restored in the root volume.

[-split-disabled <>true>] - Disable Space Efficient LUN Splitting

If you use this option and set it to `true`, space efficient LUN or NVMe namespace clone split is not allowed during the restore operation. If you use this option and set it to `false` or do not use this option, then space efficient LUN or NVMe namespace clone split is allowed during the restore operation.

[~~-ignore-streams~~ <true>] - Ignore Streams

If you use this parameter, the file is restored without its streams. By default, the streams are restored.

Examples

The following example restores a file `foo.txt` from the Snapshot copy `vol3_snap` inside the volume `vol3` contained in a Vserver `vs0`:

```
cluster1::> volume snapshot restore-file -vserver vs0 -volume vol3
-snapshot vol3_snap -path /foo.txt
```

volume snapshot restore

Restore the volume to a snapshot.

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume snapshot restore` command restores a Snapshot copy to be the read-write parent volume for the volume family. This replaces the current working copy of the volume with the Snapshot copy that results in a loss of all changes made since the Snapshot copy was created.



You should manually update all the SnapMirror relationships of a volume immediately after you restore its Snapshot copy. Not doing so can result in unusable SnapMirror relationships that must be deleted and re-created.

After the restore is complete, the size of the flexible volume will be set to either the current volume size or the snapshot size - whichever is greater.

Parameters

~~-vserver~~ <vserver name> - Vserver

This specifies the Vserver that contains the volume on which the specified Snapshot copy to be restored is saved.

~~-volume~~ <volume name> - Volume

This specifies the parent read-write volume whose Snapshot copy is to be restored to take its place.

~~-snapshot~~ <snapshot name> - Snapshot

This specifies the Snapshot copy that is to be restored to be the read-write parent volume.

[~~-force~~ <true>] - Force Restore

If you use this parameter, the Snapshot copy is restored even if the volume has one or more newer Snapshot copies which are currently used as reference Snapshot copy by SnapMirror. If a restore is done in this situation, this will cause future SnapMirror transfers to fail. The SnapMirror relationship may be repaired using [snapmirror resync](#) command if a common Snapshot copy is found between the source and destination volume. If there is no common Snapshot copy between the source and the destination volume, a baseline SnapMirror copy would be required. If you use this parameter, the operation is also allowed on

application tagged volumes.

[`-preserve-lun-ids {true|false}`] - Preserve LUN Identifiers

This option enables you to select whether the Snapshot copy restore needs to be non-disruptive to clients due to LUN or NVMe namespace identifiers changing. If you use this option and set it to `true`, or choose to not use this option at all, the `volume snapshot restore` command fails if the system determines that it cannot be non-disruptive with regards to LUN or NVMe namespace identifiers. If you use this option and set it to `false`, the restore operation proceeds even if this might cause client-visible effects. In this case, administrators should take the LUNs or NVMe namespaces offline before proceeding.

Examples

The following example restores a Snapshot copy named `vol3_snap_archive` to be the parent read-write volume for the volume family. The existing read-write volume is named `vol3` and is located on a Vserver named `vs0`:

```
cluster1::> volume snapshot restore -vserver vs0 -volume vol3
-snapshot vol3_snap_archive
```

Related Links

- [snapmirror resync](#)

volume snapshot show-delta

Computes delta between two Snapshot copies

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume snapshot show-delta` command returns the number of bytes that changed between two Snapshot copies or a Snapshot copy and the active filesystem. This is calculated from the number of blocks that differ multiplied by the block size. The command also shows the time elapsed between the Snapshot copies in seconds.

Queries that use query operators (`*`, `|`, etc.) are disabled for this command to avoid performance degradation for client requests.

Parameters

`-vserver <vserver name>` - Vserver Name

This specifies the Vserver on which the volume is located.

`-volume <volume name>` - Volume Name

This specifies the volume for which the delta is to be calculated.

`-snapshot1 <snapshot name>` - First Snapshot Name

This specifies the first Snapshot copy for the comparison.

[`-snapshot2 <snapshot name>`] - Second Snapshot Name

This specifies the second Snapshot copy for the comparison. If the field is not specified, it is assumed to be the Active File System.

Examples

The following example shows the bytes changed and the time separating the two Snapshots copies:

```
cluster1::> volume snapshot show-delta -vserver vs0 -volume vol2
-snapshot1 one -snapshot2 two
A total of 139264 bytes (34 blocks) are different. Elapsed time between
the Snapshot copies: 1s.
```

volume snapshot show

Display a list of snapshots

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume snapshot show` command displays information about Snapshot copies. The command output depends on the parameters specified with the command. If no parameters are specified, the command displays a table with the following information about all the available Snapshot copies:

- Vserver name
- Volume name
- Snapshot copy name
- State
- Size
- Percentage of total blocks in the parent volume
- Percentage of used blocks in the parent volume

To display a detailed list view with additional information, run the command and select the `-instance` view. In addition to the above mentioned information about the Snapshot copies, the detailed list view provides the following additional information:

- Creation time
- Snapshot busy
- List of the Snapshot copy's owners
- Comment associated with the Snapshot copy
- SnapMirror Label associated with the Snapshot copy
- 7-Mode Snapshot
- Constituent Snapshot
- Expiry Time

- SnapLock Expiry Time
- SnapLock Snapshot Expired
- Seconds until SnapLock Snapshot Expiry

At the advanced or higher privilege level the detailed view provides the following additional information:

- Snapshot copy's Dataset ID
- Snapshot copy's master Dataset ID
- Number of consistency points in the Snapshot copy
- Internal status of the Snapshot copy
- File system version
- File system block format
- Physical Snap ID
- Logical Snap ID
- Database record owner
- Snapshot tags
- Instance UUID
- Version UUID
- Node
- AFS used size
- Compression savings size
- Deduplication savings size
- Vbn0 savings size
- Reserved size
- Logical used size
- Performance metadata size
- Status of FlexGroup Qtree support in the Snapshot copy



For Snapshot copies whose parent volume is a FlexGroup, some information is not available and empty values will be displayed. This information includes:

- State
- Size
- Percentage of total blocks in the parent volume
- Percentage of used blocks in the parent volume

All information is available for Snapshot copies whose parent volume is a FlexGroup Constituent.

