



volume efficiency commands

ONTAP 9.5 commands

NetApp
November 11, 2022

Table of Contents

- volume efficiency commands 1
 - volume efficiency check 1
 - volume efficiency modify 1
 - volume efficiency off 3
 - volume efficiency on 4
 - volume efficiency prepare-to-downgrade 5
 - volume efficiency promote 6
 - volume efficiency revert-to 7
 - volume efficiency show 8
 - volume efficiency start 13
 - volume efficiency stat 16
 - volume efficiency stop 17
 - volume efficiency undo 18
 - volume efficiency policy create 19
 - volume efficiency policy delete 21
 - volume efficiency policy modify 21
 - volume efficiency policy show 22

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

Specifies the Vserver on which the volume is located.

{ -volume <volume name> - Volume Name

Specifies the volume on which the verify operation needs to be started.

| -path </vol/volume> - Volume Path }

Specifies the volume path on which the verify operation needs to be started.

[-d, -delete-checkpoint {true|false}] - Delete Checkpoint

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.

Description

This command is used to set or modify the schedule, policy and various other efficiency configuration options on a volume.

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* policy when there is an active background operation on the volume.

Note that schedule and policy are mutually exclusive options.

[-compression-type {none|secondary|adaptive}] - Compression Type

This option is used to specify the size of compression group on the volume. The default value is determined based on the platform.

[-compression {true|false}] - Compression

This option is used to enable and disable compression. The default value is false.

[-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 false.

You can use the *inline-only* predefined efficiency policy to run inline compression without the need of any background efficiency operations.

[`-inline-dedupe {true|false}`] - Inline Dedupe

This option is used to enable and disable inline deduplication. The default value is false.

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 false.

[`-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 false.

[`-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 false.

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.

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 vol1
```

```
cluster1::> volume efficiency off -vserver vs1 -path /vol/vol1
```

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
```

```
cluster1::> volume efficiency on -vserver vs1 -volume vol1 -needs-upgrade  
true
```

```
cluster1::> volume efficiency on -vserver vs1 -path /vol/vol1 -needs  
-upgrade true
```

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 releases prior to ONTAP 9. 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

This parameter specifies the Data ONTAP version that introduced new volume efficiency feature set.

Examples

The following example disables the 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 the features introduced in Data ONTAP 8.3.2.

```
cluster1::*> volume efficiency prepare-to-downgrade -disable-feature-set
8.3.2
```

The following example ignores offline volumes while disabling the the features introduced in Data ONTAP 8.3.2

```
cluster1::*> volume efficiency prepare-to-downgrade -disable-feature-set
8.3.2 -skip-offline-volumes true
```

The following example ignores offline volumes while disabling the 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

This specifies the Vserver on which the volume is located.

{ -volume <volume name> - Volume Name

This specifies the name of the volume on which auto scheduling needs to be restarted.

| -path </vol/volume> - Volume Path }

This specifies the volume path on which auto scheduling needs to be restarted.

Examples

The following examples promote 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 Data ONTAP. This command is not supported on FlexGroups.

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 for which volume efficiency metadata needs to be reverted.

| -path </vol/volume> - Volume Path }

This specifies the volume path for which volume efficiency metadata needs to be reverted.

[-version <revert version>] - Revert to Version

Specifies the version of Data ONTAP to which the volume efficiency metadata needs to be formatted.

[-d, -delete {true|false}] - Delete Existing Metafile on Revert

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

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

If set to `true`, this parameter Specifies that the volume efficiency metadata needs to be reverted to minor version of Data ONTAP. By default this parameter is set to `false`.

[-check-snapshot {true|false}] - Revert ignore snapshots

If set to `false`, this parameter specifies that the volume efficiency revert will not check for Snapshot copies created by previous releases of Data 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 8.3.

```
cluster1::> volume efficiency revert-to -vserver vs1 -volume vol1 -version 8.3
```

```
cluster1::> volume efficiency revert-to -vserver vs1 -path /vol/vol1 -version 8.3
```

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 Data 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 Data 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, Data 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.

| [-l]

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.

[-compression-type {none|secondary|adaptive}] - Compression Type

Displays information about the type of compression on the volume[adaptive or secondary].

[-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.

[-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.

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 00:00:14
                Type: Regular
                Schedule: sun-sat@0
Efficiency Policy Name: -
Blocks Skipped Sharing: 0
  Last Operation State: Success
Last Success Operation Begin: Mon Nov 15 20:13:26 UTC 2010
  Last Success Operation End: Mon Nov 15 20:13:26 UTC 2010
    Last Operation Begin: Mon Nov 15 20:13:26 UTC 2010
      Last Operation End: Mon Nov 15 20:13:26 UTC 2010
        Last Operation Size: 0.00B
          Last Operation Error: -
            Change Log Usage: 0%
              Logical Data Size: 156KB
                Logical Data Limit: 50.00TB
                  Logical Data Percent: 0%
                    Queued Job: -
Stale Fingerprint Percentage: 0
  Compression: false
    Inline Compression: false
      Constituent Volume: false
        Inline Dedupe: false
          Data Compaction: false
Cross Volume Inline Deduplication: false
ross Volume Background Deduplication: 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 gatherer 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)

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)

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)

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)

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)

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

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 voll -vserver vs1
```

```
cluster1::> volume efficiency start -scan-old-data -volume voll -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.

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 Data 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

This specifies the Vserver on which the volume is located.

{ -volume <volume name> - Volume Name

This specifies the volume name.

| -path </vol/volume> - Volume Path }

This specifies the volume path.

[-C, -compression <true>] - Decompress Data in the Volume

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

Undo the effects of deduplication. This requires efficiency to be disabled (by performing [volume efficiency off](#)).

[-i, -inode <integer>] - Inode Number to Undo Sharing

Remove the block sharings from a specified inode.

[-t, -undo-type {all|wrong}] - Selective Undo

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

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

Undo the effects of data compaction.

[-A, -cross-volume-dedupe <true>] - Undo Cross Volume Deduplication

Undo the effects of cross volume deduplication.

[-X, -extended-compression <true>] - Extended compression

Undo the effects of extended compression. This removes the compression savings for data that requires more resources to compress.

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 an SVM named vs1:

```
cluster1::> volume efficiency undo -vserver vs1 -volume vol1 -data  
-compaction
```

Related Links

- [volume efficiency off](#)

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 <Efficiency policy type>] - 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 <Efficiency QoS policy>] - 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 <Efficiency policy type>`] - 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 <Efficiency QoS policy>`] - 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 <Efficiency policy type>] - 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` <Efficiency QoS policy>] - 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.
```

Copyright information

Copyright © 2022 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.