



storage aggregate commands

ONTAP 9.11.1 commands

NetApp
May 23, 2023

Table of Contents

storage aggregate commands	1
storage aggregate add-disks	1
storage aggregate auto-provision	4
storage aggregate create	6
storage aggregate delete	10
storage aggregate mirror	11
storage aggregate modify	12
storage aggregate offline	16
storage aggregate online	16
storage aggregate remove-stale-record	17
storage aggregate rename	17
storage aggregate restrict	18
storage aggregate scrub	18
storage aggregate show-auto-provision-progress	20
storage aggregate show-cumulated-efficiency	21
storage aggregate show-efficiency	26
storage aggregate show-resync-status	32
storage aggregate show-scrub-status	33
storage aggregate show-space	35
storage aggregate show-spare-disks	41
storage aggregate show-status	47
storage aggregate show	49
storage aggregate verify	62
storage aggregate efficiency show	63
storage aggregate efficiency stat	64
storage aggregate efficiency cross-volume-dedupe revert-to	66
storage aggregate efficiency cross-volume-dedupe show	66
storage aggregate efficiency cross-volume-dedupe start	69
storage aggregate efficiency cross-volume-dedupe stop	70
storage aggregate encryption show-key-id	70
storage aggregate inode-upgrade resume	71
storage aggregate inode-upgrade show	72
storage aggregate object-store attach	73
storage aggregate object-store mirror	74
storage aggregate object-store modify	75
storage aggregate object-store show-freeing-status	76
storage aggregate object-store show-resync-status	77
storage aggregate object-store show-space	78
storage aggregate object-store show	79
storage aggregate object-store unmirror	81
storage aggregate object-store config create	81
storage aggregate object-store config delete	85
storage aggregate object-store config modify	85

storage aggregate object-store config rename	88
storage aggregate object-store config show	88
storage aggregate object-store profiler abort	90
storage aggregate object-store profiler show	91
storage aggregate object-store profiler start	93
storage aggregate object-store put-rate-limit modify	93
storage aggregate object-store put-rate-limit show	94
storage aggregate plex delete	95
storage aggregate plex offline	95
storage aggregate plex online	96
storage aggregate plex show	97
storage aggregate reallocation quiesce	99
storage aggregate reallocation restart	99
storage aggregate reallocation schedule	100
storage aggregate reallocation show	101
storage aggregate reallocation start	103
storage aggregate reallocation stop	104
storage aggregate relocation show	104
storage aggregate relocation start	105
storage aggregate resynchronization modify	106
storage aggregate resynchronization show	107
storage aggregate resynchronization options modify	109
storage aggregate resynchronization options show	109

storage aggregate commands

storage aggregate add-disks

Add disks to an aggregate

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

Description

The `storage aggregate add-disks` command adds disks to an existing aggregate. You must specify the number of disks or provide a list of disks to be added. If you specify the number of disks without providing a list of disks, the system selects the disks.

Parameters

-aggregate <aggregate name> - Aggregate

This parameter specifies the aggregate to which disks are to be added.

[-diskcount <integer>] - Disk Count

This parameter specifies the number of disks that are to be added to the aggregate.

{ [-T, -disktype {ATA | BSAS | FCAL | FSAS | LUN | MSATA | SAS | SSD | VMDISK | SSD-NVM | SSD-CAP | SSD-ZNS | VMLUN | VMLUN-SSD}] - Disk Type

This parameter specifies the type of disk that is to be added. It must be specified with the `-diskcount` parameter when adding disks to a Flash Pool.

Use this parameter when adding spare SSDs to an aggregate to convert it to a Flash Pool.



Only the aggregates marked as `hybrid-enabled` can be converted to Flash Pools. Use [storage aggregate modify](#) command to mark the aggregate as `hybrid-enabled`.



When this parameter is used, disk selection is not influenced by RAID options `raid.mix.hdd.disktype.capacity`, `raid.mix.hdd.disktype.performance`, or `raid.mix.disktype.solid_state`. Only disks of the specified type are considered eligible for selection.

[-C, -diskclass {capacity | performance | archive | solid-state | array | virtual}] - Disk Class

This parameter specifies the class of disk that is to be added. All disks that belong to the specified class are considered eligible for selection. The possible values are:

- `capacity` = Capacity-oriented, near-line disk types. Includes disk types FSAS, BSAS and ATA.
- `performance` = Performance-oriented, enterprise class disk types. Includes disk types FCAL and SAS.
- `archive` = Archive class SATA disks in multi-disk carrier storage shelves. Includes disk type MSATA.
- `solid-state` = Solid-state drives. Includes disk type SSD, SSD-CAP and SSD-NVM.
- `array` = Logical storage devices backed by storage arrays and used by Data ONTAP as disks. Includes disk type LUN.

- `virtual` = Virtual disks that are formatted and managed by the hypervisor. Includes disk type `VMDISK`.



When this parameter is used, disk selection is not influenced by RAID options `raid.mix.hdd.disktype.capacity`, `raid.mix.hdd.disktype.performance`, or `raid.mix.disktype.solid_state`.

[`-checksumstyle <aggrChecksumStyle>`] - Checksum Style

This parameter specifies the checksum style for the disks to be added to an aggregate. It is not applicable if `-disklist` or `-mirror-disklist` is specified. The possible values are `block` for block checksum and `advanced_zoned` for advanced zoned checksum (AZCS). By default, disks with the same checksum style as the aggregate are selected. This behavior can be overridden by using this parameter to create a mixed checksum aggregate. A mixed checksum aggregate can support only the `block` and `advanced_zoned` checksum styles.

[`-disksize <integer>`] - Disk Size(GB)

This parameter specifies the size, in GB, of the disks that are to be added to the aggregate. Disks with a usable size between 90% and 105% of the specified size are selected.

[`-d, -disklist <disk path name>,...`] - Disks

This parameter specifies a list of disks to be added. If you specify the `-disklist` parameter, you cannot further qualify the list of disks to be added by count, checksum style, size or type.

[`-mirror-disklist <disk path name>,...`] - Disks for Mirrored Plex

This parameter specifies a list of mirror disks to be added. It must contain the same number of disks specified in `-disklist` parameter. If you specify the `-mirror-disklist` parameter, you cannot further qualify the list of disks to be added by count, checksum style or type.

{ [`-ignore-pool-checks <true>`] - Don't Enforce Plex Pool Best Practices

The disks in a plex are normally required to come from the same SyncMirror pool. This behavior can be overridden with this parameter when it is set to `true`.

[`-f, -allow-mixed-rpm <true>`] - Allow Disks With Different RPM Values

This parameter specifies whether disks that have different RPM values can be added. For example, SAS disks can rotate at 10,000 or 15,000 RPM. If this parameter is set to `true` and a list of disks are provided by using the `-disklist` parameter, the disks will be added even if the SAS disks you specify have different RPM values. This parameter works similarly for ATA disks, which can rotate at 5,400 or 7,200 RPM.

+

NOTE: This parameter is applicable only when the `-disklist` or `-mirror-disklist` parameter is used.

[`-allow-same-carrier <true>`] - Allow Same RAID Group Within Carrier

This parameter can be used to allow two disks housed in the same carrier to be in the same RAID group when you add disks to an aggregate.

Having disks in the same carrier in the same RAID group is not desirable because a carrier failure can cause a simultaneous outage for two disks in the same RAID group. You can add a disk to an aggregate that causes this situation, but when an alternate disk becomes available, Data ONTAP automatically initiates a series of disk copy operations to put the disks into different RAID groups. For this reason, you should use this parameter only when necessary. When possible, allow Data ONTAP to choose disks that need to be added to the aggregate.

This parameter affects only the add-disks operation. It is not a persistent attribute of the aggregate.

[`-storage-pool <storage pool name>`] - Storage Pool

This parameter specifies the name of the SSD storage pool from which available allocation units are added to a given aggregate. This parameter cannot be used with the `-disk-list` or `-disk-count` parameters.

[`-allocation-units <integer>`] - Allocation Units }

This parameter specifies the number of allocation units to be added to a given aggregate from an SSD storage pool. Number of allocation units available and size of each unit can be found using the [storage pool show-available-capacity](#) command. This parameter works only when you also use the `-storage-pool` parameter.

[`-n, -simulate <true>`] - Simulate Addition of Disks

This parameter is used with the `disktype` and `diskcount` parameters to determine which disks would be added without actually performing the addition of disks operation.

[`-g, -raidgroup {new|all|<raidgroup>}`] - RAID Group

This parameter enables the administrator to specify which RAID group will receive the added disks. If this parameter is not used, the disks are added to the most recently created RAID group until it is full, then new RAID groups are created and filled until all the disks are added. If a RAID group name `rgX` is specified, the disks are added to that RAID group. If `new` is specified, the disks are added to a new RAID group, even if the disks would fit into an existing RAID group. If `all` is specified, the disks are added to existing RAID groups until all existing RAID groups are full. Then Data ONTAP creates one or more new RAID groups and adds the remaining disks to the new groups. If the disk type or checksum style parameters are specified with this parameter, the command operates only on the RAID groups with the matching disk type or checksum style, even if `all` is specified.

[`-cache-raid-group-size <integer>`] - RAID Group Size for Cache Tier

This parameter specifies the maximum number of disks that can be included in an SSD RAID group for this aggregate.



This parameter is applicable only when adding SSDs for the first time to a hybrid-enabled aggregate. If this parameter is not used when the first SSDs are added to the aggregate, the maximum RAID group size for the SSD cache is set to the default SSD RAID group size for the RAID type of the SSD cache.

[`-t, -raidtype {raid_tec|raid_dp|raid4|raid_ep}`] - RAID Type

This parameter specifies the type for the new RAID groups that would be created while adding disks to the aggregate. Use this parameter when you add the first RAID group comprised of SSDs to a hybrid-enabled aggregate. The values are `raid4` for RAID4, `raid_dp` for RAID Double Parity, and `raid_tec` for RAID-TEC. The default value is the type of RAID groups of the aggregate, except for RAID-TEC hybrid-enabled aggregates where the SSD tier will default to `raid_dp`. An aggregate might include a mix of different RAID types.

Examples

The following example adds 10 disks to an aggregate named `aggr0`. The disks are added to a RAID group named `rg1`:

```
cluster1::> storage aggregate add-disks -aggregate aggr0 -diskcount 10
-raidgroup rg1
```

In this example, an aggregate is converted to a Flash Pool aggregate using SSD capacity from a storage pool. The aggregate was created using RAID-DP for the hard disks and the SSDs are added using RAID4.

```
cluster1::> storage aggregate add-disks -aggregate FlashPool -storage-pool
SP1 -allocation-units 1 -raidtype raid4
```

Related Links

- [storage aggregate modify](#)
- [storage pool show-available-capacity](#)

storage aggregate auto-provision

Recommend and create new aggregates in the cluster

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

Description

This command analyzes available spare disks in the cluster, and it provides a recommendation how spare disks should be used to create aggregates according to best practices. The command prints the summary of recommended aggregates including their names and usable size. It then prompts the user whether or not the aggregates should be created as recommended. On positive response, ONTAP creates the aggregates as described in the recommendation.

The command parameters allow to restrict the command to some nodes in the cluster, print more details about recommended aggregates, and to skip the prompt.

Parameters

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

Comma separated list of node names to which the command applies. If this parameter is not used, the command applies to all nodes in the cluster.

[`-verbose` <true>] - Report More Details

Report additional details about recommended aggregates and spare disks. Per node summary shows number and total size of aggregates to create, discovered spares, and also remaining spare disks and partitions after aggregate creation. RAID group layout shows how spare disks and partitions will be used in new data aggregates to be created. The last table shows spare disks and partitions remaining unused after aggregate creation.

[`-skip-confirmation` <true>] - Skip the Confirmation and Create Recommended Aggregates

When this parameter is used, the command automatically creates the recommended aggregates. When this parameter is not used, the command checks to proceed with aggregate creation or not.



The command is not affected by the CLI session setting: `set`-confirmations` on/off`.

Examples

```
cluster1::storage aggregate> auto-provision
```

Node	New Data Aggregate	Usable Size
node1	node1_SSD_1	3.66TB
node2	node2_SSD_1	3.66TB
Total:		2 new data aggregates 7.32TB

```
Do you want to create recommended aggregates? {y|n}: n
```

```
cluster1::storage aggregate> auto-provision -verbose
```

Per node summary of new aggregates to create, discovered spares, and also remaining spare disks and partitions after aggregate creation:

Node	New Aggrs	Total New Usable Size	-Discovered Spare- Disks	-Remaining Spare- Partitions	Disks	Partitions
node1	1	3.66TB	6	48	1	0
node2	1	3.66TB	6	48	1	0
Total:	2	7.32TB	12	96	2	0

New data aggregates to create with counts of disks and partitions to be used:

Node	Usable New Data Aggregate	-Devices To Use- Mirrored	Size	Disks
node1	node1_SSD_1	false	3.66TB	5
node2	node2_SSD_1	false	3.66TB	5

RAID group layout showing how spare disks and partitions will be used in new data aggregates to be created:

RAID Group In New Data Aggregate To Be Created	Parity	Count	Disk Type	Usable Size	Disk Or Partition Data
node1_SSD_1		5		3.66TB	
node2_SSD_1		5		3.66TB	


```

-----
/node1_SSD_1/plex0/rg0      SSD      81.97GB partition  22
2
/node1_SSD_1/plex0/rg1      SSD      81.97GB partition  22
2
/node1_SSD_1/plex0/rg2      SSD      185.5GB disk       3
2
/node2_SSD_1/plex0/rg0      SSD      81.97GB partition  22
2
/node2_SSD_1/plex0/rg1      SSD      81.97GB partition  22
2
/node2_SSD_1/plex0/rg2      SSD      185.5GB disk       3
2

```

Details about spare disks and partitions remaining after aggregate creation:

Disk Node	Device	Disk Or Type	Pool Usable Size	Remaining Partition Number	Spares
node1	SSD	185.5GB disk	Pool0	1	1
node2	SSD	185.5GB disk	Pool0	1	1

Do you want to create recommended aggregates? {y|n}: y

Info: Aggregate auto provision has started. Use the "storage aggregate show-auto-provision-progress" command to track the progress.

Related Links

- [set](#)

storage aggregate create

Create an aggregate

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

Description

The `storage aggregate create` command creates an aggregate. An aggregate consists of disks. You must specify the number of disks or provide a list of disks to be added to the new aggregate. If you specify the number of disks without providing a list of disks, the system selects the disks.

When creating an aggregate, you can optionally specify the aggregate's home node, the RAID type for RAID groups on the aggregate, and the maximum number of disks that can be included in a RAID group.

Parameters

-aggregate <aggregate name> - Aggregate

This parameter specifies the name of the aggregate that is to be created.

[-chksumstyle <aggrChecksumStyle>] - Checksum Style

This parameter specifies the checksum style for the aggregate. The values are *block* for Block Checksum and *advanced_zoned* for Advanced Zoned Checksum (AZCS).

-diskcount <integer> - Number Of Disks

This parameter specifies the number of disks that are to be included in the aggregate, including the parity disks. The disks in this newly created aggregate come from the pool of spare disks. The smallest disks in this pool are added to the aggregate first, unless you specify the `-disksize` parameter.

[-R, -diskrpm <integer>] - Disk RPM

This parameter specifies the RPM of the disks on which the aggregate is to be created. The possible values include 5400, 7200, 10000, and 15000.

[-disksize <integer>] - Disk Size(GB)

This parameter specifies the size, in GB, of the disks on which the aggregate is to be created. Disks with a usable size between 90% and 105% of the specified size are selected.

{ [-T, -disktype {ATA | BSAS | FCAL | FSAS | LUN | MSATA | SAS | SSD | VMDISK | SSD-NVM | SSD-CAP | SSD-ZNS | VMLUN | VMLUN-SSD}] - Disk Type

This parameter specifies the type of disk on which the aggregate is to be created.



When this parameter is used, disk selection is not influenced by RAID options `raid.mix.hdd.disktype.capacity`, `raid.mix.hdd.disktype.performance`, or `raid.mix.disktype.solid_state`. Only disks of the specified type are considered eligible for selection.

[-C, -diskclass {capacity | performance | archive | solid-state | array | virtual}] - Disk Class

This parameter specifies the class of disks on which the aggregate is to be created. All disks that belong to the specified class are considered eligible for selection. The possible values are:

- `capacity` = Capacity-oriented, near-line disk types. Includes disk types FSAS, BSAS and ATA.
- `performance` = Performance-oriented, enterprise class disk types. Includes disk types FCAL and SAS.
- `archive` = Archive class SATA disks in multi-disk carrier storage shelves. Includes disk type MSATA.
- `solid-state` = Solid-state drives. Includes disk type SSD, SSD-CAP and SSD-NVM.
- `array` = Logical storage devices backed by storage arrays and used by Data ONTAP as disks. Includes disk type LUN.
- `virtual` = Virtual disks that are formatted and managed by the hypervisor. Includes disk type VMDISK.



When this parameter is used, disk selection is not influenced by RAID options `raid.mix.hdd.disktype.capacity`, `raid.mix.hdd.disktype.performance`, or `raid.mix.disktype.solid_state`.

[-m, -mirror <true>] - Is Mirrored

This parameter specifies that the new aggregate be mirrored (have two plexes). If this parameter is set to *true*, the specified disks are split between the two plexes. By default, the new aggregate will not be mirrored. You cannot use the *-mirror* parameter when supplying a specific list of disks with either the *-disklist* or *-mirror-disklist* parameters.

[-pool <aggrSparePool>] - Spare Pool

This parameter specifies the SyncMirror pool to be used to supply the disks for the aggregate. Valid values are Pool0 or Pool1.

[-d, -disklist <disk path name>,... - Disks for First Plex

This parameter specifies a list of disks to be added to the new aggregate. If you specify the *-disklist* parameter, you cannot further qualify the list of disks to be added by count, checksum style, type, size, or RPM. You cannot use the *-disklist* parameter when the *-mirror* parameter is set to *true*.

[-mirror-disklist <disk path name>,...] - Disks for Mirrored Plex

This parameter specifies a list of mirror disks to be added to the new mirrored aggregate. It must contain the same number of disks specified in *-disklist* parameter. If you specify the *-mirror-disklist* parameter, you cannot further qualify the list of disks to be added by count, checksum style, type, size, or RPM. You cannot use the *-mirror-disklist* parameter when the *-mirror* parameter is set to *true*.

[-ignore-pool-checks <true>] - Don't Enforce Plex Pool Best Practices

The disks in a plex are normally required to come from the same SyncMirror pool. This behavior can be overridden with this parameter when it is set to *true*. This option cannot be used when the *-mirror* option is set to *true*

[-f, -allow-mixed-rpm <true>] - Allow Disks With Different RPM Values

This parameter specifies whether the aggregate can contain disks that have different RPM values. For example, SAS disks can rotate at 10,000 or 15,000 RPM. If this parameter is set to *true* and a list of disks are provided by using the *-disklist* parameter, the aggregate will be created even if the SAS disks you specify have different RPM values. This parameter works similarly for ATA disks, which can rotate at 5,400 or 7,200 RPM.

[-allow-same-carrier <true>] - Allow Same RAID Group Within Carrier

This parameter can be used to allow two disks housed in the same carrier to be in the same RAID group when you add disks to an aggregate.

Having disks in the same carrier in the same RAID group is not desirable because a carrier failure can cause a simultaneous outage for two disks in the same RAID group. You create an aggregate with this characteristic, but when an alternate disk becomes available, Data ONTAP automatically initiates a series of disk copy operations to put the disks into different RAID groups. For this reason, you should use this parameter only when necessary. When possible, allow Data ONTAP to choose the disks from which to create the aggregate.

This parameter affects only the aggregate creation operation. It is not a persistent attribute of the aggregate.

[-node {<nodename>|local}] - Node

This parameter specifies the home node for the aggregate. If this parameter is not specified, Data ONTAP selects the node where the aggregate is created.

[-s, -maxraidsize <integer>] - Max RAID Size

This parameter specifies the maximum number of disks that can be included in a RAID group.

[-t, -raidtype {raid_tec|raid_dp|raid4|raid_ep}] - RAID Type

This parameter specifies the type for RAID groups on the aggregate. The values are *raid4* for RAID4 of disks other than *SSD-ZNS*, *raid_ep* for RAID4 of *SSD-ZNS* disks plus RAID4 inside each *SSD-ZNS* disk, *raid4* for RAID4, *raid_dp* for RAID Double Parity, and *raid_tec* for RAID Triple-Erasure-Code. The default setting is *raid_dp* for disks other than *SSD-ZNS* unless the disks are HDDs with a capacity larger than 4 TB, in which case the default will be *raid_tec*. For *SSD-ZNS* disks the default setting will be *raid_ep*. This parameter is not needed for array LUNs because they are always created with the *raid0* raidtype. *raid4* and *raid_ep* are not compatible with shared disks unless the shared disks belong to a storage pool.

[-simulate <true>] - Simulate Aggregate Provisioning Operation

This option simulates the aggregate creation and prints the layout of the new aggregate.

[-force-small-aggregate <true>] - Force the Creation of a Small Aggregate

This parameter can be used to force the creation of a 2-disk RAID4 and RAID-EP aggregates, or a 3-disk or 4-disk RAID-DP aggregate.

[-is-autobalance-eligible {true|false}] - Is Eligible for Auto Balance Aggregate

This specifies whether the aggregate will be considered by the Auto Balance Aggregate feature. If the Auto Balance Aggregate feature is not used, this field is not used. When this parameter is set to *true* the Auto Balance Aggregate feature might recommend moving volumes to or from this aggregate in order to balance system workload. When this parameter is set to *false* the aggregate will not be considered as a destination for the Auto Balance Aggregate feature allowing for predictability in data placement. The default value is *false*.

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

This parameter specifies the type of SnapLock aggregate to be created. In order to create a SnapLock Compliance aggregate, specify *compliance*. To create a SnapLock Enterprise aggregate, specify *enterprise*.

[-autobalance-unbalanced-threshold-percent <integer>] - Threshold When Aggregate Is Considered Unbalanced (%)

This parameter specifies the space used threshold percentage that will cause the Auto Balance Aggregate feature to consider an aggregate as unbalanced.

[-autobalance-available-threshold-percent <integer>] - Threshold When Aggregate Is Considered Balanced (%)

This parameter specifies the threshold percentage which will determine if an aggregate is a target destination for a move. The Auto Balance Aggregate feature will attempt to move volumes from an unbalanced aggregate until it is under this percentage.