At the admin and advanced privilege level, Snapshot copies whose parent volume is a FlexGroup Constituent

are not displayed by default. To display these, run the command and set the `is-constituent` to `true`. At the diagnostic or higher privilege level, all Snapshot copies are displayed by default.

The list view is automatically enabled if a single Snapshot copy is specified by using the `-vserver`, `-volume` and `-snapshot` options together.

A preformatted query for displaying the time-related information is available by specifying the `-time` format specifier. This displays a table that contains the following fields for all the available Snapshot copies:

- Vserver name
- Volume name
- Snapshot copy name
- Creation time

By using the `-fields` option you can choose to print only the certain fields in the output. This presents the selected fields in a table view. This is ideal when you want additional information to be different from the information that is provided by the default table view, but would like it in a format which is visually easy to compare.

You can specify additional parameters to display the information that matches only those parameters. For example, to display information only about Snapshot copies of the load-sharing volumes, run the command with the `-volume-type LS` parameter. If you specify multiple filtering parameters, only those Snapshot copies that match all the specified parameters are displayed.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-time]

If the `-time` format is specified, the command displays time related information about all entries.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver

If you use this parameter, the Snapshot copies located only on the specified Vserver will be displayed.

[-volume <volume name>] - Volume

If you use this parameter only Snapshot copies located on the specified volume will be displayed.

[-snapshot <snapshot name>] - Snapshot

If you use this parameter only Snapshot copies matching the specified name will be displayed.

[-dsid <integer>] - Snapshot Data Set ID (privilege: advanced)

If this parameter is specified, the command displays information only about the Snapshot copy that has the specified data set ID.

[`-msid <integer>`] - Snapshot Master Data Set ID (privilege: advanced)

If this parameter is specified, the command displays information only about the Snapshot copy that has the specified master data set ID.

[`-create-time <Date>`] - Creation Time

If this parameter is specified, the command displays information only about the Snapshot copies that match the specified creation time.

[`-busy {true|false}`] - Snapshot Busy

If this parameter is specified, the command displays information only about the Snapshot copies that have the specified busy status.

[`-owners <text>,...`] - List of Owners

If this parameter is specified, the command displays information only about the Snapshot copies that are owned by the specified list of owners.

[`-size {<integer>[KB|MB|GB|TB|PB]}`] - Snapshot Size

If this parameter is specified, the command displays information only about the Snapshot copies that have the specified size. The size is specified as a character specifying the unit of measurement after a number specifying the size in the mentioned unit of measurement: k (kilobytes), m (megabytes), g (gigabytes), or t (terabytes). If the unit of measurement is not specified, bytes are used as the unit, and the specified number is rounded up to the nearest 4 KB. You may also use an inequality such as `>10 MB` as input.

[`-blocks <percent>`] - Percentage of Total Blocks

If this parameter is specified, the command displays information only about the Snapshot copies that have the specified percentage of total blocks on their parent volumes. You may also use an inequality such as `>10` as input.

[`-usedblocks <percent>`] - Percentage of Used Blocks

If this parameter is specified, the command displays information only about the Snapshot copies that have the specified percentage of used blocks on their parent volumes. You may also use an inequality such as `>10` as input.

[`-cpcount <integer>`] - Consistency Point Count (privilege: advanced)

If this parameter is specified, the command displays information only about the Snapshot copies that have the specified number of consistency points. You may also use an inequality such as `<100` as input.

[`-comment <text>`] - Comment

If this parameter is specified, the command displays information only about the Snapshot copies that have the specified comment text. You may also specify an inequality such as `!"-"` as input.

[`-fs-version <text>`] - File System Version (privilege: advanced)

If you use this parameter the only Snapshot copies displayed are those that were created when the file system was of a specific release. This parameter is helpful especially when you need to upgrade to newer software release and want to know the Snapshot copies that will be impacted by the upgrade process.

[`-logical-snap-id <integer>`] - Logical Snap ID (privilege: advanced)

If you use this parameter only those Snapshot copies which have the specified logical snapshot ID will be shown. Logical snapshot ids are internal implementation details of volume Snapshot copies.

[-is-7-mode {true|false}] - 7-Mode Snapshot

If you use this parameter only those Snapshot copies which have the specified value are shown. This value is *true* for the Snapshot copies that exist on the volume that was in 7-mode configuration and then transitioned to a clustered configuration. In such a scenario, the volume is in a clustered configuration and the existing Snapshot copies are still in the 7-mode configuration.

[-snapmirror-label <text>] - Label for SnapMirror Operations

If you use this parameter, only those Snapshot copies that have the specified SnapMirror Label value are shown.

[-state {valid|invalid|partial|pre-conversion}] - Snapshot State

If you use this parameter only those Snapshot copies which have the specified state will be shown.

[-is-constituent {true|false}] - Constituent Snapshot

If you use this parameter, only those Snapshot copies whose parent volume is a constituent volume of a FlexGroup will be shown.

[-node <nodename>] - Node (privilege: advanced)

If you use this parameter only those Snapshot copies that are located on the specified storage system are shown.

[-inofile-version <integer>] - Snapshot Inofile Version (privilege: advanced)

If this parameter is specified, the command displays information only about the Snapshot copies whose inode files are at the specified version.

[-expiry-time <MM/DD/YYYY HH:MM:SS>] - Expiry Time

If you use this parameter only those Snapshot copies that have the specified expiry time are shown.

[-compression-type {none|secondary|adaptive}] - Compression Type (privilege: advanced)

If you use this parameter only those Snapshot copies that have the specified compression type are shown.

[-snaplock-expiry-time {MM/DD/YYYY HH:MM:SS [{+|-}hh:mm] | infinite}] - SnapLock Expiry Time

If you use this parameter only those Snapshot copies that have the specified snaplock expiry time are shown.

[-application-io-size <Application I/O Size>] - Application IO Size (privilege: advanced)

If you use this parameter only those Snapshot copies that have the specified application IO size are shown.

[-is-qtree-caching-enabled {true|false}] - Is Qtree Caching Support Enabled (privilege: advanced)

If you use this parameter, only those Snapshot copies of FlexGroups or origin of a FlexCache volumes that have the specified Qtree caching status are shown.

[-snaplock-snapshot-expired {true|false}] - Has the SnapLock Snapshot Expired

If this parameter is specified, the command displays information as to whether or not a locked Snapshot copy has expired.

[-seconds-until-snaplock-snapshot-expiry <integer>] - Seconds until the SnapLock Snapshot Expiry

If this parameter is specified, the command displays remaining SnapLock expiry time (in seconds) of a locked Snapshot copy.

Examples

The following example displays default information about all Snapshot copies of a volume named vol1:

```
cluster1::> volume snapshot show -volume vol1
```

				---Blocks	
Vserver	Volume	Snapshot	Size	Total%	Used%
cluster1	vol1	one	68KB	0%	33%
		two	72KB	0%	34%

2 entries were displayed.

The following example displays Snapshot copies which are older than 1 hour, limiting the output to wanted fields:

```
cluster1::> volume snapshot show -create-time <1h -fields create-time, size
```

vserver	volume	snapshot	create-time	size
cluster1	vol1	one	Mon Nov 17 10:23:42 2014	68KB
cluster1	vol1	two	Mon Nov 17 10:23:44 2014	72KB

2 entries were displayed.

The following example displays detailed information about a specific Snapshot copy, using the 'snap' alias:

```

cluster1::> snap show -volume voll1 -snapshot one -instance
Vserver: cluster1

                Volume: voll1
                Snapshot: one
                Snapshot Data Set ID: 4294968322
                Snapshot Master Data Set ID: 6442451970
                Creation Time: Mon Nov 17 10:23:42 2014
                Snapshot Busy: false
                List of Owners: -
                Snapshot Size: 68KB
                Percentage of Total Blocks: 0%
                Percentage of Used Blocks: 33%
                Consistency Point Count: 4
                Comment: -
                File System Version: 9.0
                7-Mode Snapshot: false
                Label for SnapMirror Operations: -
                Constituent Snapshot: false
                Node: node1
                Snapshot Inofile Version: 3
                Expiry Time: -
                SnapLock Expiry Time: -
                SnapLock Snapshot Expired: -
                Seconds until SnapLock Snapshot Expiry: -

```

volume snapshot autodelete modify

Modify autodelete settings

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume snapshot autodelete modify` command enables you to modify Snapshot autodelete and LUN, NVMe namespace or file clone autodelete policy settings. Based on the defined policy, automatic deletion of Snapshot copies and LUN, NVMe namespace or file clones is triggered. Automatic deletion of Snapshot copies and LUN, NVMe namespace or file clones is useful when you want to automatically reclaim space consumed by the Snapshot copies and LUN, NVMe namespace or file clones from the volume when it is low in available space. LUN, NVMe namespace or file clone autodelete follows Snapshot copy autodelete. This command works only on a read-write parent volume. You cannot setup automatic Snapshot copy deletion and automatic LUN, NVMe namespace or file clone deletion for read-only volumes.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver on which the volume is located.

-volume <volume name> - Volume Name

This specifies the volume whose autodelete policy has to be modified.

[-enabled {true|false}] - Enabled

This option specifies whether automatic deletion of Snapshot copies and LUN, NVMe namespace or file clones is enabled or disabled. If set to *true*, automatic deletion of Snapshot copies and LUN, NVMe namespace or file clones is enabled. If set to *false*, automatic deletion of Snapshot copies and LUN, NVMe namespace or file clones is disabled.

[-commitment {try|disrupt|destroy}] - Commitment

This option specifies which Snapshot copies and LUN, NVMe namespace or file clones can be automatically deleted to reclaim back space. + When set to *try*, the Snapshot copies which are not locked by any application and the LUN, NVMe namespace or file clones which are not configured as preserved are deleted. + When set to *disrupt*, the Snapshot copies which are not locked by data backing functionalities (such as volume clones, LUN clones, NVMe namespace clones and file clones) and LUN, NVMe namespace or file clones which are not configured as preserved are deleted. In the *disrupt* mode, the Snapshot copies locked by data protection utilities such as Snapmirror and Volume Move can be deleted. If such a locked Snapshot copy is deleted during the data transfer, the transfer is aborted. + When set to *destroy*, the Snapshot copies locked by the data backing functionalities are deleted. In addition, all the LUN, NVMe namespace or file clones in the volume are deleted.