[-encrypt-with-aggr-key {true|false}] - Enable Aggregate level Encryption

This parameter specifies the data encryption policy for the contained volumes. If this parameter is set to *true*, then by default, the volumes created in this aggregate will be encrypted, using the aggregate level encryption keys.

Examples

The following example creates an aggregate named `aggr0` on a home node named `node0`. The aggregate contains 20 disks and uses RAID-DP. The aggregate contains regular FlexVol volumes:

```
cluster1::> storage aggregate create -aggregate aggr0 -nodes node0
-diskcount 20 -raidtype raid_dp -volume-style flex
```

The following example creates an aggregate named `aggr0` on a home node named `node0`. The aggregate contains the disks specified and uses RAID-DP

```
cluster1::> storage aggregate create -aggregate aggr0 -nodes node0
-disklist 1.0.15,1.0.16,1.0.17,1.0.18,1.0.19 -raidtype raid_dp
```

The following example creates an aggregate named `aggr0` on a home node named `node0`. The aggregate contains 20 disks of size 6 TB and of type FSAS and uses RAID-TEC:

```
cluster1::> storage aggregate create -aggregate aggr0 -nodes node0
-diskcount 20 -raidtype raid_tec -disksize 6000 -disktype FSAS
```

The following example creates a mirrored aggregate named `aggr0` on the local node. The aggregate contains 10 disks in each plex:

```
cluster1::> storage aggregate create -aggregate aggr0 -mirror
-diskcount 20
```

storage aggregate delete

Delete an aggregate

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

Description

The `storage aggregate delete` command deletes a storage aggregate. The command fails if there are volumes present on the aggregate. If the aggregate has an object store attached to it, then in addition to deleting the aggregate the command deletes the objects in the object store as well. No changes are made to the object store configuration as part of this command.

Parameters

-aggregate <aggregate name> - Aggregate

This parameter specifies the aggregate that is to be deleted.

[`-preserve-config-data <true>`] - Delete Physical Aggregate but Preserve Configuration Data

Deletes the physical aggregate, but preserves the aggregate configuration data. The aggregate must not have any disks associated with it. If the parameter `-preserve-config-data` is specified without a value, the default value is `true`; if this parameter is not specified, the default value is `false`.

Examples

The following example deletes an aggregate named `aggr1`:

```
cluster1::> storage aggregate delete -aggregate aggr1
```

storage aggregate mirror

Mirror an existing aggregate

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

Description

The `storage aggregate mirror` command adds a plex to an existing unmirrored aggregate. You can specify a list of disks to be used for the mirrored plex. If you do not specify the disks, the system automatically selects the disks based on the aggregate's existing plex.

Parameters

`-aggregate <aggregate name>` - Aggregate

This parameter specifies the aggregate to mirror.

[`-f, -allow-mixed-rpm <true>`] - Allow Disks With Different RPM Values

This parameter specifies whether disks that have different RPM values can be used. For example, SAS disks can rotate at 10,000 or 15,000 RPM. If this parameter is set to `true` and a list of disks are provided by using the `-mirror-disklist` parameter, the disks will be added even if the SAS disks you specify have different RPM values. This parameter works similarly for ATA disks, which can rotate at 5,400 or 7,200 RPM.

+

NOTE: This parameter is only applicable when the `-mirror-disklist` parameter is used.

[`-d, -mirror-disklist <disk path name>,...`] - Disks for Mirrored Plex

This parameter specifies a list of disks to be used for the plex to be added. It must contain the same number of disks as the existing plex of the unmirrored aggregate specified using the `-aggregate` parameter.

[`-ignore-pool-checks <true>`] - Don't Enforce Plex Pool Best Practices

For maximum reliability, all disks from a plex should come from the same SyncMirror pool, and the disks for the second plex should all come from the other pool. If needed, this behavior can be overridden by setting this parameter to `true`. This parameter can be used only with the `-mirror-disklist` parameter.

[-f, -allow-same-carrier <true>] - Allow Same RAID Group Within Carrier

This parameter can be used to allow two disks housed in the same carrier to be in the same RAID group for a mirrored aggregate. Having disks in the same carrier in the same RAID group is not desirable, because a carrier failure can cause a simultaneous outage for two disks in the same RAID group. For this reason, this configuration is not allowed by default. This restriction can be overridden by setting this parameter to *true*.

+

NOTE: This parameter is accepted only when the `-mirror-disklist` parameter is used.

[-n, -simulate <true>] - Simulate Mirroring of an Existing Aggregate

This option simulates the mirroring of an existing aggregate and prints the layout of the new plex.

Examples

The following example mirrors an unmirrored aggregate `aggr1`:

```
cluster1::> storage aggregate mirror -aggregate aggr1
```

The following example mirrors an unmirrored aggregate `aggr1`. The specified disks are used for the new plex.

```
cluster1::> storage aggregate mirror -aggregate aggr1 -mirror-disklist  
1.2.12, 1.2.14, 1.2.16
```

storage aggregate modify

Modify aggregate attributes

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

Description

The `storage aggregate modify` command can be used to modify attributes of an aggregate such as RAID type and maximum RAID group size.

Changing the RAID type immediately changes the RAID group type for all RAID groups in the aggregate.

Changing the maximum RAID size does not cause existing RAID groups to grow or to shrink; rather, it affects the size of RAID groups created in the future, and determines whether more disks can be added to the RAID group that was most recently created.

Parameters

-aggregate <aggregate name> - Aggregate

This parameter specifies the storage aggregate that is to be modified.

[-T, -disktype {ATA | BSAS | FCAL | FSAS | LUN | MSATA | SAS | SSD | VMDISK | SSD-NVM | SSD-CAP | SSD-ZNS | VMLUN | VMLUN-SSD}] - Disk Type

This parameter specifies the disk type of the RAID groups to be modified. In case of Flash Pool, it specifies either the HDD tier or the SSD tier. If the HDD tier is composed of more than one type of disk, specifying any of the disk types in use causes that tier to be modified. If the current aggregate RAID type is `mixed_raid_type`, this parameter is mandatory.

`[-free-space-realloc {on|off}] - Free Space Reallocation`

This parameter specifies whether free space reallocation is enabled on the aggregate.

Free space reallocation optimizes the free space in an aggregate immediately before Data ONTAP writes data to the blocks in that aggregate. This field is not allowed on aggregates with SSD-ZNS disks.

The default setting is `off`.

`[-ha-policy {sfo|cfo}] - HA Policy`

This parameter specifies the high-availability policy to be used in the context of a root recovery procedure. Do not modify this setting unless directed to do so by a customer support representative.

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

This parameter is used to set the space reserved for Snapshot copies to the specified value. For example, to set the snapshot reserve to 5%, you should enter ``-percent-snapshot-space` 5`.

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

This optionally specifies the percentage at which the aggregate 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 aggregate nearly full space alerts.

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

This optionally specifies the percentage at which the aggregate 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 aggregate full space alerts.

`[-hybrid-enabled {true|false}] - Hybrid Enabled`

If the `hybrid-enabled` option is set to "true", the aggregate is marked as `hybrid_enabled`, that is, the aggregate can contain a mix of SSDs and HDDs (Hard Disk Drives, e.g., SAS, SATA, and/or FC). By default, aggregates cannot be marked "hybrid_enabled" if the aggregate contains FlexVols that cannot be write cached. A FlexVol cannot be write-cached if it is part of an aggregate created in Data ONTAP 7. Use `-force-hybrid-enabled` to over-ride this behavior.

`[-f, -force-hybrid-enabled <true>] - Force Marking of Aggregate as Hybrid Enabled`

By default, aggregates cannot be marked "hybrid_enabled" if the aggregate contains FlexVols that cannot be write cached. A FlexVol cannot be write-cached if it is part of an aggregate created in Data ONTAP 7. Use `-force-hybrid-enabled` to over-ride this behavior. Note that read caching will be enabled on these FlexVols, but write caching will be disabled.

`[-s, -maxraidsize <integer>] - Max RAID Size`

This parameter specifies the maximum number of disks that can be included in a RAID group for this aggregate.



For Flash Pools, this option controls the maximum size of the HDD RAID groups.

[-cache-raid-group-size <integer>] - Flash Pool SSD Tier Maximum RAID Group Size

This parameter specifies the maximum number of disks that can be included in a SSD RAID group for this Flash Pool.



This parameter is applicable only for Flash Pools.

[-t, -raidtype {raid_tec|raid_dp|raid4|raid_ep}] - RAID Type

This parameter specifies the RAID type for RAID groups on the aggregate. The possible values are *raid4* for RAID4 of disks other than *SSD-ZNS*, *raid_ep* for RAID4 of *SSD-ZNS* disks with additional RAID4 protection inside each disk, *raid_dp* for RAID-DP, and *raid_tec* for RAID-TEC. If you change the RAID type from RAID4 or RAID-EP to RAID-DP, each RAID group allocates a spare disk for the group's second parity disk and begins a reconstruction process. If you change the RAID type from RAID-DP to RAID-TEC, each RAID group allocates a spare disk for the group's third parity disk and begins a reconstruction process. Changing the RAID type from RAID4 and RAID-EP to RAID-TEC or vice-versa is not supported. To change the RAID type from RAID4 or RAID-EP to RAID-TEC, first change from RAID4 or RAID-EP to RAID-DP and then to RAID-TEC.

[-resyncsnaptime <integer>] - SyncMirror Resync Snapshot Frequency in Minutes

This parameter sets the mirror resynchronization snapshot frequency to be the given number of minutes. The default value is 60 (minutes).

[-state <aggregate state>] - State

This deprecated parameter specifies the state of the aggregate. The possible values are as follows:

- **online** - Immediately sets the aggregate online. All volumes on the aggregate are set to the state they were in when the aggregate was taken offline or restricted. The preferred command to bring an aggregate online is `storage aggregate online`.
- **offline** - Takes an aggregate offline. You cannot take an aggregate offline if any of its volumes are online. The preferred command to take an aggregate offline is `storage aggregate offline`.
- **restricted** - Restricts the aggregate. You cannot restrict an aggregate if any of its volumes are online. The preferred command to restrict an aggregate is `storage aggregate restrict`.

[-is-autobalance-eligible {true|false}] - Is Eligible for Auto Balance Aggregate

This specifies whether the aggregate is considered by the Auto Balance Aggregate feature. If the Auto Balance Aggregate feature is not used, this field is not used. When this parameter is set to *true* the Auto Balance Aggregate feature might recommend moving volumes to or from this aggregate in order to balance system workload. When this parameter is set to *false* the aggregate will not be considered as a destination for the Auto Balance Aggregate feature allowing for predictability in data placement. The default value is *false*.

[-autobalance-unbalanced-threshold-percent <integer>] - Threshold When Aggregate Is Considered Unbalanced (%)

This parameter sets the space used threshold percentage that will cause the Auto Balance Aggregate feature to consider an aggregate as unbalanced.

[-autobalance-available-threshold-percent <integer>] - Threshold When Aggregate Is Considered Balanced (%)

This parameter sets the threshold percentage which will determine if an aggregate is a target destination for a move. The Auto Balance Aggregate feature will attempt to move volumes from an unbalanced aggregate until it is under this percentage.

[`-resync-priority` {`high(fixed)`|`high`|`medium`|`low`}] - Resynchronization Priority

This parameter specifies the new resynchronization priority value for the specified aggregate. This field cannot be modified for unmirrored or Data ONTAP system aggregates.

Possible values for this parameter are:

- `high`: Mirrored data aggregates with this priority value start resynchronization first.
- `medium`: Mirrored data aggregates with this priority value start resynchronization after all the system aggregates and data aggregates with 'high' priority value have started resynchronization.
- `low`: Mirrored data aggregates with this priority value start resynchronization only after all the other aggregates have started resynchronization.

[`-single-instance-data-logging` {`off`|`on`}] - Enable SIDL

This parameter specifies whether Single Instance Data Logging feature is enabled on the aggregate and the constituent volumes on the aggregate. This feature improves user write performance by optimizing the amount of data nlogged by user writes on platforms where NVRAM and secondary storage are of same media type.

[`-is-inactive-data-reporting-enabled` {`true`|`false`}] - Inactive Data Reporting Enabled

This parameter specified whether the reporting of how much user data is inactive should be enabled on the aggregate and volumes on the aggregate. This parameter is not allowed on FabricPools.

[`-encrypt-with-aggr-key` {`true`|`false`}] - Enable Aggregate level Encryption

This parameter specifies that the volumes within the new aggregate can be encrypted with aggregate keys. If this parameter is set to `true`, the aggregate will support encryption with aggregate keys.

[`-force-disable-encrypt-with-aggr-key` <`true`>] - Force disable NAE. Skip aggregate snapshot check.

This parameter allows disabling NetApp Aggregate Encryption (NAE) on an aggregate if the user is certain there is no aggregate snapshot for that aggregate containing NAE volumes. If the parameter is set to `true`, aggregate snapshot check is skipped and NAE is disabled.

[`-azcs-read-optimization` {`on`|`off`}] - azcs read optimization

This parameter specifies whether azcs read optimization is enabled on the aggregate. This feature improves read performance on cloud platforms.

Examples

The following example changes all RAID groups on an aggregate named `aggr0` to use RAID-DP:

```
cluster1::> storage aggregate modify -aggregate aggr0 -raidtype raid_dp
```

The following example changes all RAID groups with FSAS disks in an aggregate named `aggr0` to use RAID-TEC:

```
cluster1::> storage aggregate modify -aggregate aggr0 -disktype FSAS  
-raidtype raid_tec
```

storage aggregate offline

Offline an aggregate

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

Description

The `storage aggregate offline` command takes an aggregate offline.

If you are taking a root aggregate offline, the node owning the aggregate must be in maintenance mode.

Parameters

-aggregate <aggregate name> - Aggregate

The name of the aggregate to be taken offline.

Examples

The following example takes an aggregate named `aggr1` offline:

```
cluster1::> storage aggregate offline -aggregate aggr1
```

storage aggregate online

Online an aggregate

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

Description

The `storage aggregate online` command brings an aggregate online if the aggregate is in offline or restricted state. If an aggregate is in an inconsistent state, it must be brought to a consistent state before it can be brought online. If you have an aggregate that is in an inconsistent state, contact technical support.

Parameters

-aggregate <aggregate name> - Aggregate

The name of the aggregate to be brought online.

Examples

The following example brings an aggregate named `aggr1` online:

```
cluster1::> storage aggregate online -aggregate aggr1
```

storage aggregate remove-stale-record

Remove a stale aggregate record

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

Description

The `storage aggregate remove-stale-record` command removes a stale storage aggregate record on disk. A stale aggregate record refers to an aggregate that has been removed from the storage system, but whose information remains recorded on disk. Stale aggregate records are displayed in the `nodeshell aggr status -r` command, but the `storage aggregate show` command does not show the aggregate as hosted on that node.

Parameters

-aggregate <aggregate name> - Aggregate

This parameter specifies the aggregate that corresponds to the stale aggregate record that is to be deleted.

-nodename {<nodename>|local} - Node Name

This parameter specifies the node that contains the aggregate.

Examples

The following example removes a stale aggregate record that refers to aggregate "aggr1":

```
cluster1::> storage aggregate remove-stale-record -aggregate aggr1
-nodename node1
```

storage aggregate rename

Rename an aggregate

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

Description

The `storage aggregate rename` command renames an aggregate.

Parameters

-aggregate <aggregate name> - Aggregate

This parameter specifies the aggregate to be renamed.

-newname <aggregate name> - New Name

This parameter specifies the new name for the aggregate.

Examples

The following example renames an aggregate named `aggr5` as `sales-aggr`:

```
cluster1::> storage aggregate rename -aggregate aggr5 -newname sales-aggr
```

storage aggregate restrict

Restrict an aggregate

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

Description

The `storage aggregate restrict` command puts an aggregate in restricted state to make data in the aggregate's volumes unavailable to clients. When an aggregate is in restricted state data access is not allowed. However, few operations such as aggregate copy, parity recomputation, scrub and RAID reconstruction are allowed. You can also use this command if you want the aggregate to be the target of an aggregate copy or SnapMirror replication operation.

Parameters

-aggregate <aggregate name> - Aggregate

The name of the aggregate to be restricted.

Examples

The following example restricts an aggregate named `aggr1`:

```
cluster1::> storage aggregate restrict -aggregate aggr1
```

storage aggregate scrub

Aggregate parity scrubbing

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

Description

The `storage aggregate scrub` command scrubs an aggregate for media and parity errors. Parity scrubbing compares the data disks to the parity disks in their RAID group and corrects the parity disks contents, as required. If no name is given, parity scrubbing is started on all online aggregates.



By default, scrubs are scheduled to run for a specified time on a weekly basis. However, you can use this command to run scrubs manually to check for errors and data inconsistencies.

Parameters

{ **-aggregate <aggregate name>** - **Aggregate**

This parameter specifies the aggregate to be scrubbed for errors.

[**-plex <text>**] - **Plex**

This parameter specifies the name of the plex to scrub. If this parameter is not specified, the command scrubs the entire aggregate.

[**-raidgroup <text>**] - **RAID Group**

This parameter specifies the RAID group to be scrubbed. If this parameter is not specified, the command scrubs the entire aggregate.

+

NOTE: This parameter is only applicable when the `-plex` parameter is used.

| **-node {<nodename>|local}** - **Node }**

This parameter specifies the name of the node associated with the aggregate to be scrubbed. The value `local` specifies the current node.

-action {start|stop|resume|suspend|status} - **Action**

This parameter specifies the action to be taken. The possible actions are:

- `start` - Starts a scrub.
- `stop` - Permanently stops a scrub. A stopped scrub cannot be resumed.
- `resume` - Resumes a suspended parity scrub.
- `suspend` - Suspends a parity scrub.
- `status` - Displays the current status of a scrub.

Examples

The following example starts a scrub on a RAID group named `rg0` of plex named `plex0` on an aggregate named `aggr0`:

```
cluster1::> storage aggregate scrub -aggregate aggr0 -raidgroup rg0 -plex
plex0 -action start
```

The following example queries the status of a scrub:

```
cluster1::> storage aggregate scrub -aggregate aggr0 -raidgroup rg0 -plex
plex0 -action status
```

```
Raid Group:/aggr0/plex0/rg0, Is Suspended:false, Last Scrub:Sun Nov 13
01:30:55 2011
, Percentage Completed:7%
```

The following example starts a scrub on plex1 of an aggregate named aggr1:

```
cluster1::> storage aggregate scrub -aggregate aggr1 -plex plex1 -action
start
```

The following example queries the status of plex1 of an aggregate named aggr1:

```
cluster1::> storage aggregate scrub -aggregate aggr1 -plex plex1 -action
status

Raid Group:/aggr1/plex1/rg0, Is Suspended:false, Last Scrub:Sun Nov 13
02:07:29
2011
, Percentage Completed:1%
```

The following example queries the status of all the plexes for an aggregate named aggr1:

```
cluster1::> storage aggregate scrub -aggregate aggr1 -action status

Raid Group:/aggr1/plex0/rg0, Is Suspended:false, Last Scrub:Sun Nov 13
01:58:06
2011

Raid Group:/aggr1/plex1/rg0, Is Suspended:false, Last Scrub:Sun Nov 13
02:07:29
2011
, Percentage Completed:4%
```

storage aggregate show-auto-provision-progress

Display aggregate auto provision status

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

Description

The `storage aggregate show-auto-provision-progress` command displays the status of the most recent auto provision operation. The command output displays the progress for all the aggregates included in the provisioning operation. The command displays the following information about each aggregate:

- Aggregate
- Current Usable Size
- Target Usable Size
- Provisioning Progress

Examples

The following example displays the information about all aggregates that are provisioned during the aggregate auto provision operation:

```
cluster1::> aggr auto-provision
Node                New Data Aggregate                Usable Size
-----
node1                node1_SSD_1                        23.65GB
node2                node2_SSD_1                        23.65GB
-----
Total:                2    new data aggregates            47.30GB

Do you want to create recommended aggregates? {y|n}: y

Info: Aggregate auto provision has started. Use the "storage aggregate
      show-auto-provision-progress" command to track the progress.

cluster1::> storage aggregate show-auto-provision-progress
                Current                Target
Aggregate      Usable Size  Usable Size  Provisioning
Progress
-----
node1_SSD_1    0B          23.65GB    Creating
node2_SSD_1    0B          23.65GB    Creating

cluster1::> storage aggregate show-auto-provision-progress
                Current                Target
Aggregate      Usable Size  Usable Size  Provisioning
Progress
-----
node1_SSD_1    23.65GB     23.65GB    Completed
node2_SSD_1    23.65GB     23.65GB    Completed
```

storage aggregate show-cumulated-efficiency

Display cumulated storage efficiency details

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

Description

The `storage aggregate show-cumulated-efficiency` command displays information about the cumulated storage efficiency of all the aggregates. The storage efficiency is displayed at four different levels:

- Total
- Aggregate
- Volume
- Snapshot and FlexClone 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.

| [-details]

Use this parameter to show additional Storage Efficiency Ratios.

| [-all-details]

Use this parameter to show additional Storage Efficiency Ratios and size values.

| [-instance] }

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

[-aggregates <aggregate name>,...] - List of Aggregates to cumulate Storage Efficiency ratio

If this parameter is specified, the command calculates the cumulated storage efficiency of the specified list of aggregates.

[-nodes {<nodename>|local}] - List of Aggregates to cumulate Storage Efficiency ratio

If this parameter is specified, the command calculates the cumulated storage efficiency of aggregates that are located on the specified list of node.

[-total-logical-used {<integer>[KB|MB|GB|TB|PB]}] - Logical Size Used by volumes, clones, Snapshot copies in the Aggregate

Displays the total logical size used in all the specified aggregates. This includes Volumes, Clones and Snapshots in all the specified aggregates. The logical size is computed based on physical usage and savings obtained in all the specified aggregates.

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

Displays the physical size used by all the specified aggregates.

[-total-storage-efficiency-ratio <text>] - Total Storage Efficiency Ratio

Displays the total storage efficiency ratio of the aggregate.

[-total-data-reduction-logical-used-wo-snapshots {<integer>[KB|MB|GB|TB|PB]}] - Total Data Reduction Logical Used Without Snapshots

Displays the total logical size used in all the specified aggregates excluding Snapshot copies.