[-defer-delete {scheduled|user_created|prefix|none}] - Defer Delete

This option determines the order in which Snapshot copies can be deleted. + Possible values are as follows:

- When set to *scheduled*, scheduled Snapshot copies are the last to be deleted.
- When set to *user_created*, user Snapshot copies are the last to be deleted.
- When set to *prefix*, Snapshot copies matching a certain prefix are the last to be deleted.
- When set to *none*, no defer deletion order is honored.

This option is not applicable for LUN, NVMe namespace or file clones.

[-delete-order {newest_first|oldest_first}] - Delete Order

This option specifies if the oldest Snapshot copy and the oldest LUN, NVMe namespace or file clone or the newest Snapshot copy and the newest LUN, NVMe namespace or file clone are deleted first.

[-defer-delete-prefix <text>] - Defer Delete Prefix

This option specifies the prefix string for the `-defer-delete prefix` parameter. The option is not applicable for LUN, NVMe namespace or file clones.

[-target-free-space <percent>] - Target Free Space

This option specifies the free space percentage at which the automatic deletion of Snapshot copies and LUN, NVMe namespace or file clones must stop. Depending on the `-trigger` Snapshot copies and LUN, NVMe namespace or file clones are deleted until you reach the target free space percentage.

[-trigger {volume|snap_reserve|(DEPRECATED)-space_reserve}] - Trigger

This option specifies the condition which starts the automatic deletion of Snapshot copies and LUN, NVMe

namespace or file clones. + Setting this option to *volume* triggers automatic deletion of Snapshot copies and LUN, NVMe namespace or file clones when the volume reaches threshold capacity and the volume space reserved for Snapshot copies is exceeded. + Setting the option to *snap_reserve* triggers automatic deletion of Snapshot copies and LUN, NVMe namespace or file clones when the space reserved for Snapshot copies reaches threshold capacity. + Setting the option to (DEPRECATED)-*space_reserve* triggers automatic deletion of Snapshot copies when reserved space in the volume reaches threshold capacity and the volume space reserved for Snapshot copies is exceeded. +



The option *space_reserve* is deprecated.

The threshold capacity is determined by the size of the volume as follows:

- If the volume size is less than 20 GB, the autodelete threshold is 85%.
- If the volume size is equal to or greater than 20 GB and less than 100 GB, the autodelete threshold is 90%.
- If the volume size is equal to or greater than 100 GB and less than 500 GB, the autodelete threshold is 92%.
- If the volume size is equal to or greater than 500 GB and less than 1 TB, the autodelete threshold is 95%.
- If the volume size is equal to or greater than 1 TB, the autodelete threshold is 98%.

[*-destroy-list* <text>] - Destroy List

This option specifies a comma separated list of data backing functions which are affected if the automatic deletion of the Snapshot copy backing that service is triggered. The possible values for this option are *lun_clone*, *fileclone*, *lun_clone*, *sfsr*, *vol_clone*, *cifs_share*, or *none*. Except *none*, all the other options can be combined as a comma separated list. Note that "lun_clone", "file_clone" and "sfsr" individually are not valid values. Only pairs "lun_clone,file_clone" and "lun_clone,sfsr" are supported.



For the purposes of autodelete, *lun_clone* includes both LUNs and NVMe namespaces.

If you specify *vol_clone*, the cloned volume backed by the Snapshot copy is deleted. + If you specify *lun_clone*, and a LUN or NVMe namespace is in the process of being cloned when autodelete is triggered, the cloning operation is aborted. Any access to this LUN or NVMe namespace will result in an error being reported to the client. + If you specify *file_clone*, and the file cloning operation is in progress when autodelete is triggered, the cloning operation is aborted. Any access to this file will result in an error being reported to the client. + If you specify *sfsr*, and the file restore is in progress when autodelete is triggered, the restore operation is aborted. + If the Snapshot copy is locked either by a *lun_clone* or *file_clone* or both, the *-destroy-list* must be set to *lun_clone,file_clone*. + If the Snapshot copy is locked either by a *lun_clone* or *sfsr* operation or both, the *-destroy-list* must be set to *lun_clone,file_clone*. The options *file_clone* and *sfsr* are equivalent to each other. + If you set *-destroy-list* to *lun_clone,file_clone* and the Snapshot copy is backing a file clone or *sfsr* operation, both the operations are aborted. This is also the case when you set *-destroy-list* to *lun_clone,sfsr*. + LUN, NVMe namespace or file clone autodelete is applicable only if *-destroy-list* contains *lun_clone* and *file_clone*

Examples

The following example enables Snapshot autodelete and sets the trigger to *snap_reserve* for volume *vol13* which is part of the Vserver *vs0*:


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

The following example enables Snapshot autodelete and LUN, NVMe namespace or file clone autodelete for volume `vol3` which is part of the Vserve `vs0`:

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

volume snapshot autodelete show

Display autodelete settings

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume snapshot autodelete show` command displays information about Snapshot autodelete policies. The command output depends on the parameters specified with the command. If no parameters are specified, the command displays a table with the following information about all the available Snapshot autodelete policies:

- Vserver name
- Volume name
- Option name
- Option value

To display a detailed list view with additional information, run the command and select the `-instance` view. The detailed list view provides the following information:

- Vserver name
- Volume name
- Enabled
- Commitment
- Defer Delete
- Delete Order
- Defer Delete Prefix
- Target Free Space
- Trigger
- Destroy List
- Is Constituent Volume

By using the `-fields` option you can choose to print only the certain fields in the output. This presents the

selected fields in a table view. This is ideal when you want additional information to be different from the information that is provided by the default table view, but would like it in a format which is visually easy to compare.

You can specify additional parameters to display the information that matches only those parameters. For example, to display information only about Snapshot autodelete policies which are enabled, run the command with `-enabled true` parameter. If you specify multiple filtering parameters, only those policies that match all the specified parameters are displayed.

Parameters

{ [-fields <fieldname>,...]

This option allows you to print only certain fields in the output.

| [-instance] }

This option allows you to print a detailed list view.

[-vserver <vserver name>] - Vserver Name

If this parameter and the `-volume` parameter are specified, the command displays detailed autodelete policy information about the specified volume. If this parameter is specified by itself, the command displays autodelete policy information about volumes on the specified Vserver.

[-volume <volume name>] - Volume Name

If this parameter and the `-vserver` parameter are specified, the command displays detailed autodelete policy information about the specified volume. If this parameter is specified by itself, the command displays autodelete policy information about all volumes matching the specified name.

[-enabled {true|false}] - Enabled

If this parameter is specified, the command displays information about autodelete policies that match the specified parameter value.

[-commitment {try|disrupt|destroy}] - Commitment

If this parameter is specified, the command displays information about autodelete policies that match the specified commitment value.

[-defer-delete {scheduled|user_created|prefix|none}] - Defer Delete

If this parameter is specified, the command displays information about autodelete policies that match the specified defer deletion criterion.

[-delete-order {newest_first|oldest_first}] - Delete Order

If this parameter is specified, the command displays information about autodelete policies that match the specified deletion order.

[-defer-delete-prefix <text>] - Defer Delete Prefix

If this parameter is specified, the command displays information about autodelete policies that match the prefix used for deferring deletion.

[-target-free-space <percent>] - Target Free Space

If this parameter is specified, the command displays information about autodelete policies that match the specified target free space.

[`-trigger {volume|snap_reserve| (DEPRECATED) -space_reserve}`] - Trigger

If this parameter is specified, the command displays information about autodelete policies that match the specified trigger condition.

[`-destroy-list <text>`] - Destroy List

If this parameter is specified, the command displays information about autodelete policies that match the specified value.

[`-is-constituent {true|false}`] - Is Constituent Volume

If this parameter is specified, the command displays information about autodelete policies for the constituent volumes of FlexGroups.

Examples

The following example displays Snapshot autodelete policy settings for volume vol3 which is inside the Vserver vs0:

```
cluster1::> volume snapshot autodelete show -vserver vs0 -volume vol3
```

Vserver	Volume	Option Name	Option Value
vs0	vol3	Enabled	false
		Commitment	try
		Trigger	volume
		Target Free Space	20%
		Delete Order	oldest_first
		Defer Delete	user_created
		Defer Delete Prefix	(not specified)
		Destroy List	none

volume snapshot policy add-schedule

Add a schedule to snapshot policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume snapshot policy add-schedule` command adds a schedule to a Snapshot policy. You can create a schedule by using the [job schedule cron create](#) or [job schedule interval create](#) commands.

Parameters

`-vserver <vserver name>` - Vserver Name

This specifies the Vserver on which a Snapshot policy schedule is to be added.

-policy <snapshot policy> - Snapshot Policy Name

This specifies the Snapshot policy to which a schedule is to be added.

-schedule <text> - Schedule Name

This specifies the schedule that is to be added to the Snapshot policy.

-count <integer> - Maximum Snapshot Copies for Schedule

This specifies the maximum number of Snapshot copies that can be taken by the specified schedule. The total count of all the Snapshot copies to be retained for the policy cannot be more than 1023.

[-prefix <text>] - Snapshot Copy Name Prefix for Schedule

This option specifies the prefix with which Snapshot copies will be created for the added schedule. Every schedule has only one prefix. Once a prefix gets associated with a schedule, you cannot update the prefix. If some prefix is already associated with the schedule and you do not specify this parameter, then the previously defined prefix is used. The command fails if you try to update an existing prefix for a schedule. If no prefix is associated with the schedule and you do not specify this parameter, then schedule name is be used as the prefix.

[-snapmirror-label <text>] - Label for SnapMirror Operations

This specifies the SnapMirror Label identified with a Snapshot copy when it is created for the added schedule. The SnapMirror Label is used by the Vaulting subsystem when you back up Snapshot copies to the Vault Destination.

[-retention-period <snaplock minmax period>] - SnapLock Retention Period

This specifies the retention period for Snapshot copies created by the schedule.

Examples

The following example adds a schedule named midnight to the Snapshot policy named snappolicy_nightly on Vserver vs0. The schedule can take a maximum of five Snapshot copies.

```
cluster1::> volume snapshot policy add-schedule -vserver vs0 -policy
snappolicy_nightly -schedule midnight -count 5 -retention-period "7 days"
```

Related Links

- [job schedule cron create](#)
- [job schedule interval create](#)

volume snapshot policy create

Create a new snapshot policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume snapshot policy create` command creates a Snapshot policy. A Snapshot policy includes at least one schedule, up to a maximum of five schedules, and a maximum number of Snapshot copies per

schedule. You can create a schedule by using the [job schedule cron create](#) or [job schedule interval create](#) commands. When applied to a volume, the Snapshot policy specifies the schedule on which Snapshot copies are taken and the maximum number of Snapshot copies that each schedule can take. The total count of all the Snapshot copies to be retained for the policy cannot be more than 1023.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver on which the Snapshot policy is to be created.