[-total-data-reduction-physical-used-wo-snapshots {<integer>[KB|MB|GB|TB|PB]}] - Total Data Reduction Physical Used Without Snapshots

Displays the total physical size used by all the specified aggregates excluding Snapshot copies.

[-total-data-reduction-efficiency-ratio-wo-snapshots <text>] - Total Data Reduction Efficiency Ratio Without Snapshots

Displays the total storage efficiency ratio obtained by Deduplication, Compression, Data Compaction, Pattern Detection and FlexClone data reduction technologies excluding snapshot copies on the specified aggregates.

[-total-data-reduction-logical-used-wo-snapshots-flexclones {<integer>[KB|MB|GB|TB|PB]}] - Total Data Reduction Logical Used without snapshots and flexclones

Displays the total logical size used in all the specified aggregates excluding Snapshot copies and FlexClones.

[-total-data-reduction-physical-used-wo-snapshots-flexclones {<integer>[KB|MB|GB|TB|PB]}] - Total Data Reduction Physical Used without snapshots and flexclones

Displays the total physical size used by all the specified aggregates aggregates excluding Snapshot copies and FlexClones.

[-total-data-reduction-efficiency-ratio-wo-snapshots-flexclones <text>] - Total Data Reduction Efficiency Ratio without snapshots and flexclones

Displays the total storage efficiency ratio obtained by Deduplication, Compression, Data Compaction, Pattern Detection data reduction technologies excluding snapshot copies on the specified aggregates.

[-total-performance-tier-data-reduction-physical-used-wo-snapshots-flexclones {<integer>[KB|MB|GB|TB|PB]}] - Total Data Reduction Physical Used without snapshots and flexclones in the FabricPool Performance Tier

Displays the total performance tier physical size used by all the specified aggregates aggregates excluding Snapshot copies and FlexClones.

[-volume-logical-used {<integer>[KB|MB|GB|TB|PB]}] - Logical Space Used for All volumes

Displays the total logical size used by all the volumes in all the specified aggregates.

[-volume-physical-used {<integer>[KB|MB|GB|TB|PB]}] - Physical Space Used for All volumes

Displays the total physical size used by all volumes in all the specified aggregates.

[-volume-dedupe-zero-pattern-saved {<integer>[KB|MB|GB|TB|PB]}] - Space Saved by volume Deduplication and pattern detection

Displays the total disk space that is saved by deduplication, Zero pattern detection and FlexClone for files, LUNs or NVMe namespaces by all volumes in all the specified aggregates.

[-volume-efficiency-saved-ratio <text>] - Volume Deduplication Savings ratio

Displays the storage efficiency ratio for savings by deduplication and FlexClone for files, LUNs or NVMe namespaces by all volumes in all the specified aggregates.

[-volume-compression-saved {<integer>[KB|MB|GB|TB|PB]}] - Space Saved by volume Compression

Displays the total disk space that is saved by compressing blocks by all volumes in all the specified aggregates.

[-volume-compression-saved-ratio <text>] - Volume Compression Savings ratio

Displays the storage efficiency ratio for savings by compressing blocks on all volumes in all the specified

aggregates.

[-volume-data-reduction-storage-efficiency-ratio <text>] - Volume Data Reduction SE Ratio

Displays the storage efficiency ratio of all the volumes in all the specified aggregates.

[-aggr-logical-used {<integer>[KB|MB|GB|TB|PB]}] - Logical Space Used by the Aggregate

Displays the logical size used by all the specified aggregates.

[-aggr-physical-used {<integer>[KB|MB|GB|TB|PB]}] - Physical Space Used by the Aggregate

Displays the physical size used by all the specified aggregates.

[-aggr-data-reduction-storage-efficiency-ratio <text>] - Aggregate Data Reduction SE Ratio

Displays the storage efficiency ratio of the aggregate.

[-snapshot-logical-used {<integer>[KB|MB|GB|TB|PB]}] - Logical Size Used by Snapshot copies

Displays the logical size used by all Volume Snapshots residing in all the specified aggregates.

[-snapshot-physical-used {<integer>[KB|MB|GB|TB|PB]}] - Physical Size Used by Snapshot copies

Displays the physical size used by all Volume Snapshots residing in all the specified aggregates.

[-snapshot-volume-data-reduction-storage-efficiency-ratio <text>] - Snapshot volume Data Reduction Ratio

Displays the Snapshot volume storage efficiency ratio of the aggregate.

[-flexclone-volume-logical-used {<integer>[KB|MB|GB|TB|PB]}] - Logical Size Used by FlexClone volumes

Displays the logical size used by all FlexClone volumes residing in all the specified aggregates.

[-flexclone-volume-physical-used {<integer>[KB|MB|GB|TB|PB]}] - Physical Sized Used by FlexClone volumes

Displays the physical size used by all FlexClone volumes in all the specified aggregates.

[-flexclone-volume-data-reduction-storage-efficiency-ratio <text>] - FlexClone volume Data Reduction Ratio

Displays the FlexClone volume storage efficiency ratio of the aggregate.

[-snapshot-flexclone-volume-data-reduction-storage-efficiency-ratio <text>] - Snapshot And FlexClone volume Data Reduction SE Ratio

Displays the Snapshot and FlexClone volume storage efficiency ratio of the aggregate.

[-number-of-offline-volumes <integer>] - Number of volumes Offline

Displays the number of volumes that are offline in all the specified aggregates.

[-number-of-sis-disabled-volumes <integer>] - Number of SIS Disabled volumes

Displays the number of volumes on which volume efficiency is disabled in all the specified aggregates.

[-number-of-sis-change-log-disabled-volumes <integer>] - Number of SIS Change Log Disabled volumes

Displays the number of volumes on which efficiency change log is disabled in all the specified aggregates. The scheduled background Deduplication will be disabled on these volumes.

[-number-of-skipped-aggregates <integer>] - Number of Skipped Aggregates

Displays the number of aggregates that were skipped for calculating the cumulated storage efficiency.

[-skipped-aggregates <aggregate name>, ...] - List of Aggregates skipped

Displays the list of aggregates that were skipped for calculating the cumulated storage efficiency.

Examples

The following example displays information about all aggregates that are owned by nodes in the local cluster:

```
cluster::> aggr show-cumulated-efficiency
Total Data Reduction Efficiency Ratio:  5.00:1
Total Storage Efficiency Ratio:         6.97:1

cluster::> aggr show-cumulated-efficiency -details
                Total Data Reduction Ratio: 8.44:1
                Total Storage Efficiency Ratio: 6.97:1
Aggregate level Storage Efficiency
(Aggregate Deduplication and Data Compaction): 1.00:1
                Volume Deduplication Efficiency: 1.12:1
                Compression Efficiency: 5.73:1
Snapshot Volume Storage Efficiency: 1.00:1
                FlexClone Volume Storage Efficiency: 1.00:1
                Number of Offline Volumes: 0
                Number of Skipped Aggregates: 0
                Number of Efficiency Disabled Volumes: 0

cluster::> aggr show-cumulated-efficiency -aggregates aggr1
Total Data Reduction Efficiency Ratio:  6.00:1
Total Storage Efficiency Ratio:         7.41:1
saiscluster-1::*> aggr show-cumulated-efficiency -all-details
----- Total Data Reduction Efficiency -----
   Logical    Physical           Storage
   Used      Used           Efficiency Ratio
-----
   89.11MB   10.69MB           8.34:1

----- Total Storage Efficiency -----
   Logical    Physical           Storage
   Used      Used           Efficiency Ratio
-----
   89.11MB   12.91MB           6.90:1
```

```

-- Aggregate level Storage Efficiency -----
   Logical      Physical      Storage
   Used         Used         Efficiency Ratio
-----
   12.91MB     12.91MB     1.00:1

----- Volume level Storage Efficiency -----
   Logical      Physical      Total Volume Level Data
   Used         Used         Reduction Efficiency Ratio
-----
   84.74MB     5.51MB     15.39:1

---- Deduplication ---- ----- Compression ----
   Savings      Efficiency      Savings      Efficiency
   Ratio        Ratio        Ratio        Ratio
-----
   9.27MB     1.12:1     69.96MB     5.73:1

-----Snapshot-----
   Logical      Physical      Storage
   Used         Used         Efficiency Ratio
-----
   0B          2.22MB     1.00:1

-----FlexClone-----
   Logical      Physical      Storage
   Used         Used         Efficiency Ratio
-----
   0B          0B         1.00:1
Number of Offline Volumes: 0
                        Number of Skipped Aggregates: 0
                        Number of Efficiency Disabled Volumes: 0
Number of Background Deduplicaiton Disabled Volumes: 2

```

storage aggregate show-efficiency

Display aggregate storage efficiency details

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

Description

The `storage aggregate show-efficiency` command displays information about the storage efficiency of all the aggregates. The storage efficiency is displayed at four different levels:

- Total
- Aggregate

- Volume
- Snapshot and FlexClone 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.

| [-details]

Use this parameter to show additional Storage Efficiency Ratios.

| [-advanced]

Use this parameter to show additional Storage Efficiency Ratios and size values.

| [-instance] }

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

[-aggregate <aggregate name>] - Name of the Aggregate

Displays the aggregate name. If this parameter is specified, the command displays detailed information about the storage efficiency of the specified aggregate.

[-node {<nodename>|local}] - Node where Aggregate Resides

Displays the node which owns the aggregate. If this parameter is specified, the command displays storage efficiency information only about the aggregates that are located on the specified node.

[-total-logical-used {<integer>[KB|MB|GB|TB|PB]}] - Logical Size Used by Volumes, Clones, Snapshot Copies in the Aggregate

Displays the logical size used in the aggregate. This includes Volumes, Clones and Snapshots in the aggregate. The logical size is computed based on physical usage and savings obtained in the aggregate.

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

Displays the physical size used by the aggregate.

[-total-storage-efficiency-ratio <text>] - Total Storage Efficiency Ratio

Displays the total storage efficiency ratio of the aggregate.

[-total-data-reduction-logical-used-wo-snapshots {<integer>[KB|MB|GB|TB|PB]}] - Total Data Reduction Logical Used Without Snapshots

Displays the logical size used in the aggregate excluding Snapshot copies.

[-total-data-reduction-physical-used-wo-snapshots {<integer>[KB|MB|GB|TB|PB]}] - Total Data Reduction Physical Used Without Snapshots

Displays the physical size used by the aggregate excluding Snapshot copies.

[-total-data-reduction-efficiency-ratio-wo-snapshots <text>] - Total Data Reduction Efficiency Ratio Without Snapshots

Displays the total storage efficiency ratio obtained by Deduplication, Compression, Data Compaction, Pattern Detection and FlexClone data reduction technologies excluding snapshot copies on the aggregate.

`[-total-data-reduction-logical-used-wo-snapshots-flexclones {<integer>[KB|MB|GB|TB|PB] }] - Total Data Reduction Logical Used without snapshots and flexclones`

Displays the logical size used in the aggregate excluding Snapshot copies and FlexClones.

`[-total-data-reduction-physical-used-wo-snapshots-flexclones {<integer>[KB|MB|GB|TB|PB] }] - Total Data Reduction Physical Used without snapshots and flexclones`

Displays the physical size used by the aggregate excluding Snapshot copies and FlexClones.

`[-total-data-reduction-efficiency-ratio-wo-snapshots-flexclones <text>] - Total Data Reduction Efficiency Ratio without snapshots and flexclones`

Displays the total storage efficiency ratio obtained by Deduplication, Compression, Data Compaction, Pattern Detection data reduction technologies excluding snapshot copies and flexclones on the aggregate.

`[-volume-logical-used {<integer>[KB|MB|GB|TB|PB] }] - Logical Space Used for All Volumes`

Displays the logical size used by all the volumes in the aggregate.

`[-volume-physical-used {<integer>[KB|MB|GB|TB|PB] }] - Physical Space Used for All Volumes`

Displays the physical size used by all volumes in the aggregate.

`[-volume-efficiency-saved {<integer>[KB|MB|GB|TB|PB] }] - Space Saved by Volume Deduplication`

Displays the total disk space that is saved by deduplication and FlexClone for files, LUNs or NVMe namespaces by all volumes in the aggregate.

`[-volume-dedupe-zero-pattern-saved {<integer>[KB|MB|GB|TB|PB] }] - Space Saved by Volume Deduplication and pattern detection`

Displays the total disk space that is saved by deduplication, Zero pattern detection and FlexClone for files, LUNs or NVMe namespaces by all volumes in the aggregate.

`[-volume-efficiency-saved-ratio <text>] - Volume Deduplication Savings ratio`

Displays the storage efficiency ratio for savings by deduplication and FlexClone for files, LUNs or NVMe namespaces by all volumes in the aggregate.

`[-volume-compression-saved {<integer>[KB|MB|GB|TB|PB] }] - Space Saved by Volume Compression`

Displays the total disk space that is saved by compressing blocks by all volumes in the aggregate.

`[-volume-compression-saved-ratio <text>] - Volume Compression Savings ratio`

Displays the storage efficiency ratio for savings by compressing blocks on all volumes in the aggregate.

`[-volume-vbn-zero-saved {<integer>[KB|MB|GB|TB|PB] }] - Space Saved by Inline Zero Pattern Detection`

Displays the total disk space that is saved by inline zero pattern detection by all the volumes in the aggregate.

`[-volume-data-reduction-storage-efficiency-ratio <text>] - Volume Data Reduction SE Ratio`

Displays the storage efficiency ratio of all the volumes in the aggregate.

[-aggr-logical-used {<integer>[KB|MB|GB|TB|PB] }] - Logical Space Used by the Aggregate

Displays the logical size used by the aggregate.

[-aggr-physical-used {<integer>[KB|MB|GB|TB|PB] }] - Physical Space Used by the Aggregate

Displays the physical size used by the aggregate.

[-aggr-compaction-saved {<integer>[KB|MB|GB|TB|PB] }] - Space Saved by Aggregate Data Reduction

Displays the total disk space that is saved by data compaction, cross volume sharing at the aggregate level.

[-aggr-data-reduction-storage-efficiency-ratio <text>] - Aggregate Data Reduction SE Ratio

Displays the storage efficiency ratio of the aggregate.

[-snapshot-logical-used {<integer>[KB|MB|GB|TB|PB] }] - Logical Size Used by Snapshot Copies

Displays the logical size used by all Volume Snapshots residing in the aggregate.

[-snapshot-physical-used {<integer>[KB|MB|GB|TB|PB] }] - Physical Size Used by Snapshot Copies

Displays the physical size used by all Volume Snapshots residing in the aggregate.

[-snapshot-volume-data-reduction-storage-efficiency-ratio <text>] - Snapshot Volume Data Reduction Ratio

Displays the Snapshot volume storage efficiency ratio of the aggregate.

[-flexclone-volume-logical-used {<integer>[KB|MB|GB|TB|PB] }] - Logical Size Used by FlexClone Volumes

Displays the logical size used by all FlexClone volumes residing in the aggregate.

[-flexclone-volume-physical-used {<integer>[KB|MB|GB|TB|PB] }] - Physical Sized Used by FlexClone Volumes

Displays the physical size used by all FlexClone volumes in the aggregate.

[-flexclone-volume-data-reduction-storage-efficiency-ratio <text>] - FlexClone Volume Data Reduction Ratio

Displays the FlexClone volume storage efficiency ratio of the aggregate.

[-snapshot-flexclone-volume-data-reduction-storage-efficiency-ratio <text>] - Snapshot And FlexClone Volume Data Reduction SE Ratio

Displays the Snapshot and FlexClone volume storage efficiency ratio of the aggregate.

[-number-of-offline-volumes <integer>] - Number of Volumes Offline

Displays the number of volumes that are offline in the aggregate.

[-number-of-sis-disabled-volumes <integer>] - Number of SIS Disabled Volumes

Displays the number of volumes on which volume efficiency is disabled in the aggregate.

[-number-of-sis-change-log-disabled-volumes <integer>] - Number of SIS Change Log Disabled Volumes

Displays the number of volumes on which efficiency change log is disabled in the aggregate. The scheduled

background Deduplication will be disabled on these volumes.

Examples

The following example displays information about all aggregates that are owned by nodes in the local cluster:

```
cluster::*> aggr show-efficiency
```

```
Aggregate: aggr1
```

```
Node: node1
```

```
Total Data Reduction Efficiency Ratio: 3.29:1
```

```
Total Storage Efficiency Ratio: 4.29:1
```

```
Aggregate: aggr2
```

```
Node: node1
```

```
Total Data Reduction Efficiency Ratio: 4.50:1
```

```
Total Storage Efficiency Ratio: 5.49:1
```

```
cluster::*> aggr show-efficiency -details
```

```
Aggregate: aggr1
```

```
Node: node1
```

```
Total Data Reduction Ratio: 2.39:1
```

```
Total Storage Efficiency Ratio: 4.29:1
```

```
Aggregate level Storage Efficiency
```

```
(Aggregate Deduplication and Data Compaction): 1.00:1
```

```
Volume Deduplication Efficiency: 5.03:1
```

```
Compression Efficiency: 1.00:1
```

```
Snapshot Volume Storage Efficiency: 8.81:1
```

```
FlexClone Volume Storage Efficiency: 1.00:1
```

```
Number of Efficiency Disabled Volumes: 1
```

```
Aggregate: aggr2
```

```
Node: node1
```

```
Total Data Reduction Ratio: 2.39:1
```

```
Total Storage Efficiency Ratio: 4.29:1
```

```
Aggregate level Storage Efficiency
```

```
(Aggregate Deduplication and Data Compaction): 1.00:1
```

```
Volume Deduplication Efficiency: 5.03:1
```

```
Compression Efficiency: 1.00:1
```

```
Snapshot Volume Storage Efficiency: 8.81:1
```

```
FlexClone Volume Storage Efficiency: 1.00:1
```

```
Number of Efficiency Disabled Volumes: 1
```

storage aggregate show-resync-status

Display aggregate resynchronization status

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

Description

The `storage aggregate show-resync-status` command displays resync status information for each plex. 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 aggregates:

- Aggregate Name
- Resyncing Plex Name
- Resyncing 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.

[-aggregate <aggregate name>] - Aggregate

This parameter specifies the name of the aggregate.

[-plex <text>] - Plex Name

This parameter specifies the name of the plex.

[-status <text>] - Status

Displays plex status. Possible values are:

- *normal*
- *failed*
- *empty*
- *invalid*
- *uninitialized*
- *failed assimilation*
- *limbo*
- *active*
- *inactive*
- *resyncing*

These values may appear by themselves or in combination separated by commas; for example, "*normal,active*".

[-is-online {true|false}] - Is Online

Indicates whether the plex is online.

[-in-progress {true|false}] - Resync is in Progress

Indicates whether the plex is currently resyncing.

[-resyncing-percent <percent>] - Resyncing Percentage

Displays the resynchronization completion percentage if the plex is currently being resynced, '-' otherwise.

[-resync-level <integer>] - Resync Level

Displays the resync level if the plex is currently being resynced, '-' otherwise.

[-pool <integer>] - Pool

The pool number to which the majority of disks in the plex belong.

Examples

The following example displays resynchronization status for all the aggregates:

```
cluster1::> storage aggregate show-resync-status
Aggregate Resyncing Plex           Complete
-----
aggr0      plex0                       -
aggr1      plex0                       -
aggr1      plex1                      10.00
aggr2      plex0                       -
aggr2      plex2                       -
5 entries were displayed.
```

storage aggregate show-scrub-status

Display aggregate scrubbing status

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

Description

The `storage aggregate show-scrub-status` command displays the following information about the scrub status of aggregates:

- Aggregate name
- RAID groups
- Whether the scrub is suspended

- Percentage of the scrub that is completed
- Last scrub time of the aggregate

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.

[-aggregate <aggregate name>] - Aggregate

If this parameter is specified, the command displays detailed scrub-status information about the specified aggregate.

[-raidgroup <text>] - RAID Group

If this parameter is specified, the command displays information only about the aggregate that contains the specified RAID group.

[-node {<nodename>|local}] - Node

If this parameter is specified, the command displays information only about the aggregates on the specified node. The value `local` specifies the current node.

[-suspended {true|false}] - Is Suspended

If this parameter is specified, the command displays information only about the aggregates that have the specified scrub-suspension state (true or false).

[-complete-percentage <percent>] - Percentage Completed

If this parameter is specified, the command displays information only about the aggregates whose scrubs have the specified completed percentage.

[-last-scrub-time <MM/DD/YYYY HH:MM:SS>] - Last Scrub Time

If this parameter is specified, the command displays information only about the aggregates that have the specified last-scrub time, in the format `MM/DD/YYYY HH:MM:SS`.

Examples

The following example displays scrub-status information for all the aggregates:

```

cluster1::> storage aggregate show-scrub-status
Aggregate RAID Groups           Suspended Percentage Last Scrub Time
-----
aggr0    /aggr0/plex0/rg0    true           0% 3/31/2011  21:23:02
aggr1    /aggr1/plex0/rg1    true           45% 3/30/2011  01:05:00
aggr2    /aggr2/plex0/rg0    true           33% 3/30/2011  23:43:34
aggr3    /aggr3/plex0/rg1    true           79% 3/29/2011  00:34:36
4 entries were displayed.

```

The following example displays detailed information about the aggregate named aggr1:

```

cluster1::> storage aggregate show-scrub-status -instance -aggregate aggr1
Aggregate: aggr1
RAID Group: /aggr1/plex0/rg0
Is Suspended: false
Percentage Completed: 2%
Last Scrub Time: 3/31/2011 22:02:50

```

storage aggregate show-space

Display details of space utilization within an aggregate.

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

Description

The `storage aggregate show-space` command displays information about space utilization within aggregates and any attached external capacity tier. The command output breaks down space usage in the specified aggregate by feature. If no parameters are specified, the command displays this information about all aggregates. Note that used percentage for an external capacity tier will be non-zero only if a size limit was set for that aggregate's attached tier.

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.

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

If this parameter is specified, the command displays information only about the specified aggregates.

[-bin-num <integer>] - Bin Number

If this parameter is specified, the command displays information only about the aggregates whose bin number for the storage tier matches the specified value. Typically, bin 0 refers to the performance tier or active file system and bin numbers greater than 0 refer to the external capacity tiers attached to the aggregate.

[-tier-name <text>] - Tier Name For Show Command

If this parameter is specified, the command displays information only about the aggregates whose attached storage tier name matches the specified value.

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

If this parameter is specified, the command displays information only about space used in the specified aggregate or aggregates.