-policy <snapshot policy> - Snapshot Policy Name

This specifies the Snapshot policy that is to be created.

-enabled {true|false} - Snapshot Policy Enabled

This specifies whether the Snapshot policy is enabled.

[-comment <text>] - Comment

This option specifies a text comment for the Snapshot policy.

-schedule1 <text> - Schedule1 Name

This specifies the name of the first schedule associated with the Snapshot policy.

-count1 <integer> - Maximum Snapshot Copies for Schedule1

This specifies the maximum number of Snapshot copies that can be taken by the first schedule.

[-prefix1 <text>] - Snapshot Copy Name Prefix for Schedule1

This option specifies the prefix associated with the first schedule. Every schedule has only one prefix. The command fails if you try to update an existing prefix. If you do not specify this parameter and there is no prefix associated with the schedule, the schedule name is used as the prefix. If you do not specify this parameter and there is already a prefix associated with the schedule from a previous invocation of the command, then that prefix is used.

[-snapmirror-label1 <text>] - Label for SnapMirror Operations for Schedule1

This specifies the SnapMirror Label of the first schedule associated with the Snapshot policy. Once specified, all Snapshot copies created for that schedule have the SnapMirror Label assigned to them. The SnapMirror Label is used by the Vaulting subsystem when you back up Snapshot copies to the Vault Destination.

[-retention-period1 <snaplock minmax period>] - SnapLock Retention Period

This specifies the retention period for Snapshot copies that can be taken by the first schedule. If this parameter is specified, the `-schedule1` parameter must also be specified. The default is no retention period.

[-schedule2 <text>] - Schedule2 Name

This option specifies the name of the second schedule associated with the Snapshot policy. If this parameter is specified, the `-count2` parameter must also be specified.

[-count2 <integer>] - Maximum Snapshot Copies for Schedule2

This option specifies the maximum number of Snapshot copies that can be taken by the second schedule. If this parameter is specified, the `-schedule2` parameter must also be specified.

[-prefix2 <text>] - Snapshot Copy Name Prefix for Schedule2

This option specifies the prefix associated with the second schedule. If this parameter is specified, `-schedule2` and `-count2` parameters must also be specified. Every schedule has only one prefix. The command fails if you try to update an existing prefix. If you do not specify this parameter and there is no prefix associated with the schedule, the schedule name is used as the prefix. If you do not specify this parameter and there is already a prefix associated with the schedule from a previous invocation of the command, then that prefix is used.

[-snapmirror-label2 <text>] - Label for SnapMirror Operations for Schedule2

This specifies the SnapMirror Label of the second schedule associated with the Snapshot policy. Once specified, all Snapshot copies created for that schedule have the SnapMirror Label assigned to them. The SnapMirror Label is used by the Vaulting subsystem when you back up Snapshot copies to the Vault Destination.

[-retention-period2 <snaplock minmax period>] - SnapLock Retention Period

This specifies the retention period for Snapshot copies that can be taken by the second schedule. If this parameter is specified, the `-schedule2` parameter must also be specified. The default is no retention period.

[-schedule3 <text>] - Schedule3 Name

This option specifies the name of the third schedule associated with the Snapshot policy. If this parameter is specified, the `-count3` parameter must also be specified.

[-count3 <integer>] - Maximum Snapshot Copies for Schedule3

This option specifies the maximum number of Snapshot copies that can be taken by the third schedule. If this parameter is specified, the `-schedule3` parameter must also be specified.

[-prefix3 <text>] - Snapshot Copy Name Prefix for Schedule3

This option specifies the prefix associated with the third schedule. If this parameter is specified, `-schedule3` and `-count3` parameters must also be specified. Every schedule has only one prefix. The command fails if you try to update an existing prefix. If you do not specify this parameter and there is no prefix associated with the schedule, the schedule name is used as the prefix. If you do not specify this parameter and there is already a prefix associated with the schedule from a previous invocation of the command, then that prefix is used.

[-retention-period3 <snaplock minmax period>] - SnapLock Retention Period

This specifies the retention period for Snapshot copies that can be taken by the third schedule. If this parameter is specified, the `-schedule3` parameter must also be specified. The default is no retention period.

[-snapmirror-label3 <text>] - Label for SnapMirror Operations for Schedule3

This specifies the SnapMirror Label of the third schedule associated with the Snapshot policy. Once specified, all Snapshot copies created for that schedule have the SnapMirror Label assigned to them. The SnapMirror Label is used by the Vaulting subsystem when you back up Snapshot copies to the Vault Destination.

[-schedule4 <text>] - Schedule4 Name

This option specifies the name of the fourth schedule associated with the Snapshot policy. If this parameter is specified, the `-count4` parameter must also be specified.

[`-count4 <integer>`] - Maximum Snapshot Copies for Schedule4

This option specifies the maximum number of Snapshot copies that can be taken by the fourth schedule. If this parameter is specified, the `-schedule4` parameter must also be specified.

[`-prefix4 <text>`] - Snapshot Copy Name Prefix for Schedule4

This option specifies the prefix associated with the fourth schedule. If this parameter is specified, `-schedule4` and `-count4` parameters must also be specified. Every schedule has only one prefix. The command fails if you try to update an existing prefix. If you do not specify this parameter and there is no prefix associated with the schedule, the schedule name is used as the prefix. If you do not specify this parameter and there is already a prefix associated with the schedule from a previous invocation of the command, then that prefix is used.

[`-retention-period4 <snaplock minmax period>`] - SnapLock Retention Period

This specifies the retention period for Snapshot copies that can be taken by the fourth schedule. If this parameter is specified, the `-schedule4` parameter must also be specified. The default is no retention period.

[`-snapmirror-label4 <text>`] - Label for SnapMirror Operations for Schedule4

This specifies the SnapMirror Label of the fourth schedule associated with the Snapshot policy. Once specified, all Snapshot copies created for that schedule have the SnapMirror Label assigned to them. The SnapMirror Label is used by the Vaulting subsystem when you back up Snapshot copies to the Vault Destination.

[`-schedule5 <text>`] - Schedule5 Name

This option specifies the name of the fifth schedule associated with the Snapshot policy. If this parameter is specified, the `-count5` parameter must also be specified.

[`-count5 <integer>`] - Maximum Snapshot Copies for Schedule5

This option specifies the maximum number of Snapshot copies that can be taken by the fifth schedule. If this parameter is specified, the `-schedule5` parameter must also be specified.

[`-prefix5 <text>`] - Snapshot Copy Name Prefix for Schedule5

This option specifies the prefix associated with the fifth schedule. If this parameter is specified, `-schedule5` and `-count5` parameters must also be specified. Every schedule has only one prefix. The command fails if you try to update an existing prefix. If you do not specify this parameter and there is no prefix associated with the schedule, the schedule name is used as the prefix. If you do not specify this parameter and there is already a prefix associated with the schedule from a previous invocation of the command, then that prefix is used.

[`-retention-period5 <snaplock minmax period>`] - SnapLock Retention Period

This specifies the retention period for Snapshot copies that can be taken by the fifth schedule. If this parameter is specified, the `-schedule5` parameter must also be specified. The default is no retention period.

[`-snapmirror-label5 <text>`] - Label for SnapMirror Operations for Schedule5

This specifies the SnapMirror Label of the fifth schedule associated with the Snapshot policy. Once specified, all Snapshot copies created for that schedule have the SnapMirror Label assigned to them. The SnapMirror Label is used by the Vaulting subsystem when you back up Snapshot copies to the Vault Destination.

Examples

The following example creates a Snapshot policy named `snappolicy_4hrs` on a Vserver named `vs0`. The policy runs on a single schedule named `4hrs` with a prefix `every_4_hour` and has a maximum number of five Snapshot copies.

```
cluster1::> volume snapshot policy create -vserver vs0 -policy
snappolicy_4hrs
-schedule1 4hrs -count1 5 -prefix1 every_4_hour -retention-period1 "8
hours"
```

Related Links

- [job schedule cron create](#)
- [job schedule interval create](#)

volume snapshot policy delete

Delete a snapshot policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume snapshot policy delete` command deletes a Snapshot policy.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver on which the Snapshot policy is to be deleted.

-policy <snapshot policy> - Snapshot Policy Name

This specifies the Snapshot policy that is to be deleted.

Examples

The following example deletes a Snapshot policy named `snappolicy_hourly` on Vserver `vs0`:

```
cluster1::> volume snapshot policy delete -vserver vs0 -policy
snappolicy_hourly
```

volume snapshot policy modify-schedule

Modify a schedule within snapshot policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume snapshot policy modify-schedule` command modifies the maximum number of Snapshot copies that can be taken by a Snapshot policy's schedule.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver on which a Snapshot policy schedule is to be modified.

-policy <snapshot policy> - Snapshot Policy Name

This specifies the Snapshot policy whose schedule is to be modified.

-schedule <text> - Schedule Name

This specifies the schedule that is to be modified.

[-newcount <integer>] - Maximum Snapshot Copies for Schedule

This specifies the maximum number of Snapshot copies that can be taken by the specified schedule. The total count of all the Snapshot copies to be retained for the policy cannot be more than 1023.

[-newsnapmirror-label <text>] - Label for SnapMirror Operations

This specifies the SnapMirror Label identified with a Snapshot copy when it is created for the specified schedule. The SnapMirror Label is used by the Vaulting subsystem when you back up Snapshot copies to the Vault Destination. If an empty label ("") is specified, the existing label will be deleted.

[-newretention-period <snaplock minmax period>] - SnapLock Retention Period

This specifies the retention period for Snapshot copies created by the schedule.

Examples

The following example changes the maximum number of Snapshot copies from five to four for a schedule named `midnight` on a Snapshot policy named `snappolicy_nightly` on Vserver `vs0`:

```
cluster1::> volume snapshot policy modify-schedule -vserver vs0 -policy
snappolicy_nightly -schedule midnight -newcount 4 -newretention-period "7
days"
```

volume snapshot policy modify

Modify a snapshot policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume snapshot policy modify` command enables you to modify the description associated with a Snapshot policy and whether the policy is enabled or disabled.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver on which the Snapshot policy is to be modified.

-policy <snapshot policy> - Snapshot Policy Name

This specifies the Snapshot policy that is to be modified.

[-enabled {true|false}] - Snapshot Policy Enabled

This optionally specifies whether the Snapshot policy is enabled.

[-comment <text>] - Comment

This specifies the comment text for the Snapshot policy.