[-aggregate-uuid <UUID>] - Uuid of the Aggregate

If this parameter is specified, the command displays information only about the aggregates whose UUID matches the specified value.

[-volume-footprints {<integer>[KB|MB|GB|TB|PB]}] - Volume Footprints

If this parameter is specified, the command displays information only about the aggregate or aggregates that have the specified amount of space in use by volume footprints. A volume's footprint is the overall amount of space that a volume occupies in the aggregate, including the volume metadata and data.

[-volume-footprints-percent <percent_no_limit>] - Volume Footprints Percent

If this parameter is specified, the command displays information only about the aggregate or aggregates whose volume footprints occupy the specified percentage of space.

[-snap-size-total {<integer>[KB|MB|GB|TB|PB]}] - Total Space for Snapshot Copies in Bytes

If this parameter is specified, the command displays information only about the aggregate or aggregates that have the specified amount of space in use by aggregate Snapshot copies. This field includes the space that is reserved for Snapshot copies and is not available to volumes or aggregate data and metadata. It is set to 0 by default.

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

If this parameter is specified, the command displays information only about the aggregate or aggregates that have the specified percentage of space in use by aggregate Snapshot copies.

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

If this parameter is specified, the command displays information only about the aggregate or aggregates that have the specified amount of space in use by aggregate metadata.

[-aggregate-metadata-percent <percent_no_limit>] - Aggregate Metadata Percent

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

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

If this parameter is specified, the command displays information only about the aggregate or aggregates that have the specified amount of space in use in the aggregate.

It is important to note that this parameter treats the entire Snapshot reserve as used space since it is not available for volumes.

[-used-including-snapshot-reserve-percent <percent_no_limit>] - Total Used Percent

If this parameter is specified, the command displays information only about the aggregate or aggregates that have the specified percentage of space in use in the aggregate and its Snapshot reserve.

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

If this parameter is specified, the command displays information only about the aggregate or aggregates that have the specified size.

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

If this parameter is specified, the command displays information only about the aggregate or aggregates 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 aggregate will make this space available.

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

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

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

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

This differs from *total-used* space by the space that is guaranteed for future writes. 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 aggregate or aggregates that have the specified percentage of physical space in use in the aggregates.

[-performance-tier-inactive-user-data {<integer>[KB|MB|GB|TB|PB] }] - Performance Tier Inactive User Data

If this parameter is specified, the command displays information only about the aggregates whose amount of inactive user data in the performance tier matches the specified value. The inactive user data can be tiered out to a capacity tier if the aggregate is a FabricPool.

[-performance-tier-inactive-user-data-percent <percent>] - Performance Tier Inactive User Data Percent

If this parameter is specified, the command displays information only about the aggregates whose percentage of inactive user data in the performance tier matches the specified value.

[-cross-volume-dedupe-metadata {<integer>[KB|MB|GB|TB|PB] }] - Aggregate Dedupe Metadata

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

[-cross-volume-dedupe-metadata-percent <percent_no_limit>] - Aggregate Dedupe Metadata Percent

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

[`-cross-volume-dedupe-temp-metadata` {<integer>[KB|MB|GB|TB|PB]}] - Aggregate Dedupe Temporary Metadata

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

[`-cross-volume-dedupe-temp-metadata-percent` <percent_no_limit>] - Aggregate Dedupe Temporary Metadata Percent

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

[`-total-provisioned-space` {<integer>[KB|MB|GB|TB|PB]}] - Total Space Provisioned inside Aggregate

If this parameter is specified, the command displays information about the total provisioned space across all the volumes in the aggr.

[`-total-provisioned-space-percent` <percent_no_limit>] - Percentage Space Provisioned inside Aggregate

If this parameter is specified, the command displays information about the percentage of total provisioned space across all the volumes against aggr's size.

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

If this parameter is specified, the command displays information only about the aggregates whose physical space use in the attached object store matches the specified value.

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

If this parameter is specified, the command displays information only about aggregates whose physical space in use in the attached object store as a percentage of the license limit matches the specified value.

[`-object-store-referenced-capacity` {<integer>[KB|MB|GB|TB|PB]}] - Total Object Store Logical Referenced Capacity

If this parameter is specified, the command displays information only about the aggregates whose logical reference capacity space in use in the attached object store matches the specified value.

[`-object-store-referenced-capacity-percent` <percent_no_limit>] - Object Store Logical Referenced Capacity Percentage

If this parameter is specified, the command displays information only about aggregates whose logical reference capacity space in use in the attached object store as a percentage of the license limit matches the specified value.

[`-object-store-metadata` {<integer>[KB|MB|GB|TB|PB]}] - (DEPRECATED)-Object Store Metadata

This parameter is deprecated in Data ONTAP 9.10.1 and later. If this parameter is specified, the command displays information only about the aggregates whose metadata space in use in the attached object store matches the specified value.

[`-object-store-metadata-percent` <percent_no_limit>] - (DEPRECATED)-Object Store Metadata Percent

This parameter is deprecated in Data ONTAP 9.10.1 and later. If this parameter is specified, the command displays information only about aggregates whose metadata space in use in the attached object store as a percentage of the license limit matches the specified value.

`[-object-store-unreclaimed-space {<integer>[KB|MB|GB|TB|PB] }] - (DEPRECATED)-Total Unreclaimed Space`

This parameter is deprecated in Data ONTAP 9.10.1 and later. If this parameter is specified, the command displays information only about the aggregates whose unreclaimed space in use in the attached object store matches the specified value.

`[-object-store-unreclaimed-space-percent <percent_no_limit>] - (DEPRECATED)-Object Store Unreclaimed Space Percentage`

This parameter is deprecated in Data ONTAP 9.10.1 and later. If this parameter is specified, the command displays information only about aggregates whose unreclaimed space in use in the attached object store as a percentage of the license limit matches the specified value.

`[-object-store-size {<integer>[KB|MB|GB|TB|PB] }] - Object Store Size`

If this parameter is specified, the command displays information only about the aggregates whose attached object store size limit matches the specified value.

`[-object-store-sis-space-saved {<integer>[KB|MB|GB|TB|PB] }] - Object Store Space Saved by Storage Efficiency`

If this parameter is specified, the command displays information only about the aggregates whose amount of space saved by storage efficiency matches the specified value.

`[-object-store-sis-space-saved-percent <percent_no_limit>] - Object Store Space Saved by Storage Efficiency Percentage`

If this parameter is specified, the command displays information only about the aggregates whose percentage of space saved by storage efficiency matches the specified value.

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

If this parameter is specified, the command displays information only about the aggregates whose logical space use in the attached object store matches the specified value.

`[-object-store-logical-used-percent <percent_no_limit>] - Logical Used Percentage`

If this parameter is specified, the command displays information only about aggregates whose logical space in use in the attached object store as a percentage of the license limit matches the specified value.

`[-object-store-logical-unreferenced-capacity {<integer>[KB|MB|GB|TB|PB] }] - Logical Unreferenced Capacity`

If this parameter is specified, the command displays information only about the aggregates whose logical unreferenced capacity in use in the attached object store matches the specified value.

`[-object-store-logical-unreferenced-capacity-percent <percent_no_limit>] - Logical Unreferenced Percentage`

This parameter is deprecated in Data ONTAP 9.10.1 and later. If this parameter is specified, the command displays information only about aggregates whose logical unreferenced capacity in use in the attached object store as a percentage of the license limit matches the specified value.

Examples

The following example displays information about all aggregates:

```

cluster1::> storage aggregate show-space
Aggregate : aggr0
Feature                               Used      Used%
-----
Volume Footprints                     5.75GB    91%
Aggregate Metadata                     380KB     0%
Snapshot Reserve                       325.3MB   5%
Total Used                             6.07GB    96%
Total Physical Used                     221.9MB   3%
Aggregate : aggr1
Feature                               Used      Used%
-----
Volume Footprints                     2.03GB    33%
Aggregate Metadata                     304KB     0%
Total Used                             2.03GB    33%
Total Physical Used                     2.23MB    0%

2 entries were displayed.

```

The following example displays information about all the aggregates in a system including the ones that have an object store attached to them.

```

cluster-1::> storage aggregate show-space
Aggregate : aggr0
Feature                               Used          Used%
-----
Volume Footprints                     2.87GB        90%
Aggregate Metadata                     328KB         0%
Snapshot Reserve                       162.6MB       5%
Total Used                             3.03GB        95%
Total Physical Used                    2.08GB        65%
Aggregate : aggr1
Performance Tier
Feature                               Used          Used%
-----
Volume Footprints                     1.25GB        13%
Aggregate Metadata                     540KB         0%
Snapshot Reserve                       0B            0%
Total Used                             1.25GB        13%
Total Physical Used                    1.23GB        13%
Aggregate : aggr1
Object Store: my-store
Feature                               Used          Used%
-----
Referenced Capacity                    811.2MB       0%
Metadata                               0B            0%
Unreclaimed Space                      0B            0%
Space Saved by Storage Efficiency       0B            0%
Total Physical Used                    811.2MB       0%

```

storage aggregate show-spare-disks

Display spare disks

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

Description

The command `storage aggregate show-spare-disks` displays information about spare disks. The command output depends on the parameter or parameters specified with the command. If no parameters are specified, the command displays information about all spare disks 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.

| [-partition-info]

Displays the following information about root-data and root-data1-data2 partitioned spares.

- Disk
- Type
- Class
- RPM
- Checksum
- Local Data Usable
- Local Data1 Usable
- Local Data2 Usable
- Local Root Usable
- Physical Size
- Status

| [-instance] }

If this parameter is specified, the command displays detailed information about each spare disk.

[-original-owner <text>] - Original Owner

Selects the spare disks that match this parameter value.

[-disk <disk path name>] - Disk Name

Selects the spare disks that match this parameter value.

[-checksum-style {advanced_zoned | block | none}] - Checksum Style

Selects the spare disks that match this parameter value. Possible values are:

- block — Supports block checksum
- advanced_zoned — Supports advanced zone checksum
- none — No checksum support

[-disk-type {ATA | BSAS | FCAL | FSAS | LUN | MSATA | SAS | SSD | VMDISK | SSD-NVM | SSD-CAP | SSD-ZNS | VMLUN | VMLUN-SSD}] - Disk Type

Selects the spare disks that match this parameter value.

[-effective-disk-type {ATA | BSAS | FCAL | FSAS | LUN | MSATA | SAS | SSD | VMDISK | SSD-NVM | SSD-CAP | SSD-ZNS | VMLUN | VMLUN-SSD}] - Effective Disk Type

Selects the spare disks that match this parameter value.

Hard disk drives with the same `effective-disk-type` value may be mixed together in the same aggregate depending upon the system's `raid.mix.hdd.disktype.capacity` and `raid.mix.hdd.disktype.performance` option settings. Solid state drives with the same `effective-disk-type` value may be mixed together in the same aggregate depending upon the system's `raid.mix.disktype.solid_state` option setting.

[`-standard-disk-type` {SATA | FC | NL-SAS | LUN | SAS | SCSI | SSD | VM-DISK | NVMe-SSD | SSD-CAP | SSD-ZNS | VM-LUN | VM-LUN-SSDS}] - Standard Disk Type

Selects the spare disks that match this parameter value.

[`-disk-class` {capacity | performance | archive | solid-state | array | virtual}] - Disk Class

Selects the spare disks that match this parameter value. Possible values are:

- `capacity` — Capacity-oriented, near-line disk types. Includes disk types FSAS, BSAS and ATA.
- `performance` — Performance-oriented, enterprise class disk types. Includes disk types FCAL and SAS.
- `archive` — Archive class SATA disks in multi-disk carrier storage shelves. Includes disk type MSATA.
- `solid-state` — Solid-state drives. Includes disk type SSD, SSD-CAP and SSD-NVM.
- `array` — Logical storage devices backed by storage arrays and used by Data ONTAP as disks. Includes disk type LUN.
- `virtual` — Virtual disks that are formatted and managed by the hypervisor. Includes disk type VMDISK.

Disks with the same `disk-class` value are compatible for use in the same aggregate.

[`-disk-rpm` <integer>] - Disk RPM

Selects the spare disks that match this parameter value.

[`-effective-disk-rpm` <integer>] - Effective Disk RPM

Selects the spare disks that match this parameter value.

Hard disk drives with the same `effective-disk-rpm` value may be mixed together in the same aggregate depending upon the system's `raid.mix.hdd.rpm.capacity` and `raid.mix.hdd.rpm.performance` option settings.

[`-syncmirror-pool` <text>] - Pool Number

Selects the spare disks that match this parameter value.

[`-owner-name` {<nodename>|local}] - Current Owner

Selects the spare disks that match this parameter value.

[`-home-owner-name` {<nodename>|local}] - Home Owner

Selects the spare disks that match this parameter value.

[`-dr-owner-name` {<nodename>|local}] - DR Home Owner

Selects the spare disks that match this parameter value.

[`-usable-size-blks` <integer>] - Disk Usable Size in 4K blocks

Selects the spare disks that match this parameter value.

[`-local-usable-data-size-blks` <integer>] - Local Node Data Usable Size in 4K blocks

Selects the spare disks that match this parameter value.

Disks that have two partitions can be used for one root aggregate and one data aggregate.

Disks that have three partitions can be used for one root aggregate and one or two data aggregates.

This value describes the data partition size (of root-data partitioned disk) or the combined data1 + data2 partition size (of root-data1-data2 partitioned disk) in 4KB blocks.

[-local-usable-root-size-blks <integer>] - Local Node Root Usable Size in 4K blocks

Selects the spare disks that match this parameter value.

Disks that have two partitions can be used for one root aggregate and one data aggregate.

Disks that have three partitions can be used for one root aggregate and one or two data aggregates.

This value describes the root partition size in 4KB blocks.

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

Selects the spare disks that match this parameter value.

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

Selects the spare disks that match this parameter value.

[-local-usable-data-size {<integer>[KB|MB|GB|TB|PB]}] - Local Node Data Usable Size

Selects the spare disks that match this parameter value.

Disks that have two partitions can be used for one root aggregate and one data aggregate.

Disks that have three partitions can be used for one root aggregate and one or two data aggregates.

This value describes the data partition size (of root-data partitioned disk) or the combined data1 + data2 partition size (of root-data1-data2 partitioned disk) in auto-scaled units.

[-local-usable-root-size {<integer>[KB|MB|GB|TB|PB]}] - Local Node Root Usable Size

Selects the spare disks that match this parameter value.

Disks that have two partitions can be used for one root aggregate and one data aggregate.

Disks that have three partitions can be used for one root aggregate and one or two data aggregates.

This value describes the root partition size in auto-scaled units.

[-is-disk-zeroed {true|false}] - Is Disk Zeroed?

Selects the spare disks that match this parameter value.

When disks are zeroed, they can be provisioned directly into aggregates which avoids a lengthy zeroing process.

[-is-disk-zeroing {true|false}] - Is Disk Zeroing?

Selects the spare disks that match this parameter value.

[-zeroing-percent <percent>] - Zeroing Percentage Completed

Selects the spare disks that match this parameter value.

[-is-sparecore {true|false}] - Sparecore Disk?

Selects the spare disks that match this parameter value.

[-sparecore-status <Spare core status>] - Sparecore Status

Selects the spare disks that match this parameter value.

[-sparecore-percent <percent>] - Sparecore Percentage Completed

Selects the spare disks that match this parameter value.

[-is-disk-shared {true|false}] - Is Disk Shared?

Selects the spare disks that match this parameter value.

Shared disks have partitions that allow them to be used in multiple aggregates and between nodes in an HA pair. When set to *true*, this parameter selects shared disks in which the root partition and/or the data partition is a spare. When set to *false* only spare disks without partitions are displayed. When this parameter is not used, all spare disks are displayed.

[-is-disk-offline {true|false}] - Is Disk Offline?

Selects the spare disks that match this parameter value.

Disk offline events are typically temporary events which allow Data ONTAP to perform background error recovery activity.

[-is-disk-sick {true|false}] - Is Disk Sick?

Selects the spare disks that match this parameter value.

A sick disk triggers Rapid RAID Recovery to copy data to a spare drive. At the end of the process the sick disk is marked as *broken*.

[-is-disk-left-behind {true|false}] - Is Disk Left Behind Spare?

Selects the spare disks that match this parameter value.

Disks are left behind if they are not responding during a giveback or switchback event.

[-local-usable-data1-size-blks <integer>] - Local Node Data1 Usable Size in 4K blocks

Selects the spare disks that match this parameter value.

Disks that have two partitions can be used for one root aggregate and one data aggregate.

Disks that have three partitions can be used for one root aggregate and one or two data aggregates.

This value describes the data1 partition size of a root-data1-data2 partitioned disk in 4KB blocks.

[-local-usable-data2-size-blks <integer>] - Local Node Data2 Usable Size in 4K blocks

Selects the spare disks that match this parameter value.

Disks that have two partitions can be used for one root aggregate and one data aggregate.

Disks that have three partitions can be used for one root aggregate and one or two data aggregates.

This value describes the data2 partition size of a root-data1-data2 partitioned disk in 4KB blocks.

[-local-usable-data1-size {<integer>[KB|MB|GB|TB|PB]}] - Local Node Data1 Usable Size

Selects the spare disks that match this parameter value.

Check on the progress of a previous disk zeroing command.

```
cluster1::> storage aggregate show-spare-disks -owner-name node-b -zeroing
-percent >0
```

```
Original Owner: node-b
```

```
Pool0
```

```
Spare Pool
```

```
Usable Physical
```

```
Disk          Type  Class          RPM  Checksum          Size
Size Status
```

```
-----
```

```
1.1.13          BSAS  capacity      7200  block             827.7GB
```

```
828.0GB zeroing, 17% done
```

```
1.1.15          BSAS  capacity      7200  block             413.2GB
```

```
414.0GB zeroing, 28% done
```

```
2 entries were displayed.
```

storage aggregate show-status

Display aggregate configuration

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

Description

The `storage aggregate show-status` command displays the RAID layout and disk configuration of aggregates. The command output depends on the parameter or parameters specified with the command. If no parameters are specified, the command displays information about all aggregates in the cluster.



This command does not use pagination. You can reduce the output by filtering with the parameters below.

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] }

This parameter currently has no effect.

[-aggregate <text>] - Aggregate Name

Selects the aggregates that match this parameter value.

[-node <nodename>] - Node

Selects the aggregates that match this parameter value.

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

Selects the aggregates that match this parameter value.

Examples

Display the RAID layout of a Flash Pool aggregate.

```
cluster1::> storage aggregate show-status -aggregate nodeB_flashpool_1
```

Owner Node: node-b

Aggregate: nodeB_flashpool_1 (online, raid_dp, hybrid) (block checksums)

Plex: /nodeB_flashpool_1/plex0 (online, normal, active, pool0)

RAID Group /nodeB_flashpool_1/plex0/rg0 (normal, block checksums)

Usable

Physical

Position	Disk	Pool	Type	RPM	Size
----------	------	------	------	-----	------

Size Status

-----	-----	-----	-----	-----	-----
dparity	1.1.7	0	BSAS	7200	827.7GB
828.0GB	(normal)				
parity	1.1.8	0	BSAS	7200	827.7GB
828.0GB	(normal)				
data	1.1.10	0	BSAS	7200	827.7GB
828.0GB	(normal)				
data	1.1.11	0	BSAS	7200	827.7GB
828.0GB	(normal)				
data	1.1.12	0	BSAS	7200	827.7GB
828.0GB	(normal)				

RAID Group /nodeB_flashpool_1/plex0/rg1 (normal, block checksums) (Storage Pool: SP2)

Usable

Physical

Position	Disk	Pool	Type	RPM	Size
----------	------	------	------	-----	------

Size Status

-----	-----	-----	-----	-----	-----
shared	1.0.22	0	SSD	-	186.2GB
745.2GB	(normal)				
shared	1.0.20	0	SSD	-	186.2GB
745.2GB	(normal)				
shared	1.0.18	0	SSD	-	186.2GB
745.2GB	(normal)				
shared	1.0.16	0	SSD	-	186.2GB
745.2GB	(normal)				

storage aggregate show

Display a list of aggregates

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

Description

The `storage aggregate show` command displays information about aggregates. 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 aggregates:

- Aggregate name
- Size
- Available size
- Percentage used
- State
- Number of volumes
- Node on which the aggregate is located
- RAID status

To display detailed information about a single aggregate, use the `-aggregate` 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.

| [-checksum]

If this parameter is specified, the command displays information about the checksum for all aggregates in the cluster:

- Aggregate name
- Checksum status (active, off, reverting, none, unknown, initializing, reinitializing, reinitialized, upgrading_phase1, upgrading_phase2)
- Checksum style (none, advanced_zoned, block, mixed, WAFL, or unknown)

| [-disk]

If this parameter is specified, the command displays disk names for all aggregates in the cluster:

- Aggregate name
- Number and names of disks in the aggregate

| [-raid-info]

If this parameter is specified, the command displays information about RAID groups, RAID type, maximum RAID size, checksum state, checksum style and whether the RAID status is inconsistent.

| [-instance] }

If this parameter is specified, the command displays detailed information about all aggregates in the cluster.

[-aggregate <aggregate name>] -Aggregate

If this parameter is specified, the command displays detailed information about the specified aggregate.

[-storage-type {hdd | hybrid | lun | ssd | vmdisk}] - Storage Type

If this parameter is specified, the command displays information only about the aggregates with the specified storage type. The possible values are hdd, hybrid, lun, ssd and vmdisk.

[-chksumstyle <aggrChecksumStyle>] - Checksum Style

If this parameter is specified, the command displays information only about the aggregates that use the specified checksum style.

[-diskcount <integer>] - Number Of Disks

If this parameter is specified, the command displays information only about the aggregates that have the specified number of disks.

[-m, -mirror <true>] - Is Mirrored

If this parameter is specified, the command displays information only about the aggregates that have the specified mirrored value.

[-d, -disklist <disk path name>,...] - Disks for First Plex

If this parameter is specified, the command displays information only about the aggregates that have the specified disk or disks.

[-mirror-disklist <disk path name>,...] - Disks for Mirrored Plex

If this parameter is specified, the command displays information only about the aggregates that have the specified disk or disks present in the mirrored plex.

[-node {<nodename>|local}] - Node

If this parameter is specified, the command displays information only about the aggregates that are located on the specified node.

[-free-space-realloc {on|off}] - Free Space Reallocation

If this parameter is specified, the command displays whether free space reallocation is enabled on the specified aggregate.

[-ha-policy {sfo|cfo}] - HA Policy

This optionally specifies the high-availability policy to be used in the context of a root recovery procedure. Do not modify this setting unless directed to do so by a customer support representative.

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

If this parameter is specified, the command displays information only about the aggregates that have the specified space reserved for Snapshot copies.

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

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

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

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

[-hybrid-enabled {true|false}] - Hybrid Enabled

If this parameter is specified, the command displays information only about the aggregates that are eligible

to contain both SSD and non-SSD RAID groups.

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

If this parameter is specified, the command displays information only about the aggregates that have the specified available size.

[`-checksumenabled` {true|false}] - Checksum Enabled

If this parameter is specified, the command displays information only about the aggregates that have the specified checksum setting.

[`-checksumstatus` <text>] - Checksum Status

If this parameter is specified, the command displays information only about the aggregates that have the specified checksum status. The possible values for checksum status include the following: active, off, reverting, none, unknown, initializing, reinitializing, reinitialized, upgrading_phase1, and upgrading_phase2.

[`-cluster` <text>] - Cluster

If this parameter is specified, the command displays information only about the aggregates that are owned by nodes in the specified cluster. By default, only local cluster aggregates are displayed.

[`-cluster-id` <UUID>] - Home Cluster ID

If this parameter is specified, the command displays information only about the aggregates that are owned by nodes in the cluster specified by the cluster UUID. By default, only local cluster aggregates are displayed.

[`-dr-home-id` <integer>] - DR Home ID

If this parameter is specified, the command displays information only about the aggregates whose Disaster Recovery home node has the specified system ID.

[`-dr-home-name` <text>] - DR Home Name

If this parameter is specified, the command displays information only about the aggregates whose Disaster Recovery home is the specified node.

[`-inofile-version` <integer>] - Inofile Version

If this parameter is specified, the command displays information only about the aggregates whose inode files are at the specified version.

[`-has-mroot` {true|false}] - Has Mroot Volume

If this parameter is specified, the command displays information about only the aggregates that contain their owning node's management root directory.

[`-has-partner-mroot` {true|false}] - Has Partner Node Mroot Volume

If this parameter is specified, the command displays information about only the aggregates that contain the management root directory of their owning node's failover partner.

[`-home-id` <integer>] - Home ID

If this parameter is specified, the command displays information only about the aggregates whose home node has the specified system ID.

[`-home-name` <text>] - Home Name

If this parameter is specified, the command displays information only about the aggregates whose home node is the specified node.

[-hybrid-cache-size-total <integer>[KB|MB|GB|TB|PB]] - Total Hybrid Cache Size

If this parameter is specified, the command displays information only about the aggregates that have the specified total cache size in a Flash Pool.

[-hybrid {true|false}] - Hybrid

If this parameter is specified, the command displays information only about the aggregates that currently contain both SSD and non-SSD RAID groups.

[-inconsistent {true|false}] - Inconsistent

If this parameter is specified, the command displays information only about the aggregates that have the specified consistency.

[-is-home {true|false}] - Is Aggregate Home

If this parameter is specified, the command displays information only about the aggregates whose home node and owner node have the same system ID.

[-s, -maxraidsize <integer>] - Max RAID Size

If this parameter is specified, the command displays information only about the aggregates that have the specified maximum number of disks for RAID groups.



For Flash Pools, this option controls the maximum size of the HDD RAID groups.

[-cache-raid-group-size <integer>] - Flash Pool SSD Tier Maximum RAID Group Size

If this parameter is specified, the command displays information about the maximum RAID group size for the SSD tier for Flash Pools.



This parameter is applicable only for Flash Pools.

[-owner-id <integer>] - Owner ID

If this parameter is specified, the command displays information only about the aggregates that are owned by the node with the specified system ID.

[-owner-name <text>] - Owner Name

If this parameter is specified, the command displays information only about the aggregates that are owned by the specified node.

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

If this parameter is specified, the command displays information only about the aggregates that have the specified used size, as a percentage.

[-plexes <text>,...] - Plexes

If this parameter is specified, the command displays information only about the aggregates that have the specified plex or plexes.

[-raidgroups <text>,...] - RAID Groups

If this parameter is specified, the command displays information only about the aggregates that have the specified RAID group or groups.

[-raidstatus <text>] - RAID Status

If this parameter is specified, the command displays information only about the aggregates that have the specified RAID status. The possible values for RAID status are normal, copying, ironing, degraded, mirror degraded, growing, initializing, invalid, needs check, partial, reconstruct, raid4, raid_ep, raid0, raid_dp, raid_tec, redirect, and waf1 inconsistent. You can specify multiple values (for example, reconstruct and growing).

[-t, -raidtype {raid_tec|raid_dp|raid4|raid_ep}] - RAID Type

If this parameter is specified, the command displays information only about the aggregates that use the specified RAID type. The possible values are *raid0* for RAID 0, *raid4* for RAID4 of disks other than *SSD-ZNS*, *raid_ep* for RAID4 of *SSD-ZNS* disks with additional RAID4 protection inside each disk, *raid_dp* for RAID-DP, *raid_tec* for RAID-TEC, and *mixed_raid_type* for aggregates that include a mix of RAID types.

[-resyncsnaptime <integer>] - SyncMirror Resync Snapshot Frequency in Minutes

If this parameter is specified, the command displays information only about the aggregates whose SyncMirror Resynchronization Snapshot Frequency is the specified value.

[-root {true|false}] - Is Root

If this parameter is specified, the command displays information about only the root aggregates in the cluster.

[-sis-metadata-space-used {<integer>[KB|MB|GB|TB|PB]}] - Space Used by Metadata for Volume Efficiency

If this parameter is specified, the command displays information about only the aggregates with the specified space used by A-SIS metafiles for volume efficiency. This parameter is deprecated in Data ONTAP 8.2 and later. Use the volume-footprint-list-info API for details related to space usage by deduplication metadata

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

If this parameter is specified, the command displays information only about the aggregates that have the specified size. The size of the aggregate is reported as the size available for use by WAFL, excluding WAFL reserve and aggregate Snapshot reserve capacity. Use the [storage aggregate show-space](#) command to see the details of space utilization within an aggregate.

[-state <aggregate state>] - State

If this parameter is specified, the command displays information only about the aggregates that have the specified state.

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

If this parameter is specified, the command displays information only about the aggregates that have the specified used size.

[-uses-shared-disks {true|false}] - Uses Shared Disks

Selects the aggregates that match this parameter value. This parameter is used to list all the aggregates that use shared HDDs or shared SSDs.

[-uuid <text>] - UUID String

If this parameter is specified, the command displays information only about the aggregate that has the specified UUID. This parameter is available only at the advanced privilege level and higher.

[-volcount <integer>] - Number Of Volumes

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

[-is-autobalance-eligible {true|false}] - Is Eligible for Auto Balance Aggregate

If this parameter is specified, the command displays information only about the aggregates that are considered by the Auto Balance Aggregate feature.

[-autobalance-state <Auto Balance Aggregate state>] - State of the aggregate being balanced

If this parameter is specified, the command displays information only about the aggregates that have the specified state.

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

If this parameter is specified, the command displays information only about the aggregates that have the specified physical used size. This differs from *total-used* space by the space that is guaranteed for future writes. 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 aggregates that have the specified physical used percent.

[-autobalance-state-change-counter <integer>] - State Change Counter for Auto Balancer

If this parameter is specified, the command displays information only about the aggregates that have the specified number of state change caused by the Auto Balance Aggregate feature.

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

If this parameter is specified, the command displays information only about the aggregates that have the specified snaplock-type.

[-is-nve-capable {true|false}] - Is NVE Capable

This parameter indicates whether or not the aggregate is capable of supporting NVE (NetApp volume encryption).

[-is-cft-precommit {true|false}] - Is in the precommit phase of Copy-Free Transition

Selects the aggregates that are set with this parameter value. This parameter lists all the aggregates that are in the precommit phase of a Copy-Free Transition workflow.

[-is-transition-out-of-space {true|false}] - Is a 7-Mode transitioning aggregate that is not yet committed in clustered Data ONTAP and is currently out of space

Selects the aggregates that match this parameter value. This parameter is used to list all the 7-mode transitioning aggregates that are not yet committed in clustered Data ONTAP, and are currently out of space.

[-autobalance-unbalanced-threshold-percent <integer>] - Threshold When Aggregate Is Considered Unbalanced (%)

If this parameter is specified, the command displays information only about the aggregates that have the specified unbalanced threshold percentage.

[`-autobalance-available-threshold-percent <integer>`] - Threshold When Aggregate Is Considered Balanced (%)

If this parameter is specified, the command displays information only about the aggregates that have the specified available threshold percentage.

[`-resync-priority {high(fixed)|high|medium|low}`] - Resynchronization Priority

This parameter indicates the relative priority that is used to decide whether a mirrored aggregate can start a resynchronization operation or not. This field is not set for unmirrored aggregates.

Use the [storage aggregate resynchronization modify](#) command to modify this field for mirrored aggregates.

The valid values for this field are:

- `high(fixed)`: This value is reserved for Data ONTAP system aggregates, which cannot have any other value for this field. It cannot be explicitly set on a data aggregate. These aggregates always start their resynchronization operation at the first available opportunity.
- `high`: Mirrored data aggregates with this priority value start resynchronization first.
- `medium`: Mirrored data aggregates with this priority value start resynchronization after all the system aggregates and data aggregates with 'high' priority value have started resynchronization.
- `low`: Mirrored data aggregates with this priority value start resynchronization only after all the other aggregates have started resynchronization.

[`-data-compaction-space-saved {<integer>[KB|MB|GB|TB|PB]}`] - Space Saved by Data Compaction

This parameter indicates the amount of the space saved by Data Compaction in bytes.

[`-data-compaction-space-saved-percent <percent>`] - Percentage Saved by Data Compaction

This parameter indicates the percentage of space saved in the aggregate by Data Compaction.

[`-data-compacted-count {<integer>[KB|MB|GB|TB|PB]}`] - Amount of compacted data

This parameter indicates the number of bytes occupied by compacted data inside this aggregate.

[`-creation-timestamp <MM/DD/YYYY HH:MM:SS>`] - Timestamp of Aggregate Creation

This parameter indicates the date and time the aggregate was created.

[`-single-instance-data-logging {off|on}`] - Enable SIDL

If this parameter is specified, the command displays whether Single Instance Data Logging feature is enabled on the specified aggregate.

[`-composite {true|false}`] - Composite

If this parameter is specified, the command displays information only about aggregates whose classification as a FabricPool matches the specified value. A FabricPool has an external capacity tier attached to it.

[`-is-fabricpool-mirrored {true|false}`] - Is FabricPool Mirrored

If this parameter is specified, the command displays information only about FabricPools whose classification as mirrored matches the specified value. A mirrored FabricPool has a second external capacity tier attached to it.

[`-composite-capacity-tier-used` {<integer>[KB|MB|GB|TB|PB]}] - Capacity Tier Used Size

If this parameter is specified, the command displays the amount of space in use in the attached external capacity tier.

[`-sis-space-saved` {<integer>[KB|MB|GB|TB|PB]}] - Space Saved by Storage Efficiency

This parameter indicates the total amount of space saved by storage efficiency in bytes.

[`-sis-space-saved-percent` <percent>] - Percentage of Space Saved by Storage Efficiency

This parameter indicates the percentage of space saved by storage efficiency.

[`-sis-shared-count` {<integer>[KB|MB|GB|TB|PB]}] - Amount of Shared bytes count by Storage Efficiency

This parameter indicates the number of bytes shared by storage efficiency.

[`-is-inactive-data-reporting-enabled` {true|false}] - Inactive Data Reporting Enabled

If this parameter is specified, the command displays whether reporting of inactive user data is enabled. This parameter is not allowed on FabricPools

[`-inactive-data-reporting-start-timestamp` <MM/DD/YYYY HH:MM:SS>] - Timestamp when Inactive Data Reporting was Enabled

If this parameter is specified, the command displays the timestamp at which inactive data reporting was enabled on the aggregate. This parameter is not allowed on FabricPools.

[`-encrypt-with-aggr-key` {true|false}] - Enable Aggregate level Encryption

Selects the aggregates that are encrypted with aggregate keys.

[`-drive-protection-enabled` {true|false}] - Aggregate uses data protected SEDs

If this parameter is specified, the command displays whether this aggregate is entirely composed of self-encrypting drives that have data protection enabled.

[`-azcs-read-optimization` {on|off}] - azcs read optimization

If this parameter is specified, the command displays whether azcs-with-compression feature is enabled.

[`-space-required-for-revert` {<integer>[KB|MB|GB|TB|PB]}] - Metadata Reserve Space Required For Revert

If this parameter is specified, the command displays the additional amount of free space needed for metadata reserve in the aggregate if the system is required to revert to an ONTAP release earlier than 9.11.1x.

Examples

The following example displays information about all aggregates that are owned by nodes in the local cluster:

```

cluster1::> storage aggregate show
Aggregate      Size Available Used% State   #Vols Nodes           RAID
Status
-----
-----
aggr0          6.21TB      1.78TB   71% online    49 cluster1-01
raid_dp,
normal
aggr1          56.04MB     55.89MB    0% online     0 cluster1-02
raid_dp,
mirrored,
normal
aggr2          1.77TB      1.63TB    8% online     1 cluster1-01
raid_dp,
normal
aggr3          1.77TB      1.73TB    2% online     2 cluster1-02
raid_dp,
normal
4 entries were displayed.

```

The following example displays information about an aggregate name aggr1:

```

cluster1::> storage aggregate show -aggregate aggr1
Aggregate: aggr1
Checksum Style: block
Number Of Disks: 6
Mirror: true
Nodes: cluster1-02
Disks for First Plex: 1.1.2,
1.1.10,
1.1.11
Disks for Mirrored Plex: 1.1.6,
1.1.8,
1.1.9
Free Space Reallocation: off
HA Policy: sfo
Space Reserved for Snapshot Copies: 5%
Hybrid Enabled: false
Available Size: 53.10MB
Block Type: 64-bit
Checksum Enabled: true
Checksum Status: active
Cluster: cluster1
Home Cluster ID: 686964a0-2172-11e3-

```

837d-123478563412

DR Home ID: -
DR Home Name: -
Has Mroot Volume: false
Has Partner Node Mroot Volume: false
Home ID: 4050409551
Home Name: cluster1-02
Total Hybrid Cache Size: 0B
Hybrid: false
Inconsistent: false
Is Aggregate Home: true
Max RAID Size: 16
Hybrid Aggregate SSD Tier Maximum RAID Group Size: -
Owner ID: 4050409551
Owner Name: cluster1-02
Used Percentage: 0%
Plexes: /aggr1/plex0,
/aggr1/plex1
RAID Groups: /aggr1/plex0/rg0
(block) /aggr1/plex1/rg0
(block)
RAID Status: raid_dp, mirrored,
normal
RAID Type: raid_dp
SyncMirror Resync Snapshot Frequency in Minutes: 60
Is Root: false
Space Used By metadata for Volume Efficiency: 0B
Size: 53.24MB
SnapLock Type of the Aggregate: -
State: online
Used Size: 144KB
Number Of Volumes: 0
Is Flash Pool Caching: -
Is Eligible for Auto Balance Aggregate: false
State of the aggregate being balanced: ineligible
State Change Counter for Auto Balancer: 0
Is Encrypted: true
Encryption Key ID:
40004FE300000000030300000000000436F5DB53445FD603FB5A8A64937AA7B
Is in the precommit phase of Copy-Free Transition: false
Is a 7-Mode transitioning aggregate that is not yet committed in clustered
Data ONTAP and is currently out of space: false
Threshold When Aggregate Is Considered Unbalanced (%): 70
Threshold When Aggregate Is Considered Balanced (%): 40
Resynchronization Priority: -

```

Space Saved by Data Compaction: 99.24MB
  Percentage Saved by Data Compaction: 7%
  Amount of compacted data: 99.24MB
  Timestamp of Aggregate Creation: 1/3/2017 23:38:06
  Enable SIDL: off
  Composite: false
  Capacity Tier Used Size: 0B
Space Saved by Storage Efficiency: 99.24MB
  Percentage of Space Saved by Storage Efficiency: 7%
  Amount of Shared bytes count by Storage Efficiency: 99.24MB

```

The following example displays information about aggregates that are owned by nodes in cluster1:

```

cluster1::> storage aggregate show -cluster cluster1

cluster1:
Aggregate      Size Available Used% State  #Vols  Nodes           RAID
Status
-----
-----
aggr0          6.04GB    3.13GB   48% online    2 cluster1-01
raid_dp,
mirrored,
normal
aggr1          53.24MB   12.59MB  76% online    2 cluster1-02
raid_dp,
mirrored,
normal
2 entries were displayed.

```

The following example displays information about aggregates that are owned by nodes in the remote cluster named cluster2:

```
cluster1::> storage aggregate show -cluster cluster2
```

```
cluster2:
```

Aggregate	Size	Available	Used%	State	#Vols	Nodes	RAID
-----------	------	-----------	-------	-------	-------	-------	------

```
-----
```

aggr2	-	-	-	remote_cluster	-	-	-
-------	---	---	---	----------------	---	---	---

aggr3	-	-	-	remote_cluster	-	-	-
-------	---	---	---	----------------	---	---	---

```
2 entries were displayed.
```

The following example displays information about aggregates that are owned by nodes in all the clusters:

```
cluster1::> storage aggregate show -cluster *
```

```
cluster2:
```

Aggregate	Size	Available	Used%	State	#Vols	Nodes	RAID
-----------	------	-----------	-------	-------	-------	-------	------

```
-----
```

aggr2	-	-	-	remote_cluster	-	-	-
-------	---	---	---	----------------	---	---	---

aggr3	-	-	-	remote_cluster	-	-	-
-------	---	---	---	----------------	---	---	---

```
cluster1:
```

Aggregate	Size	Available	Used%	State	#Vols	Nodes	RAID
-----------	------	-----------	-------	-------	-------	-------	------

```
-----
```

aggr0	6.04GB	3.14GB	48%	online	2	cluster1-01	
-------	--------	--------	-----	--------	---	-------------	--

```
raid_dp,
```

```
mirrored,
```

```
normal
```

aggr1	53.24MB	12.59MB	76%	online	2	cluster1-02	
-------	---------	---------	-----	--------	---	-------------	--

```
raid_dp,
```

```
mirrored,
```

```
normal
```

```
4 entries were displayed.
```


Related Links

- [storage aggregate show-space](#)
- [storage aggregate resynchronization modify](#)

storage aggregate verify

Verify an aggregate

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

Description

The `storage aggregate verify` command verifies the two plexes of an aggregate. It compares the data in the two plexes to ensure that the plexes are identical. It can be used whenever the administrator needs to ensure that the two plexes are completely synchronized with each other. To view any discrepancies, use the following command:

```
event log show -message-name raid.mirror.verify.mismatch
```

Parameters

-aggregate <aggregate name> - Aggregate

This parameter specifies the aggregate to be verified. If no aggregate is specified then the action specified by the parameter `-action` will be taken on all the aggregates.

-action {start|stop|resume|suspend|status} - Action

This parameter specifies the action to be taken. The possible actions are:

- `start` - Starts a verify.
- `stop` - Permanently stops a verify. A stopped verify cannot be resumed.
- `resume` - Resumes a suspended verify.
- `suspend` - Suspends a verify.
- `status` - Displays the current status of a verify.

[-plex-to-fix <text>] - Plex to be Corrected in Case of Mismatches

This parameter specifies the name of a plex to fix in case the two plexes of the aggregate do not match. The default behavior is to log any discrepancies instead of fixing them.



This parameter is only applicable when the command is used to start a verify.

Examples

The following example starts a verify on an aggregate named `aggr1`.

```
cluster1::> storage aggregate verify -aggregate aggr1 -action start
```

The following example queries the status of a verify on an aggregate named aggr1.

```
cluster1::> storage aggregate verify -aggregate aggr1 -action status
Aggregate:aggr1, Is Suspended:false, Percentage Completed:19.03%
```

The following example starts a verify on all the aggregates.

```
cluster1::> storage aggregate verify -action start
```

storage aggregate efficiency show

Display aggregate storage efficiency details

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

Description

The `storage aggregate efficiency show` command displays information about the different storage efficiency of all the aggregates. If no parameters are specified, the command displays the following information for all aggregates:

- Aggregate
- Node
- Cross-vol-background-dedupe State (Enabled, Disabled)
- Cross-vol-inline-dedupe State (Enabled, Disabled)

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.

[-aggregate <aggregate name>] - Aggregate

Displays the aggregate name. If this parameter is specified, the command displays detailed information about the storage efficiency of the specified aggregate.

[-node {<nodename>|local}] - Node

Displays the node which owns the aggregate. If this parameter is specified, the command displays storage efficiency information only about the aggregates that are located on the specified node.

[`-cross-volume-background-dedupe {true|false}`] - Cross Volume Background Deduplication

Displays whether the cross volume background deduplication is enabled/disabled in the aggregate.

[`-cross-volume-inline-dedupe {true|false}`] - Cross Volume Inline Deduplication

Displays whether the cross volume inline deduplication is enabled/disabled in the aggregate.

[`-cross-volume-dedupe-savings {true|false}`] - Has Cross Volume Deduplication Savings

Displays whether the aggregate may have savings from cross volume deduplication.

[`-auto-adaptive-compression-savings {true|false}`] - Has Auto Adaptive Compression Savings

Displays whether the aggregate may have auto adaptive compression savings.

Examples

The following example displays information about all aggregates that are owned by nodes in the local cluster:

```
cluster:::> storage aggregate efficiency show

Aggregate: aggr0
  Node: vivek6-vsims2

Has Cross Volume Deduplication Savings:           false
Cross Volume Background Deduplication:           false
Cross Volume Inline Deduplication:                false

Aggregate: aggr1
  Node: vivek6-vsims2

Has Cross Volume Deduplication Savings:           true
Cross Volume Background Deduplication:           true
Cross Volume Inline Deduplication:                true
2 entries were displayed.
```

storage aggregate efficiency stat

Display aggregate storage efficiency statistics

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

Description

The `storage aggregate efficiency stat` command displays storage efficiency statistics.

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.

[-aggregate <aggregate name>] - Aggregate

If this parameter is specified, the command displays storage efficiency statistics for the specified aggregate.

[-node {<nodename>|local}] - Node

If this parameter is specified, the command displays storage efficiency statistics for the aggregates that are located on the specified node.