[-snapmirror-labels <text>,...] - Label for SnapMirror Operations

This optionally specifies a comma separated list of SnapMirror labels that are applied to the schedules in the Snapshot policy. Each label in the list applies to only one schedule in the Snapshot policy (maximum of 5 SnapMirror labels), the first label applying to the first schedule, the second label applying to the second schedule, and so on. You can have a maximum of five SnapMirror labels, which corresponds to the maximum number of schedules a Snapshot policy can have. If an empty string ("") is specified, the existing labels will be deleted from all the schedules.

Examples

The following example changes the description of a Snapshot policy named `snappolicy_wknd` on Vserver `vs0` to "Runs only on weekends":

```
cluster1::> volume snapshot policy modify -vserver vs0 -policy
snappolicy_wknd -comment "Runs only on weekends"
```

volume snapshot policy remove-schedule

Remove a schedule from snapshot policy

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume snapshot policy remove-schedule` command removes a schedule from a Snapshot policy.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver on which a Snapshot policy schedule is to be removed.

-policy <snapshot policy> - Snapshot Policy Name

This specifies the Snapshot policy from which a schedule is to be removed.

-schedule <text> - Schedule Name

This specifies the schedule that is to be removed from the Snapshot policy.

Examples

The following example removes a schedule named hourly from a Snapshot policy named snappolicy_daily on Vserver vs0:

```
cluster1::> volume snapshot policy remove-schedule -vserver vs0 -policy
snappolicy_daily -schedule hourly
```

volume snapshot policy show

Show snapshot policies

Availability: This command is available to *cluster* and *Vserver* administrators at the *admin* privilege level.

Description

The `volume snapshot policy show` command displays the following information about Snapshot policies:

- Vserver name
- Snapshot policy name
- Number of schedules in the policy
- Comment for the policy
- Individual schedule names
- Maximum number of Snapshot copies associated with each schedule
- Snapshot copy name prefixes for the schedules
- SnapMirror Labels associated with the schedules

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-revert-incompatible] (privilege: advanced)

If this parameter is specified, the command displays Snapshot policies that are not supported in Data ONTAP 8.2. The total Snapshot copy count in the policy needs to be reduced to be equal to or less than the supported count for the revert operation to succeed.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-vserver <vserver name>] - Vserver Name

If this parameter is specified, the command displays Snapshot policies on the specified Vserver.

[`-policy <snapshot policy>`] - Snapshot Policy Name

If this parameter is specified, the command displays detailed information about the specified Snapshot policy.

[`-enabled {true|false}`] - Snapshot Policy Enabled

If this parameter is specified, the command displays detailed information only about the Snapshot policy or policies that have the specified enabled value.

[`-comment <text>`] - Comment

If this parameter is specified, the command displays information only about the Snapshot policy or policies that have the specified comment.

[`-total-schedules <integer>`] - Total Number of Schedules in This Policy

If this parameter is specified, the command displays information only about the Snapshot policy or policies that have the specified total number of schedules.

[`-schedules <text>,...`] - Schedule Name

If this parameter is specified, the command displays information only about the Snapshot policy or policies that have the specified list of schedules.

[`-counts <integer>,...`] - Maximum Snapshots for the Schedule

If this parameter is specified, the command displays information only about the Snapshot policy or policies that have the specified list of maximum numbers of Snapshot copies per schedule.

[`-prefixes <text>,...`] - Prefix Name

If this parameter is specified, the command displays information only about the Snapshot policy or policies that have the specified list of prefixes.

[`-snapmirror-labels <text>,...`] - Label for SnapMirror Operations

If this parameter is specified, the command displays information only about the Snapshot policies that have the specified SnapMirror Label. When you specify a list of SnapMirror labels, the command displays all the Snapshot policies that contain any of the SnapMirror Labels specified in the list.

[`-policy-owner <text>`] - Owner of the policy

If this parameter is specified, the command displays information only about the Snapshot policies that have the specified policy owner.

[`-total-count <integer>`] - Total Number of Snapshots in This Policy

If this parameter is specified, the command displays information only about the Snapshot policies that have the specified total number of Snapshot copies.

[`-retention-periods <snaplock minmax period>,...`] - SnapLock Retention Period

If this parameter is specified, the command displays information only about the Snapshot policy or policies that have the specified retention-period set.

Examples

The following example displays information about all Snapshot policies:

```
cluster1::> volume snapshot policy show
```

Vserver: cm

Policy Name	Number of Schedules	Is Enabled	Comment
default	3	false	Default policy with hourly, daily weekly schedules.

Schedule	Count	Prefix	SnapMirror
----------	-------	--------	------------

hourly	6	hourly	-
daily	2	daily	-
weekly	2	weekly	-

default-1weekly	3	false	Default policy with 6 hourly, 2 daily 1 weekly schedule.
-----------------	---	-------	--

Schedule	Count	Prefix	SnapMirror
----------	-------	--------	------------

hourly	6	hourly	-
daily	2	daily	-
weekly	1	weekly	-

none	0	false	Policy for no automatic snapshots.
------	---	-------	------------------------------------

Schedule	Count	Prefix	SnapMirror
----------	-------	--------	------------

-	-	-	-
---	---	---	---

Vserver: vs0

Policy Name	Number of Schedules	Is Enabled	Comment
-------------	---------------------	------------	---------

p1	1	false	-
----	---	-------	---

Schedule	Count	Prefix	SnapMirror
----------	-------	--------	------------

weekly	2	weekly	-
--------	---	--------	---

p2	2	true	-
----	---	------	---

Schedule	Count	Prefix	SnapMirror
Label			
-----	-----	-----	

hourly	6	hourly	-
daily	2	daily	-

5 entries were displayed.

Copyright information

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

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

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

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

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

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

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

Trademark information

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