[-total-blocks-recompressed <integer>] - Total Number of Blocks Recompressed

This field represents the total number of blocks recompressed. If this parameter is specified, the command displays storage efficiency statistics that match the specified block count.

[-active-blocks-recompressed <integer>] - Number of Active Blocks Recompressed

This field represents the total number of hot blocks recompressed. If this parameter is specified, the command displays storage efficiency statistics that match the specified block count.

[-inactive-blocks-recompressed <integer>] - Number of Inactive Blocks Recompressed

This field represents the total number of cold blocks recompressed. If this parameter is specified, the command displays storage efficiency statistics that match the specified block count.

Examples

The following example displays storage efficiency statistics for all the aggregates that are owned by nodes in the local cluster:

```
cluster1::> storage aggregate efficiency stat
Aggregate: aggr1
                                         Node: node1
  Total Number of Blocks Recompressed: 512
    Number of Active Blocks Recompressed: 120
    Number of Inactive Blocks Recompressed: 392
Aggregate: aggr2
                                         Node: node1
  Total Number of Blocks Recompressed: 0
    Number of Active Blocks Recompressed: 0
    Number of Inactive Blocks Recompressed: 0

2 entries were displayed.
```

storage aggregate efficiency cross-volume-dedupe revert-to

Reverts the cross volume deduplication savings on an aggregate

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

Description

The `storage aggregate efficiency cross-volume-dedupe revert-to` command is used to revert cross volume deduplication savings on an aggregate.

Parameters

-aggregate <aggregate name> - Aggregate

This specifies the aggregate on which cross volume deduplication savings should be reverted. If no aggregate is specified then it will revert the savings on all aggregates

[-c, -clean-up {true|false}] - Delete Previously Downgraded Metafiles

This specifies whether downgrade metafile needs to be removed so that other efficiency operations can start on that aggregate.

Examples

The following example displays information for reverting cross volume background deduplication on aggregate "aggr1":

```
cluster:::> storage aggregate efficiency cross-volume-dedupe revert-to
-aggregate aggr1
The revert operation started on aggregate "aggr1".

cluster:::> storage aggregate efficiency cross-volume-dedupe revert-to
-aggregate aggr1 -clean-up true
The revert operation started on aggregate "aggr1".
```

storage aggregate efficiency cross-volume-dedupe show

Display aggregate cross volume deduplication efficiency details

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

Description

The `storage aggregate efficiency cross-volume-dedupe show` command displays information in detail about the different storage efficiency of all the aggregates. If no parameters are specified, the command displays the following information for all aggregates:

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.

[-aggregate <aggregate name>] - Aggregate

Displays the aggregate name. If this parameter is specified, the command displays detailed information about the storage efficiency of the specified aggregate.

[-node {<nodename>|local}] - Node

Displays the node which owns the aggregate. If this parameter is specified, the command displays storage efficiency information only about the aggregates that are located on the specified node.

[-background-progress <text>] - Progress

Displays the information for the aggregates that match the specified progress.

[-background-op-status <text>] - Operation Status

Displays the information for the aggregates that match the specified operation status.

[-background-last-op-state <text>] - Last Operation State

Displays the information for the aggregates that match the specified last operation state.

[-background-last-success-op-begin <Date>] - Last Success Operation Begin Time

Displays the information for the aggregates that match the specified last successful operation begin time.

[-background-last-success-op-end <Date>] - Last Success Operation End Time

Displays the information for the aggregates that match the specified last successful operation end time.

[-background-last-op-begin <Date>] - Last Operation Begin Time

Displays the information for the aggregates that match the specified last operation begin time.

[-background-last-op-end <Date>] - Last Operation End Time

Displays the information for the aggregates that match the specified last operation end time.

[-background-last-op-error <text>] - Last Operation Error

Displays the information for the aggregates that match the specified last operation error.

[-background-stage <text>] - Stage

Displays the information for the aggregates that match the specified stage.

[-background-checkpoint-time <Date>] - Checkpoint Time

Displays the information for the aggregates that match the specified checkpoint time.

[-background-checkpoint-op-type <text>] - Checkpoint Operation Type

Displays the information for the aggregates that match the specified checkpoint operation type.

[`-background-checkpoint-stage <text>`] - Checkpoint Stage

Displays the information for the aggregates that match the specified checkpoint stage.

[`-background-dedupe {true|false}`] - Background State

Displays the information for the aggregates that match the specified cross volume background dedupe state.

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

Displays the information for the aggregates that match the specified cross volume inline dedupe state.

[`-dedupe-savings {true|false}`] - Has Cross Volume Deduplication Savings

Displays the information for the aggregates that has some savings from cross volume deduplication.

Examples

The following example displays information about all aggregates that are owned by nodes in the local cluster:

```
cluster:::> storage aggregate efficiency cross-volume-dedupe show

Aggregate: aggr0
  Node: vivek6-vsimg2

Has Cross Volume Deduplication Savings:                               false

-----:Cross Volume Background Deduplication Status:-----
State:                                                                    false
Progress:                                                                    -
Operation Status:                                                            Idle
Last Operation State:                                                        Success
Last Success Operation Begin Time:                                           -
Last Success Operation End Time:                                             -
Last Operation Begin Time:                                                   -
Last Operation End Time:                                                     -
Last Operation Error:                                                         Operation
succeeded
Stage:                                                                        -
Checkpoint Time:                                                             -
Checkpoint Operation Type:                                                   -
Checkpoint Stage:                                                            -

-----:Cross Volume Inline Deduplication Status:-----
State:                                                                    false
Aggregate: aggr1
  Node: vivek6-vsimg2

Has Cross Volume Deduplication Savings:                               true
```

```

-----:Cross Volume Background Deduplication Status:-----
State: true
Progress: -
Operation Status: Idle
Last Operation State: Success
Last Success Operation Begin Time: Wed Aug 30
06:31:50 2017
Last Success Operation End Time: Wed Aug 30
06:31:50 2017
Last Operation Begin Time: Wed Aug 30
06:31:50 2017
Last Operation End Time: Wed Aug 30
06:31:50 2017
Last Operation Error: Operation
succeeded
Stage: Cross
volume sharing Done
Checkpoint Time: -
Checkpoint Operation Type: -
Checkpoint Stage: -

-----:Cross Volume Inline Deduplication Status:-----
State: true

2 entries were displayed.

```

storage aggregate efficiency cross-volume-dedupe start

Starts the cross volume background deduplication on an aggregate

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

Description

The `storage aggregate cross-volume-dedupe start` command is used to start cross volume background deduplication on an aggregate.

Parameters

-aggregate <aggregate name> - Aggregate

This specifies the aggregate on which cross volume background deduplication should be started. If no aggregate is specified then it will start on all aggregates

[-s, -scan-old-data <true>] - Scan Old Data

This option processes all the existing data on all volumes on the aggregate. It prompts for user confirmation before proceeding. Default value is `false`.

Examples

The following example displays information for starting cross volume background deduplication on aggregate "aggr1":

```
cluster:::> storage aggregate efficiency cross-volume-dedupe start
-aggregate aggr1
The efficiency operation on aggregate "aggr1" has started.

cluster:::> storage aggregate efficiency cross-volume-dedupe start
-aggregate aggr1 -scan-old-data true
The efficiency operation on aggregate "aggr1" has started.
```

storage aggregate efficiency cross-volume-dedupe stop

Stops the cross volume background deduplication on an aggregate

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

Description

The `storage aggregate cross-volume-dedupe stop` command is used to stop cross volume background deduplication on an aggregate.

Parameters

-aggregate <aggregate name> - Aggregate

This specifies the aggregate on which cross volume background deduplication should be stopped. If no aggregate is specified then it will stop on all aggregates

Examples

The following example displays information for stopping cross volume background deduplication on aggregate "aggr1":

```
cluster:::> storage aggregate efficiency cross-volume-dedupe stop
-aggregate aggr1
The efficiency operation on aggregate "aggr1" is being stopped.
```

storage aggregate encryption show-key-id

Display encrypted aggregate information

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

Description

The `storage aggregate encryption show-key-id` command displays the key IDs of all NAE (NetApp Aggregate Encryption) aggregates.

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.

[-aggregate <text>] - Aggregate

If this parameter is specified, the command displays information only about the specific NAE (NetApp Aggregate Encryption) aggregate.

[-aggrID <UUID>] - Aggregate UUID

If this parameter is specified, the command displays the key ID of the specified NAE (NetApp Aggregate Encryption) aggregate ID.

[-keyid-index-zero <text>,...] - 0th Index Keyid

If this parameter is specified, the command displays the 0th index key ID of NAE (NetApp Aggregate Encryption) aggregates.

storage aggregate inode-upgrade resume

Resume suspended inode upgrade

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

Description

The `storage aggregate inode-upgrade resume` command resumes a suspended inode upgrade process. The inode upgrade process might have been suspended earlier due to performance reasons.

Parameters

-node {<nodename>|local} - Node Name

If this parameter is specified, the command resumes the upgrade process of an aggregate that is located on the specified node.

-aggregate <aggregate name> - Aggregate Name

This specifies the aggregate for which the inode upgrade process is to be resumed.

Examples

The following example resumes an aggregate upgrade process:

```
cluster1::> storage aggregate inode-upgrade resume -aggregate aggr1
```

storage aggregate inode-upgrade show

Display inode upgrade progress

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

Description

The `storage aggregate inode-upgrade show` command displays information about aggregates undergoing 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 aggregates undergoing the inode upgrade process. The default fields are:

- aggregate
- status
- scan-percent
- remaining-time
- space-needed
- scanner-progress

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.

[-aggregate <text>] - Aggregate Name

If this parameter is specified with the `-node` parameter, the command displays detailed information about the specified aggregate. If only this parameter is specified, the command displays information about all aggregates that match the specified name.

[-node <nodename>] - Node Name

If this parameter is specified, the command displays information only about the aggregate or aggregates that are located on the specified node.

[-status {pending|scanning|suspended-initializing|suspended|cleanup-pending|cleanup|cleanup-done|suspended-aborting|suspended-removing|suspended-while-removing|suspended-ironing}] - Upgrade Status

If this parameter is specified, the command displays information only about the aggregate or aggregates that match the specified inode upgrade status.

[`-scan-percent` <percent>] - Upgrade Scan Percent Complete

If this parameter is specified, the command displays information only about the aggregate or aggregates that match the specified inode upgrade progress percentage.

[`-space-needed` {<integer>[KB|MB|GB|TB|PB]}] - Space Needed to Complete Upgrade

If this parameter is specified, the command displays information only about the aggregate or aggregates where the space needed to complete the upgrade process matches the specified size.

[`-remaining-time` [<integer>h][<integer>m][<integer>s]>] - Remaining Upgrade Time

If this parameter is specified, the command displays information only about the aggregate or aggregates where the remaining time to complete the inode upgrade process matches the specified time.

[`-scanner-progress` <text>] - Scanner Progress

If this parameter is specified, the command displays information only about the aggregate or aggregates where the progress of the inode upgrade process matches the input.

Examples

The following example displays information about all aggregates undergoing the inode upgrade process:

```
cluster1::> storage aggregate inode-upgrade show
Aggregate Status      %Complete  Time Remaining  Space Needed  Inode Progress
-----
aggr0      pending    0%             -              20.36MB      Public : Inode 0
out of 65562
aggr1      pending    0%             -              19.84MB      Public : Inode 0
out of 63714
```

storage aggregate object-store attach

Attach an object store to an aggregate

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

Description

The `storage aggregate object-store attach` command attaches an object store to an aggregate to create a FabricPool. This command requires two parameters to create a FabricPool - an aggregate and a configuration to attach an object-store to the aggregate. This command verifies whether the object store is accessible through the intercluster LIF both from the node on which the aggregate is present as well as its High Availability (HA) partner node. The command fails if the object store is not accessible. Once an object store is attached to an aggregate, it cannot be detached.

Parameters

`-aggregate` <text> - Name of the Aggregate

This parameter specifies the aggregate to which the object store must be attached to create a FabricPool.

-object-store-name <text> - Object Store Configuration Name

This parameter specifies the object store configuration that describes the object store to be attached. The object store configuration has information about object store server name, port, access credentials, and provider type.

[-allow-flexgroup {true|false}] - Allow Existing FlexGroup Constituents in the Aggregate on Attach

This optional parameter allows attaching object store to an aggregate already containing FlexGroup constituents. The default value is false. Mixing FabricPools and non-FabricPools within a FlexGroup is not recommended. All aggregates hosting constituents of a FlexGroup should be attached to the object store.

Examples

The following example attaches an object store to aggregate aggr1:

```
cluster1::>storage aggregate object-store attach -aggregate aggr1 -object
-store-name my-store
```

storage aggregate object-store mirror

Attaches a second object store to a FabricPool aggregate to create a mirror

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

Description

The `storage aggregate object-store mirror` command attaches a second object store to a FabricPool aggregate to create a mirror. This command verifies whether the object store is accessible through the intercluster LIF from both the node on which the aggregate is present and from its high-availability (HA) partner node. The command fails if the object store is not accessible.

Parameters

-aggregate <text> - Name of the Aggregate

This parameter specifies the aggregate to which the object store must be attached in order to create a mirror.

-object-store-name <text> - Object Store Configuration Name

This parameter specifies the name of the new object store configuration to be attached to the aggregate. The object store configuration has information about the object store server name, port, access credentials, and provider type.

Examples

The following example shows how to create a mirror to aggregate aggr1:

```
cluster1::>storage aggregate object-store mirror -aggregate aggr1 -object
-store-name my-store-2
```

storage aggregate object-store modify

Modify attributes of object stores attached to an aggregate

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

Description

The `storage aggregate object-store modify` command is used to update one or more object store parameters.

Parameters

-aggregate <text> - Aggregate Name

This parameter identifies the aggregate to which the object store to be modified is attached.

-object-store-name <text> - ONTAP Name for this Object Store Config

This parameter identifies the configuration name of the object store to be modified.

[-unreclaimed-space-threshold <percent>] - Threshold for Reclaiming Unreferenced Space

This optional parameter specifies the usage threshold below which Data ONTAP reclaims unused space from objects in the object store. When Data ONTAP writes data to the object store, it packages multiple file system blocks into one object. Over time, blocks stored in an object can be freed, leaving part of the object unused. When the percentage of used blocks in an object falls below this threshold, a background task moves the blocks which are still used to a new object. Afterwards, Data ONTAP frees the original object to reclaim the unused space. Valid values are between 0% and 99%. The default value depends on the object store's provider type. It is 20% for *AWS_S3*, 25% for *Azure_Cloud*, 40% for *SGWS*, 20% for *IBM_COS*, 20% for *AliCloud*, 20% for *GoogleCloud* and 40% for *ONTAP_S3*. Consult the FabricPool best practices guidelines for more information.

[-tiering-fullness-threshold <percent>] - Aggregate Fullness Threshold Required for Tiering

This optional parameter specifies the percentage of space in the performance tier which must be used before data is tiered out to the capacity tier.

[-mirror-type {primary|mirror}] - Object Store Mirror Type

This parameter specifies the object store mirror type. Valid mirror types are *primary* or *mirror*.

[-force-tiering-on-metrocluster {true|false}] - Force Tiering with no Mirror in a MetroCluster Configuration

This parameter specifies force tiering option enabled or not on primary object stores for aggregates in a MetroCluster configuration.

Examples

The following example modifies the unreclaimed space threshold of an object store attached to an aggregate named `aggr1`:

```
cluster1::> storage aggregate object-store modify -aggregate aggr1 -object
-store-name my-store -unreclaimed-space-threshold 20%
```

storage aggregate object-store show-freeing-status

Show status of background object freeing work after aggregate delete

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

Description

The `storage aggregate object-store show-freeing-status` command displays status information about the background work that frees an aggregate's objects from an object store after a [storage aggregate 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.

[-bin-uuid <UUID>] - UUID of the Bin

If this parameter is specified, the command displays information only about the aggregate attached to the specified bin UUID.

[-config-id <integer>] - Object Store Config ID

If this parameter is specified, the command displays information only about the aggregate attached to the object-store with specified config ID.

[-object-store-name <text>] - Object Store Configuration Name

If this parameter is specified, the command displays information only about object stores whose configuration name matches the specified names.

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

If this parameter is specified, the command displays information only about the specified aggregates that were deleted.

[-request-state {queued|running|cleaning-up|finishing}] - Request State

If this parameter is specified, the command displays information only about the object stores that have the specified object freeing request state.

[-num-objects-freed <integer>] - Num Objects Freed

If this parameter is specified, the command displays information only about the object stores that have the specified number of objects that have been freed.

`[-last-error <text>]` - The Last Error Encountered

If this parameter is specified, the command displays information only about the object stores that have the specified last error encountered.

Related Links

- [storage aggregate delete](#)

storage aggregate object-store show-resync-status

Display object store mirror resync progress

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

Description

The `storage aggregate object-store show-resync-status` command displays information about the progress of resync for mirror object stores attached to a FabricPool.

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.

`[-aggregate <aggregate name>]` - Aggregate

If this parameter is specified, the command displays information only about the specified aggregates.

`[-primary-object-store-name <text>]` - Primary Object Store Name

If this parameter is specified, the command displays information only about aggregates whose primary object store configuration name matches the specified names.

`[-mirror-object-store-name <text>]` - Mirror Object Store Name

If this parameter is specified, the command displays information only about aggregates whose mirror object store configuration name matches the specified names.

`[-percent-complete <percent>]` - Resync Complete Percentage

If this parameter is specified, the command displays information only about aggregates whose mirror object store resync progress percentage matches the specified value.

Examples

The following example displays resync progress for all aggregates with mirror object stores which are not in sync:


```
cluster1::>storage aggregate object-store show-resync-status
```

storage aggregate object-store show-space

Display space utilization of object stores attached to an aggregate

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

Description

The `storage aggregate object-store show-space` command displays information about the amount of space used in the object store for each of the aggregates in FabricPool. The used space is displayed in both absolute size as well as a percentage of the FabricPool license 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.

[-aggregate <text>] - Aggregate Name

If this parameter is specified, the command displays space information only about object stores that are attached to the specified aggregates.

[-object-store-name <text>] - ONTAP Name for this Object Store Config

If this parameter is specified, the command displays space information only about object stores whose configuration name matches the specified names.

[-object-store-availability <object store availability>] - Availability of the Object Store

If this parameter is specified, the command displays space information about the object stores whose availability status matches the specified value. Supported values with this parameter are - available and unavailable.

[-object-store-unavailable-reason <text>] - Reason why Object Store is Unavailable

If this parameter is specified, the command displays information only about object stores whose unavailability reason matches the specified reason.

[-provider-type <providerType>] - Type of the Object Store Provider

If this parameter is specified, the command displays information only about object store configurations whose provider type matches the specified value.

[-license-used-percent <percent_no_limit>] - License Space Used Percent

If this parameter is specified, the command displays space information only about object stores whose space used by the associated aggregate as a percentage of the license limit matches the specified value. If the object store does not require a license, then this field is not set.

[`-unreclaimed-space-threshold <percent>`] - Threshold for Reclaiming Unreferenced Space

If this parameter is specified, the command displays information only about object stores whose threshold for reclaiming unused space from objects in the object store matches the specified value.

[`-tiering-fullness-threshold <percent>`] - Aggregate Fullness Threshold Required for Tiering

If this parameter is specified, the command displays information only about object stores whose performance tier fullness threshold for tiering matches the specified value.

[`-mirror-type {primary|mirror}`] - Object Store Mirror Type

If this parameter is specified, the command displays information about object stores with a mirror-type that matches the specified value.

[`-is-mirror-degraded {true|false}`] - This object store is in mirror degraded mode

If this parameter is specified, the command displays information only about mirror object stores which have the specified mirror degraded state.

[`-force-tiering-on-metrocluster {true|false}`] - Force Tiering with no Mirror in a MetroCluster Configuration

If this parameter is specified, the command displays information only about primary object stores for which force tiering is toggled on for aggregates in a MetroCluster configuration.

[`-cluster <Cluster name>`] - The name of the Cluster to which the bin belongs

If this parameter is specified, the command displays information only about object stores for which cluster matches the specified value.

Examples

The following example displays space information about all object stores:

```
cluster1::>storage aggregate object-store show-space
```

storage aggregate object-store show

Display the details of object stores attached to an aggregate

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

Description

The `storage aggregate object-store show` command displays information about all the object stores in the system.

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.

[`-aggregate <text>`] - Aggregate Name

If this parameter is specified, the command displays information only about the object stores that are attached to the specified aggregates.

[`-object-store-name <text>`] - ONTAP Name for this Object Store Config

If this parameter is specified, the command displays information only about object stores whose configuration name matches the specified names.

[`-object-store-availability <Object Store Availability>`] - Availability of the Object Store

If this parameter is specified, the command displays information only about object stores whose availability status matches the specified value. Supported values with this parameter are `- available` and `- unavailable`.

[`-object-store-unavailable-reason <text>`] - Reason why Object Store is Unavailable

If this parameter is specified, the command displays information only about object stores whose unavailability reason matches the specified reason.

[`-provider-type <providerType>`] - Type of the Object Store Provider

If this parameter is specified, the command displays information only about object store configurations whose provider type matches the specified value.

[`-license-used-percent <percent_no_limit>`] - License Space Used Percent

If this parameter is specified, the command displays information only about object stores whose space used by the aggregate as a percentage of the license limit matches the specified value.

[`-unreclaimed-space-threshold <percent>`] - Threshold for Reclaiming Unreferenced Space

If this parameter is specified, the command displays information only about object stores whose threshold for reclaiming unused space from objects in the object store matches the specified value.

[`-tiering-fullness-threshold <percent>`] - Aggregate Fullness Threshold Required for Tiering

If this parameter is specified, the command displays information only about object stores whose performance tier fullness threshold for tiering matches the specified value.

[`-mirror-type {primary|mirror}`] - Object Store Mirror Type

If this parameter is specified, the command displays information about object stores with a mirror-type that matches the specified value.

[`-is-mirror-degraded {true|false}`] - This object store is in mirror degraded mode

If this parameter is specified, the command displays information only about mirror object stores which have the specified mirror degraded state. When a mirror object store is attached to a FabricPool, it is initially degraded because it does not contain a copy of all the data in the primary object store. While the mirror is degraded, all reads are served from the primary object store, and the mirror cannot be promoted to become the primary. After the resync process copies all data from the primary object store to the mirror, the mirror is no longer degraded. From that point on the mirror is always kept in sync with the primary and never becomes degraded again.

[`-force-tiering-on-metrocluster {true|false}`] - Force Tiering with no Mirror in a MetroCluster Configuration

If this parameter is specified, the command displays information only about primary object stores for which force tiering is toggled on for aggregates in a MetroCluster configuration.

[-cluster <Cluster name>] - The name of the Cluster to which the bin belongs

If this parameter is specified, the command displays information only about object stores for which cluster matches the specified value.

Examples

The following example displays all information about all object stores:

```
cluster1::>storage aggregate object-store show
```

storage aggregate object-store unmirror

Remove the second object store from a mirrored FabricPool

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

Description

The `storage aggregate object-store unmirror` command detaches a second object store from a mirrored FabricPool. This command verifies whether the primary object store is accessible through the intercluster LIF from both the node on which the aggregate is present and from its high-availability (HA) partner node. The command fails if the primary object store is not available.

Parameters

-aggregate <text> - Name of the Aggregate

This parameter specifies the aggregate from which the mirrored object store must be removed.

Examples

The following example shows how to unmirror a mirrored FabricPool `aggr1`:

```
cluster1::>storage aggregate object-store unmirror -aggregate aggr1
```

storage aggregate object-store config create

Define the configuration for an object store

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

Description

The `storage aggregate object-store config create` command is used by a cluster administrator to tell Data ONTAP how to connect to an object store. Following pre-requisites must be met before creating an

object store configuration in Data ONTAP.

- A valid data bucket or container must be created with the object store provider. This assumes that the user has valid account credentials with the object store provider to access the data bucket.
- The Data ONTAP node must be able to connect to the object store. This includes
- Fast, reliable connectivity to the object store.
- An inter-cluster LIF (Logical Interface) must be configured on the cluster. Data ONTAP will verify connectivity prior to saving this configuration information.
- If SSL/TLS authentication is required, then valid certificates must be installed.
- FabricPool license (required for Amazon S3 and Azure object stores).

An object-store configuration once created must not be reassociated with a different object-store or container. See [storage aggregate object-store config modify](#) command for more information. If neither the `access-key` nor the `secret-password` are provided while setting up a configuration for `AWS_S3` object store in Cloud Volumes ONTAP, then the access key (access key ID), the secret password (secret access key), and the session token will be retrieved from EC2 instance metadata for the AWS Identity and Access Management (IAM) role associated with the EC2 instance. If Data ONTAP is unable to create a object store configuration, then the command will fail explaining the reason for failure.

Parameters

-object-store-name <text> - Object Store Configuration Name

This parameter specifies the name that will be used to identify the object store configuration. The name can contain the following characters: `"`, `-`, `A-Z`, `a-z`, and `0-9`. *The first character must be one of the following: `"`, `A-Z`, or `a-z`.*

-provider-type <providerType> - Type of the Object Store Provider

This parameter specifies the type of object store provider that will be attached to the aggregate. Valid options are: `AWS_S3` (Amazon S3 storage), `Azure_Cloud` (Microsoft Azure Cloud), `SGWS` (StorageGrid WebScale), `IBM_COS` (IBM Cloud Object Storage), `AliCloud` (Alibaba Cloud Object Storage Service), `GoogleCloud` (Google Cloud Storage) and `ONTAP_S3`.

[-auth-type <object_store_auth_type>] - Authentication Used to Access the Object Store

This parameter specifies where the system obtains credentials for authentication to an object store. The available choices depend on the platform (Cloud Volumes ONTAP or not) and `provider-type` (`AWS_S3` or not). The `keys` value is always applicable, and if selected means that the `access-key` and `secret-password` are provided by the system administrator. In Cloud Volumes ONTAP, the `EC2-IAM` value is also applicable. It means that the IAM role is associated with the EC2 instance, and that the `access-key`, `secret-password` and session token are retrieved from EC2 instance metadata for this IAM role. Note that `-use-iam-role` and `-auth-type` are mutually exclusive, `-auth-type EC2-IAM` is an equivalent of `-use-iam-role true`, and `-auth-type key` is an equivalent of `-use-iam-role false`. In Cloud Volumes ONTAP, the `GCP-SA` value may also be applicable. It means that a session token is retrieved from the GCP instance metadata for the Service Account associated with the GCP instance. Similarly, `Azure-MSI` means that a session token is retrieved from the Azure instance metadata for the Managed Service Identity (MSI) associated with the Azure instance. For the `AWS_S3` provider, the `CAP` (C2S Authentication Portal) value is also applicable. This should only be used when accessing C2S (Commercial Cloud Services). If the `CAP` value is specified, then the `-cap-url` must be specified. See `cap-url`.

[-cap-url <text>] - URL to Request Temporary Credentials for C2S Account

This parameter is available only when `-auth-type` is `CAP`. It specifies a full URL of the request to a CAP

server for retrieving temporary credentials (access-key, secret-password and session token) for accessing the object store server. The CAP URL may look like: <https://123.45.67.89:1234/CAP/api/v1/credentials?agency=myagency=mymission=myrole>

-server <Remote InetAddress> - Fully Qualified Domain Name of the Object Store Server

This parameter specifies the Fully Qualified Domain Name (FQDN) of the remote object store server. For Amazon S3, server name must be an AWS regional endpoint in the format `s3.amazonaws.com` or `s3-<region>.amazonaws.com`, for example, `s3-us-west-2.amazonaws.com`. The region of the server and the bucket must match. For more information on AWS regions, refer to 'Amazon documentation on AWS regions and endpoints'. For Azure, if the `-server` is a "blob.core.windows.net" or a "blob.core.usgovcloudapi.net", then a value of `-azure-account` followed by a period will be added in front of the server.

[-is-ssl-enabled {true|false}] - Is SSL/TLS Enabled

This parameter indicates whether a secured SSL/TLS connection will be used during data access to the object store. The default value is `true`.

[-port <integer>] - Port Number of the Object Store

This parameter specifies the port number on the remote server that Data ONTAP will use while establishing connection to the object store.

-container-name <text> - Data Bucket/Container Name

This parameter specifies the data bucket or container that will be used for read and write operations.



This name cannot be modified once a configuration is created.

{ [-access-key <text>] - Access Key ID for S3 Compatible Provider Types

This parameter specifies the access key (access key ID) required to authorize requests to the AWS S3, SGWS, IBM COS object stores and ONTAP_S3. For an Azure object store see `-azure-account`.

[-secret-password <text>] - Secret Access Key for S3 Compatible Provider Types

This parameter specifies the password (secret access key) to authenticate requests to the AWS S3, SGWS, IBM COS object stores and ONTAP_S3. If the `-access-key` is specified but the `-secret-password` is not, then one will be asked to enter the `-secret-password` without echoing the input. For an Azure object store see `-azure-private-key`.

[-azure-account <text>] - Azure Account

This parameter specifies the account required to authorize requests to the Azure object store. For other object store providers see `access-key`.



The value of this field cannot be modified once a configuration is created.

[-ask-azure-private-key {true|false}] - Ask to Enter the Azure Access Key without Echoing

If this parameter is `true` then one will be asked to enter `-azure-private-key` without echoing the input. Default value: `true`.

[-azure-private-key <text>] - Azure Access Key }

This parameter specifies the access key required to authenticate requests to the Azure object store. See also `ask-azure-private-key`. For other object store providers see `-secret-password`.

[`-ipSPACE <IPspace>`] - IPspace to Use in Order to Reach the Object Store

This optional parameter specifies the IPspace to use to connect to the object store. Default value: *Default*

[`-is-certificate-validation-enabled {true|false}`] - Is SSL/TLS Certificate Validation Enabled

This parameter indicates whether an SSL/TLS certificate of an object store server is validated whenever an SSL/TLS connection to an object store server is established. This parameter is only applicable when `is-ssl-enabled` is `true`. The default value is `true`. It is recommended to use the default value to make sure that Data ONTAP connects to a trusted object store server, otherwise identities of an object store server are not verified.

[`-use-http-proxy {true|false}`] - Use HTTP Proxy

This optional parameter indicates whether an HTTP proxy will be used for connecting to an object store. Note that an HTTP proxy is configured using the `vserver http-proxy` commands at the `diagnostic` privilege level. Default value: *false*.

[`-cluster <Cluster name>`] - The Name of the Cluster to which the Configuration Belongs

This optional parameter should only be specified in MetroCluster switched-over mode and specifies the name of the cluster for which the configuration must be created. By default the configuration is created for the local cluster.

[`-server-side-encryption {none | SSE-S3}`] - Encryption of Data at Rest by the Object Store Server

This parameter specifies if AWS or other S3 compatible object store server must encrypt data at rest. The available choices depend on provider-type. `none` encryption (no encryption required) is supported by all types of S3 (non-Azure) object store servers. `SSE-S3` encryption is supported by and is a default for all types of S3 (non-Azure) object store servers except `ONTAP_S3`. This is an advanced property. In most cases it is best not to change default value of "sse_s3" for object store servers which support SSE-S3 encryption. The encryption is in addition to any encryption done by ONTAP at a volume or at an aggregate level.

[`-url-style {path-style | virtual-hosted-style}`] - URL Style Used to Access S3 Bucket

This parameter specifies the URL style used to access S3 bucket. This option is only available for non-Azure object store providers. The available choices and default value depend on provider-type.

Examples

The following example creates an object store configuration in Data ONTAP:

```
cluster1::>storage aggregate object-store config create -object-store-name
  my_aws_store -provider-type AWS_S3 -server s3.amazonaws.com
  -container-name my-aws-bucket -access-key DXJRXHPXHYXA9X31X3JX
```

Related Links

- [storage aggregate object-store config modify](#)

storage aggregate object-store config delete

Delete the configuration of an object store

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

Description

The `storage aggregate object-store config delete` command removes an existing object store configuration in Data ONTAP. The configuration cannot be deleted if it is used by any aggregates or if the system is still freeing objects from the object store from a previously executed [storage aggregate delete](#) command. The command [storage aggregate object-store show](#) can be used to view the aggregates attached to the object store before issuing the delete command.



The [storage aggregate object-store show](#) command will not display aggregates that have been previously deleted but still has objects in the object store.

Parameters

-object-store-name <text> - Object Store Configuration Name

This parameter specifies the object store configuration to be deleted.

Examples

The following example deletes an object store configuration named my-store:

```
cluster1::>storage aggregate object-store config delete -object-store-name  
my-store
```

Related Links

- [storage aggregate delete](#)
- [storage aggregate object-store show](#)

storage aggregate object-store config modify

Modify object store configuration attributes

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

Description

The `storage aggregate object-store config modify` command is used to update one or more of object store configuration parameters. This command must not be used to reassociate an existing valid object-store configuration to a new object-store or container. The modifiable parameters include object store name, server name, port, access_key, secret-password, ipspace and is-ssl-enabled.

Parameters

-object-store-name <text> - Object Store Configuration Name

This parameter identifies the configuration to be modified.

[-new-object-store-name <text>] - Object Store Configuration New Name

This optional parameter specifies the new name for the object store configuration.

[-auth-type <object_store_auth_type>] - Authentication Used to Access the Object Store

This optional parameter specifies where the system obtains credentials for authentication to an object store. The available choices depend on the platform (Cloud Volumes ONTAP or not) and provider-type (AWS_S3 or not). The *keys* value is always applicable, and if selected means that the *access-key* and *secret-password* are provided by the system administrator. In Cloud Volumes ONTAP, the *EC2-IAM* value is also applicable. It means that the IAM role is associated with the EC2 instance, and that the *access-key*, *secret-password* and session token are retrieved from EC2 instance metadata for this IAM role. Note that *-use-iam-role* and *-auth-type* are mutually exclusive, *-auth-type EC2-IAM* is an equivalent of *-use-iam-role true*, and *-auth-type key* is an equivalent of *-use-iam-role false*. In Cloud Volumes ONTAP, the *GCP-SA* value may also be applicable. It means that a session token is retrieved from the GCP instance metadata for the Service Account associated with the GCP instance. Similarly, *Azure-MSI* means that a session token is retrieved from the Azure instance metadata for the Managed Service Identity (MSI) associated with the Azure instance. For the *AWS_S3* provider, the *CAP* (C2S Authentication Portal) value is also applicable. This should only be used when accessing C2S (Commercial Cloud Services). If the *CAP* value is specified, then the `-cap-url` must be specified. See `cap-url`.

[-cap-url <text>] - URL to Request Temporary Credentials for C2S Account

This parameter is available only when *-auth-type* is *CAP*. It specifies a full URL of the request to a CAP server for retrieving temporary credentials (*access-key*, *secret-pasword* and session token) for accessing the object store server. The CAP URL may look like: <https://123.45.67.89:1234/CAP/api/v1/credentials?agency=myagency=mymission=myrole>

[-server <Remote InetAddress>] - Fully Qualified Domain Name of the Object Store Server

This optional parameter specifies the new Fully Qualified Domain Name (FQDN) of the same object store server. For Amazon S3, server name must be an AWS regional endpoint in the format `s3.amazonaws.com` or `s3-<region>.amazonaws.com`, for example, `s3-us-west-2.amazonaws.com`. The region of the server and the bucket must match. For more information on AWS regions, refer to 'Amazon documentation on AWS regions and endpoints'. For Azure, if the *-server* is a "blob.core.windows.net" or a "blob.core.usgovcloudapi.net", then the value of *azure-account* in the configuration followed by a period will be added in front of the server. Note that the value of *azure-account* cannot be modified.

[-is-ssl-enabled {true|false}] - Is SSL/TLS Enabled

This optional parameter indicates whether a secured SSL/TLS connection will be used during data access to the object store.

[-port <integer>] - Port Number of the Object Store

This optional parameter specifies a new port number to connect to the object store server indicated in the *-server* parameter.

[-access-key <text>] - Access Key ID for S3 Compatible Provider Types

This optional parameter specifies a new access key (access key ID) for the AWS S3, SGWS, IBM COS object stores and ONTAP S3.

[`-secret-password <text>`] - Secret Access Key for S3 Compatible Provider Types

This optional parameter specifies a new password (secret access key) for the AWS S3, SGWS, IBM COS object stores and ONTAP S3. For an Azure object store see `-azure-private-key`. If the `-access-key` is specified but the `-secret-password` is not then one will be asked to enter the `-secret-password` without echoing the input.

[`-ask-azure-private-key {true|false}`] - Ask to Enter the Azure Access Key without Echoing

If this optional parameter is true then one will be asked to enter the `-azure-private-key` without echoing the input.

[`-azure-private-key <text>`] - Azure Access Key

This optional parameter specifies a new access key for Azure object store. For other object store providers see `secret-password`. See also `ask-azure-private-key`.

[`-ipspace <IPspace>`] - IPspace to Use in Order to Reach the Object Store

This optional parameter specifies new ipspace values for the configuration.

[`-is-certificate-validation-enabled {true|false}`] - Is SSL/TLS Certificate Validation Enabled

This optional parameter indicates whether an SSL/TLS certificate of an object store server is validated whenever an SSL/TLS connection to an object store server is established. This parameter is only applicable when `is-ssl-enabled` is `true`. It is recommended to keep the default value which is `true` to make sure that Data ONTAP connects to a trusted object store server, otherwise identities of an object store server are not verified.

[`-use-http-proxy {true|false}`] - Use HTTP Proxy

This optional parameter indicates whether an HTTP proxy will be used for connecting to an object store. Note that an HTTP proxy is configured using the `vserver http-proxy` commands at the `diagnostic` privilege level.

[`-server-side-encryption {none | SSE-S3}`] - Encryption of Data at Rest by the Object Store Server

This parameter specifies if AWS or other S3 compatible object store server must encrypt data at rest. The available choices depend on `provider-type`. `none` encryption (no encryption required) is supported by all S3 (non-Azure) object store servers. `SSE-S3` encryption is supported by all S3 (non-Azure) object store servers except `ONTAP_S3`. This is an advanced property. In most cases it is best not to change default value of `"sse_s3"` for object store servers which support SSE-S3 encryption. The encryption is in addition to any encryption done by ONTAP at a volume or at an aggregate level. Note that changing this option does not change encryption of data which already exist in the object store.

[`-url-style {path-style | virtual-hosted-style}`] - URL Style Used to Access S3 Bucket

This parameter specifies the URL style used to access S3 bucket. This option is only available for non-Azure object store providers. The available choices and default value depend on `provider-type`.

Examples

The following example modifies two parameters (port number and `is-ssl-enabled`) of an object store configuration named `my-store`:

```
cluster1::>storage aggregate object-store config modify -object-store-name
my-store -port 1235 -is-ssl-enabled true
```

storage aggregate object-store config rename

Rename an existing object store configuration

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

Description

The `storage aggregate object-store config rename` command is used to rename an existing object store configuration.

Parameters

-object-store-name <text> - Object Store Configuration Name

This parameter identifies an existing object store configuration.

-new-object-store-name <text> - Object Store Configuration New Name

This parameter specifies the new object store configuration name.

Examples

The following example renames an object store configuration from my-store to ms1:

```
cluster1::>storage aggregate object-store config rename -object-store-name
my-store -new-object-store-name ms1
```

storage aggregate object-store config show

Display a list of object store configurations

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

Description

The `storage aggregate object-store config show` command displays information about all existing object store configurations 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.

[`-object-store-name <text>`] - Object Store Configuration Name

If this parameter is specified, the command displays information only about object store configurations whose name matches the specified names.

[`-object-store-uuid <UUID>`] - UUID of the Object Store Configuration

If this parameter is specified, the command displays information only about object store configurations whose UUID matches the specified UUID values.

[`-provider-type <providerType>`] - Type of the Object Store Provider

If this parameter is specified, the command displays information only about object store configurations whose provider type matches the specified value.

[`-auth-type <object_store_auth_type>`] - Authentication Used to Access the Object Store

If this parameter is specified, the command displays information only about object store configurations whose authentication type matches the specified value.

[`-cap-url <text>`] - URL to Request Temporary Credentials for C2S Account

If this parameter is specified, the command displays information only about object store configurations whose CAP URL matches the specified value.

[`-server <Remote InetAddress>`] - Fully Qualified Domain Name of the Object Store Server

If this parameter is specified, the command displays information only about object store configurations whose server name matches the specified value. The server name is specified as a Fully Qualified Domain Name (FQDN).

[`-is-ssl-enabled {true|false}`] - Is SSL/TLS Enabled

If this parameter is specified, the command displays information only about object store configurations whose status about the use of secured communication over the network matches the specified value.

[`-port <integer>`] - Port Number of the Object Store

If this parameter is specified, the command displays information only about object store configurations whose port numbers matches the specified value.

[`-container-name <text>`] - Data Bucket/Container Name

If this parameter is specified, the command displays information only about object store configurations whose container name matches the specified value. Data ONTAP uses this container name or object store data bucket while accessing data from the object store.

[`-access-key <text>`] - Access Key ID for S3 Compatible Provider Types

If this parameter is specified, the command displays information only about AWS S3, SGWS, IBM COS object store configurations and ONTAP S3 whose access key matches the specified value. Data ONTAP requires the access key for authorized access to the object store.

[`-azure-account <text>`] - Azure Account

If this parameter is specified, the command displays information only about Azure object store configurations whose account matches the specified value. Data ONTAP requires the Azure account for

authorized access to the Azure object store.

[`-ipSpace <IPspace>`] - IPspace to Use in Order to Reach the Object Store

If this parameter is specified, the command displays information only about object store configurations whose IPspace matches the specified value. Data ONTAP uses the IPspace value to connect to the object store.

[`-iam-role <text>`] - IAM Role for AWS Cloud Volumes ONTAP

If this parameter is specified, the command displays information only about object store configurations whose IAM (Identity and Access Management) role matches the specified value.

[`-is-certificate-validation-enabled {true|false}`] - Is SSL/TLS Certificate Validation Enabled

If this parameter is specified, the command displays information only about object store configurations whose status about the validation of SSL/TLS certificate matches the specified value.

[`-use-http-proxy {true|false}`] - Use HTTP Proxy

If this parameter is specified, the command displays information only about object store configurations for which usage of HTTP proxy matches the specified value.

[`-cluster <Cluster name>`] - The Name of the Cluster to which the Configuration Belongs

If this parameter is specified, the command displays information only about object store configurations for which cluster matches the specified value.

[`-server-side-encryption {none | SSE-S3}`] - Encryption of Data at Rest by the Object Store Server

If this parameter is specified, the command displays information only about object store configurations whose server-side encryption matches the specified value.

[`-url-style {path-style | virtual-hosted-style}`] - URL Style Used to Access S3 Bucket

If this parameter is specified, the command displays information only about object store configurations whose URL style matches the specified value.

Examples

The following example displays all available object store configuration in the cluster:

```
cluster1::>storage aggregate object-store config show
```

storage aggregate object-store profiler abort

Abort Object Store Profiler

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

Description

The `storage aggregate object-store profiler abort` command will abort an ongoing object store profiler run. This command requires two parameters - an object store configuration and a node on which the

profiler is currently running.

Parameters

-node {<nodename>|local} - Node on Which the Profiler Should Run

This parameter specifies the node on which the object store profiler is running.

-object-store-name <text> - Object Store Configuration Name

This parameter specifies the object store configuration that describes the object store. The object store configuration has information about the object store server name, port, access credentials, and provider type.

Examples

The following example aborts the object store profiler :

```
cluster1::>storage aggregate object-store profiler abort -object-store
-name my-store -node my-node
```

storage aggregate object-store profiler show

Show Object Store Profiler Status

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

Description

The `storage aggregate object-store profiler show` command is used to monitor progress and results of the [storage aggregate object-store profiler start](#) command.

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>] - Node Name

This parameter specifies the node on which the profiler was started.

[-object-store-name <text>] - ONTAP Name for this Object Store Configuration

This parameter specifies the object store configuration that describes the object store. The object store configuration has information about the object store server name, port, access credentials, and provider type.

[-object-name-prefix <UUID>] - Bin UUID

This parameter specifies the object name prefix.

[-profiler-status <text>] - Profiler Status

Current status of the profiler.

[-start-time <MM/DD/YYYY HH:MM:SS>] - Profiler Start Time

Time at which profiler run started.

[-op-name <text>] - Operation Name - PUT/GET

Name of the operation. Possible values are PUT or GET.

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

Size of the PUT or GET operation.

[-op-count <integer>] - Number of Operations Performed

Number of operations issued to the object store.

[-op-failed <integer>] - Number of Operations Failed

Number of operations that failed.

[-op-latency-minimum <integer>] - Minimum Latency for Operation in Milliseconds

Minimum latency of the operation in milliseconds, as measured from the filesystem layer.

[-op-latency-maximum <integer>] - Maximum Latency for Operation in Milliseconds

Maximum latency of the operation in milliseconds, as measured from the filesystem layer.

[-op-latency-average <integer>] - Average Latency for Operation in Milliseconds

Average latency of the operation in milliseconds, as measured from the filesystem layer.

[-op-throughput {<integer>[KB|MB|GB|TB|PB] }] - Throughput per Second for the operation

Throughput per second for the operation.

[-op-errors <text>,...] - Error Reasons and Count

Error reasons and count for failed operation.

[-op-latency-histogram <text>,...] - Latency Histogram

Latency histogram for the operation.

Examples

The following example displays the results of [storage aggregate object-store profiler start](#) :

```
cluster1::>storage aggregate object-store profiler show
```

Related Links

- [storage aggregate object-store profiler start](#)

storage aggregate object-store profiler start

Start the object store profiler to measure latency and throughput

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

Description

The `storage aggregate object-store profiler start` command writes objects to an object store and reads those objects to measure latency and throughput of an object store. This command requires two parameters - an object store configuration and node from which to send the PUT/GET/DELETE operations. This command verifies whether the object store is accessible through the intercluster LIF of the node on which it runs. The command fails if the object store is not accessible. The command will create a 10GB dataset by doing 2500 PUTs for a maximum time period of 60 seconds. Then it will issue GET operations of different sizes - 4KB, 8KB, 32KB, 256KB for a maximum time period of 180 seconds. Finally it will delete the objects it created. This command can result in additional charges to your object store account. This is a CPU intensive command. It is recommended to run this command when the system is under 50% CPU utilization.

Parameters

-node {<nodename>|local} - Node on Which the Profiler Should Run

This parameter specifies the node from which PUT/GET/DELETE operations are sent.

-object-store-name <text> - Object Store Configuration Name

This parameter specifies the object store configuration that describes the object store. The object store configuration has information about the object store server name, port, access credentials, and provider type.

Examples

The following example starts the object store profiler :

```
cluster1::>storage aggregate object-store profiler start -object-store  
-name my-store -node my-node
```

storage aggregate object-store put-rate-limit modify

Modify the maximum, per node, FabricPool put rate

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

Description

The `storage aggregate object-store put-rate-limit modify` command enables the modification of values that can be used to limit the FabricPool put rate or to allow it to run more aggressively. Use the

[storage aggregate object-store put-rate-limit show](#) command to display the current maximum put rate, if any, in effect.

Parameters

-node {<nodename>|local} - Node Name

This parameter specifies the Node name.

[-default {true|false}] - Set Default Values

Setting this field to true will reset the put-rate-limit to the system default.

[-put-rate-bytes-limit {<integer>[KB|MB|GB|TB|PB]}] - Limit Max (B/s) Allowed

This field is the maximum put rate in bytes per second. The field can be used to reduce the put rate to an object store.

Related Links

- [storage aggregate object-store put-rate-limit show](#)

storage aggregate object-store put-rate-limit show

Display the per node FabricPool Put Rate Limit

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

Description

The `storage aggregate object-store put-rate-limit show` command displays the maximum put rate for FabricPool.

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

This parameter specifies the Node name.

[-put-rate-bytes-limit {<integer>[KB|MB|GB|TB|PB]}] - Put Rate Limit (B/s)

This field represents an upper boundary on the allowed put rate. The put rate will not be allowed to exceed this value but can be throttled to values less than it to accommodate client operations and a high average system CPU. Some platforms are intentionally configured with a lower maximum put rate in an effort to reduce the impact of the FabricPool put subsystem on the overall system health. A limit in effect, lower than the default, may keep cold data in the performance tier for a long period of time and may require an increase in space requirements. A value of -1 indicates no upper bound is in effect.

[`-put-rate` {<integer>[KB|MB|GB|TB|PB]}] - Estimated Put Rate (B/s)

The `put-rate` displays the maximum computed bytes per second the put subsystem assumes is possible. This put rate is measured in real-time and is in relation to the node for all attached object stores which have put traffic to them.

[`-max-put-rate-in-effect` {<integer>[KB|MB|GB|TB|PB]}] - Current Maximum Put Rate In Effect (B/s)

This field displays the maximum allowed put rate. If it is less than the `max-throughput` then it indicates that throttling is in effect. The throttling is the outcome of either staying at `put-rates-byte-limit` or an effort to accommodate client operations on a system with high CPU.

[`-max-throughput` {<integer>[KB|MB|GB|TB|PB]}] - Estimated Maximum Throughput (B/s)

This field displays the maximum put throughput rate in bytes per second that the put subsystem will throttle down from or up to. It's the greater of the `put-rate` and the node default value.

storage aggregate plex delete

Delete a plex

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

Description

The `storage aggregate plex delete` command deletes the specified plex. The aggregate specified with then `-aggregate` will be unmirrored and contain the remaining plex. The disks in the deleted plex become spare disks.

Parameters

`-aggregate` <aggregate name> - Aggregate

Name of an existing aggregate which contains the plex specified with the `-plex` parameter.

`-plex` <text> - Plex

Name of a plex which belongs to the aggregate specified with the `-aggregate` parameter.

Examples

The following example deletes `plex0` of aggregate `aggr1`:

```
cluster1::> storage aggregate plex delete -aggregate aggr1 -plex plex0
```

storage aggregate plex offline

Offline a plex

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

Description

The `storage aggregate plex offline` command takes the specified plex offline. The aggregate specified with the `-aggregate` parameter must be a mirrored aggregate and both plexes must be online. Prior to taking a plex offline, the system will flush all internally-buffered data associated with the plex and create a snapshot that is written out to both plexes. The snapshot allows for efficient resynchronization when the plex is subsequently brought back online.

Parameters

-aggregate <aggregate name> - Aggregate

Name of an existing aggregate which contains the plex specified with the `-plex` parameter.

-plex <text> - Plex

Name of a plex which belongs to the aggregate specified with the `-aggregate` parameter.

Examples

The following example takes plex0 of aggregate aggr1 offline:

```
cluster1::> storage aggregate plex offline -aggregate aggr1 -plex plex0
```

storage aggregate plex online

Online a plex

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

Description

The `storage aggregate plex online` command brings the specified plex online. The aggregate specified with the `-aggregate` parameter must be an online mirrored aggregate. The system will initiate resynchronization of the plex as part of online processing.

Parameters

-aggregate <aggregate name> - Aggregate

Name of an existing aggregate which contains the plex specified with the `-plex` parameter.

-plex <text> - Plex

Name of a plex which belongs to the aggregate specified with the `-aggregate` parameter.

Examples

The following example brings plex0 of aggregate aggr1 online:

```
cluster1::> storage aggregate plex online -aggregate aggr1 -plex plex0
```

storage aggregate plex show

Show plex details

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

Description

The `storage aggregate plex show` command displays information for the specified plex. By default, the command displays the following information about all plexes:

- Aggregate Name
- Plex Name
- Is Online
- Is Resyncing
- Resyncing Percentage
- Plex Status

To display detailed information about a single plex, use the `-aggregate` and `-plex` 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.

[-aggregate <aggregate name>] - Aggregate

Name of an existing aggregate which contains the plex specified with the `-plex` parameter.

[-plex <text>] - Plex Name

Name of a plex which belongs to the aggregate specified with the `-aggregate` parameter.

[-status <text>] - Status

Displays plex status. Possible values are:

- *normal*
- *failed*
- *empty*
- *invalid*

- *uninitialized*
- *failed assimilation*
- *limbo*
- *active*
- *inactive*
- *resyncing*

These values may appear by themselves or in combination separated by commas, for example, "*normal,active*".

[-is-online {true|false}] - Is Online

Selects the plexes that match this parameter value.

[-in-progress {true|false}] - Resync is in Progress

Selects the plexes that match this parameter value.

[-resyncing-percent <percent>] - Resyncing Percentage

Selects the plexes that match this parameter value.

[-resync-level <integer>] - Resync Level

Selects the plexes that match this parameter value.

[-pool <integer>] - Pool

Selects the plexes that match this parameter value.

Examples

The following example displays information about all the plexes for all the aggregates:

```
cluster1::> storage aggregate plex show
Aggregate Plex      Is      Is      Resyncing
                Online Resyncing Percent Status
-----
aggr0    plex0    true    false
aggr1    plex0    true    false
aggr1    plex1    true    false
aggr2    plex0    true    false
aggr2    plex2    true    false
5 entries were displayed.
```

The following example displays information about plex1 of aggregate aggr1:

```
cluster1::> storage aggregate plex show -aggregate aggr1 -plex plex1
Aggregate: aggr1
    Plex Name: plex1
    Status: normal,active
    Is Online: true
Resync is in Progress: false
Resyncing Percentage: -
    Resync Level: -
    Pool: 1
```

storage aggregate reallocation quiesce

Quiesce reallocate job on aggregate

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 [storage aggregate reallocation restart](#) command to restart a job that is quiesced.

There is no limit to how long a job can remain in the quiesced (paused) state.

Parameters

-aggregate <aggregate name> - Aggregate Name

Specifies the aggregate on which you want to temporarily pause the job.

Examples

```
cluster1::> storage aggregate reallocation quiesce
    -aggregate aggr0
```

Temporarily stops (pauses) any reallocation job running on aggregate aggr0.

Related Links

- [storage aggregate reallocation restart](#)

storage aggregate reallocation restart

Restart reallocate job on aggregate

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

Description

Starts a reallocation job. Use this command to restart a quiesced (temporarily stopped) job or a scheduled scan that is idle for the aggregate.

Parameters

-aggregate <aggregate name> - Aggregate Name

Specifies the aggregate 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 was stopped. If you specify this parameter without a value, it is set to true and the scan restarts at the beginning.

Examples

```
cluster1::> storage aggregate reallocation restart
             -aggregate aggr0 -ignore-checkpoint true
```

Restarts reallocation job on aggregate aggr0 from the beginning.

storage aggregate reallocation schedule

Modify schedule of reallocate job on aggregate

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 [storage aggregate 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 [storage aggregate reallocation start](#) command.

Parameters

-aggregate <aggregate name> - Aggregate Name

Specifies the aggregate on which you want to schedule reallocation jobs.

[-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::> storage aggregate reallocation schedule -aggregate aggr0 -cron  
"0 23 6 *"
```

Schedules a reallocation job to run at 11:00 p.m. every Saturday on aggr0.

Related Links

- [storage aggregate reallocation start](#)

storage aggregate reallocation show

Show reallocate job status for improving free space layout

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, aggregate and scan id. If you do not specify the *id* for a particular reallocation scan, the command displays information about all the existing reallocation scans.

Parameters

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

Displays the value of relevant field that you specify for the reallocation scans that are present.

| [-v]

Specify this parameter to display the output in a verbose format.

[`-instance`] }

Displays information about reallocation scans on aggregates in a list format.

[`-id` <integer>] - Job ID

Specify this parameter to display the reallocation scan that matches the reallocation job ID that you specify.

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

Specify this parameter to display the reallocation scan that matches the aggregate 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 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::> storage aggregate reallocation show
Job ID          Aggregate          Schedule          State
-----          -
23              aggr0              reallocate_0 23 * 6  Queued
```

Displays the job ID, aggregate, schedule, and state for the reallocation scans.

Related Links

- [job schedule show](#)

storage aggregate reallocation start

Start reallocate job on aggregate

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

Description

Begins a reallocation scan on a specified aggregate.

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 aggregate.

You can define the reallocation scan job so that it runs at a specific interval, or you can use the [storage aggregate reallocation schedule](#) command to schedule reallocation jobs.

Parameters

-aggregate <aggregate name> - Aggregate Name

Specify this parameter to specify the target aggregate on which to start a reallocation scan.

{ [-i, -interval <text>] - Interval Schedule

Specified the schedule in a single string with four fields:

- minute is a value from 0 to 59.
- hour is a value from 0 (midnight) to 23 (11:00 p.m.).
- day of month is a value from 1 to 31.



If you specify 31 as the value for the day of the month, reallocation scans will not run in any of the months with fewer than 31 days.

- day of the week is a value from 0 (Sunday) to 6 (Saturday).

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.

| [-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.

Examples

```
cluster1:> storage aggregate reallocation start -aggregate aggr0
-interval "0 23 * 6"
```

Starts a reallocation job on aggregate `aggr0` at 11:00 p.m. every Saturday.

Related Links

- [storage aggregate reallocation schedule](#)

storage aggregate reallocation stop

Stop reallocate job on aggregate

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

Description

Stops and deletes any reallocation scan running on the specified aggregate. This command stops and deletes in-progress, scheduled, and quiesced scans.

Parameters

-aggregate <aggregate name> - Aggregate Name

Specify this parameter to specify the target aggregate on which to stop and delete a reallocation scan.

Examples

```
cluster1::> storage aggregate reallocation stop -aggregate aggr0
```

Stops and deletes the reallocation scan on aggregate `aggr0`.

storage aggregate relocation show

Display relocation status of an aggregate

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

Description

The `storage aggregate relocation show` command displays status of aggregates which were relocated in the last instance of relocation operation.

Parameters

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

If you specify the `-fields <fieldname>, ...` parameter, the command only displays the fields that you specify.

| [-instance] }

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

[-node {<nodename>|local}] - Node

Selects aggregates from the specified source node.

[-aggregate <text>] - Aggregate Name

Selects the aggregates that match this parameter value.

[-relocation-status <text>] - Aggregates Relocation Status

Selects the aggregates whose relocation status matches this parameter value.

[-destination <text>] - Destination for Relocation

Selects the aggregates that are designated for relocation on the specified destination node.

Examples

The following example displays the relocation status of aggregates on all nodes in the cluster:

```

cluster1::> storage aggregate relocation show
Source          Aggregate  Destination  Relocation Status
-----
node0
                -          -              Not attempted yet
node1
                aggr1     node0         Done
                aggr2     node0         In progress
                aggr3     node0         Not attempted yet
4 entries were displayed.

```

storage aggregate relocation start

Relocate aggregates to the specified destination

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

Description

The `storage aggregate relocation start` command initiates the relocation of the aggregates from one node to the partner node in a high-availability (HA) pair.

Parameters

-node {<nodename>|local} - Name of the Node that currently owns the aggregate

This specifies the source node where the aggregates to be relocated reside.

-destination {<nodename>|local} - Destination node

This specifies the destination node where aggregates are to be relocated.

-aggregate-list <aggregate name>,... - List of Aggregates to be relocated

This specifies the list of aggregate names to be relocated from source node to destination node.

[-override-vetoes {true|false}] - Override Vetoes

This specifies whether to override the veto checks for relocation operation. Initiating aggregate relocation with vetoes overridden will result in relocation proceeding even if the node detects outstanding issues that would make aggregate relocation dangerous or disruptive. The default value is false.

[-relocate-to-higher-version {true|false}] - Relocate To Higher Version

This specifies if the aggregates are to be relocated to a node which is running on a higher version of Data ONTAP than the source node. If an aggregate is relocated to this destination then that aggregate cannot be relocated back to the source node till the source is also upgraded to the same or higher Data ONTAP version. This option is not required if the destination node is running on higher minor version, but the same major version. The default value is false.

[-override-destination-checks {true|false}] - Override Destination Checks

This specifies if the relocation operation should override the check done on destination node. This option could be used to force a relocation of aggregates even if the destination has outstanding issues. Note that this could make the relocation dangerous or disruptive. The default value is false.

[-ndo-controller-upgrade {true|false}] - Relocate Aggregates for NDO Controller Upgrade

This specifies if the relocation operation is being done as a part of non-disruptive controller upgrade process. Aggregate relocation will not change the home ownerships of the aggregates while relocating as part of controller upgrade. The default value is false.

Examples

The following example relocates aggregates name aggr1 and aggr2 from source node node0 to destination node node1:

```
cluster1::> storage aggregate relocation start -node node0 -destination
node1 -aggregate-list aggr1, aggr2
```

storage aggregate resynchronization modify

Modify aggregate resynchronization priorities

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

Description

The `storage aggregate resynchronization modify` command can be used to modify the resynchronization priority of an aggregate.

When the number of aggregates pending resynchronization is higher than the maximum number of concurrent resynchronization operations allowed on a node, the aggregates get resynchronized in the order of their "resync-priority" values.

For example, let the `max-concurrent-resync` under the `storage aggregate resynchronization`

`options` directory for a node be set to two. If there are three aggregates waiting to be resynchronized, where their respective `resync-priority` values are `high`, `medium`, and `low`, then the third aggregate is not allowed to start resynchronization until one of the first two aggregates has completed resynchronizing.

Parameters

-aggregate <aggregate name> - Aggregate

This parameter specifies the aggregate that is to be modified.

[-resync-priority {high(fixed)|high|medium|low}] - Resynchronization Priority

This parameter specifies the new resynchronization priority value for the specified aggregate. This field cannot be modified for unmirrored or Data ONTAP system aggregates.

Possible values for this parameter are:

- `high`: Mirrored data aggregates with this priority value start resynchronization first.
- `medium`: Mirrored data aggregates with this priority value start resynchronization after all the system aggregates and data aggregates with 'high' priority value have started resynchronization.
- `low`: Mirrored data aggregates with this priority value start resynchronization only after all the other aggregates have started resynchronization.

Examples

The following example changes the `resync-priority` of a specified aggregate to `medium`:

```
cluster1::> storage aggregate resynchronization modify -aggregate aggr1
-resync-priority medium
```

storage aggregate resynchronization show

Display aggregate resynchronization priorities

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

Description

The `storage aggregate resynchronization show` command displays the relative resynchronization priority for each aggregate in the cluster. When a particular node restricts how many resync operations can be active concurrently, these priorities are used to prioritize the operations. The maximum concurrent resync operations for a node is displayed in the [storage aggregate resynchronization options show](#) command. If no parameters are specified, the command displays the following information about all the aggregates in the cluster:

- Aggregate name
- Node that owns the aggregate
- Resync priority for the aggregate

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.

[-aggregate <aggregate name>] - Aggregate

If this parameter is specified, the command displays the resynchronization priority only for the specified aggregate.

[-node {<nodename>|local}] - Node

If this parameter is specified, the command displays the resynchronization priority only for the aggregates owned by the specified node.

[-resync-priority {high(fixed)|high|medium|low}] - Resynchronization Priority

If this parameter is specified, the command displays only the resynchronization priority that matches the specified value. Possible values for this parameter are:

- `high(fixed)`: This value is reserved for Data ONTAP system aggregates, which cannot have any other value for this field. These aggregates always start their resynchronization operation at the first available opportunity. This value cannot be assigned to a data aggregate.
- `high`: Mirrored data aggregates with this priority value start resynchronization first.
- `medium`: Mirrored data aggregates with this priority value start resynchronization after all the system aggregates and data aggregates with 'high' priority value have started resynchronization.
- `low`: Mirrored data aggregates with this priority value start resynchronization only after all the other aggregates have started resynchronization.

When the number of aggregates waiting for resynchronization is higher than the maximum number of resynchronization operations allowed on a node, then the `resync-priority` field is used to determine which aggregate starts resynchronization first. This field is not set for unmirrored aggregates.

Examples

The following command displays the resynchronization priorities for all the aggregates in the cluster:

```
cluster1::> storage aggregate resynchronization show
Aggregate Node           Resync Priority
-----
aggr0_n1 cluster1-01      high(fixed)
aggr0_n2 cluster1-02      high(fixed)
aggr1   cluster1-01      low
aggr2   cluster1-01      high
aggr3   cluster1-01      medium
4 entries were displayed.
```

Related Links

- [storage aggregate resynchronization options show](#)

storage aggregate resynchronization options modify

Modify node specific aggregate resynchronization options

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

Description

The `storage aggregate resynchronization options modify` command can be used to modify the options that govern the resynchronization of aggregates on a given cluster node.

Modifying the `max-concurrent-resyncs` option changes the number of aggregates that are allowed to resynchronize concurrently. When the number of aggregates waiting for resynchronization is higher than this value, the aggregates are resynchronized in the order of their "`resync-priority``". This value can be modified using the [storage aggregate resynchronization modify](#) command while specifying the `-resync-priority` parameter.

Parameters

-node {<nodename>|local} - Node

This parameter specifies the node for which the option is to be modified.

[-max-concurrent-resync <integer>] - Maximum Concurrent Resynchronizing Aggregates

This parameter specifies the new value for the maximum number of concurrent resync operations allowed on a node. This option must be specified along with the `-node` parameter. When a node has active resync operations, setting this parameter to a value that is lower than the number of currently resyncing aggregates will trigger a user confirmation.

Examples

The following example changes the maximum concurrent resync operations for the specified node to `ten` :

```
cluster1::> storage aggregate resynchronization options modify -node node1
-max-concurrent-resyncs 10
```

Related Links

- [storage aggregate resynchronization modify](#)

storage aggregate resynchronization options show

Display node specific aggregate resynchronization options

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

Description

The `storage aggregate resynchronization options show` command displays all the options that govern the resynchronization of aggregates on a given cluster node. If no parameters are specified, the command displays the following information about all nodes:

- Node for which the information is being displayed
- Maximum number of concurrent resynchronizing aggregates allowed

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

If this parameter is specified, the command displays resynchronization options only for the specified node.

[-max-concurrent-resync <integer>] - Maximum Concurrent Resynchronizing Aggregates

If this parameter is specified, the command displays only the resynchronization option that matches the specified value.

Examples

The following example displays the maximum number of concurrent resyncs allowed for each node in the cluster:

```
cluster1::> storage aggregate resynchronization options show
Node           Maximum Concurrent Resynchronizing Aggregates
-----
cluster1-01    15
cluster1-02     4
2 entries were displayed.
```

The following example displays the maximum number of concurrent resyncs allowed for a specified node:

```
cluster1::> storage aggregate resynchronization options show -node node1
Node           Maximum Concurrent Resynchronizing Aggregates
-----
cluster1-01    15
```

The following example displays all the nodes that allow more than five concurrent resync operations:

```
cluster1::> storage aggregate resynchronization options show -max  
-concurrent-resyncs >5
```

Node	Maximum Concurrent Resynchronizing Aggregates
-----	-----
cluster1-01	15

Copyright information

Copyright © 2023 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.