



storage commands

ONTAP 9.8 commands

NetApp
February 11, 2024

Table of Contents

- storage commands 1
 - storage aggregate commands 1
 - storage array commands 107
 - storage automated-working-set-analyzer commands 123
 - storage disk commands 129
 - storage dqp commands 181
 - storage encryption commands 183
 - storage errors commands 196
 - storage failover commands 198
 - storage firmware commands 240
 - storage iscsi-initiator commands 247
 - storage path commands 252
 - storage pool commands 262
 - storage port commands 275
 - storage raid-options commands 285
 - storage shelf commands 287
 - storage tape commands 341

storage commands

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 | 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`.

`[-chksumstyle <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 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:

| New Node | Total Aggrs | 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:

| Is Node | Usable New Data Aggregate | -Devices To Use- Mirrored | Size | Disks |
|---------|---------------------------|---------------------------|--------|-------|
| node1 | node1_SSD_1 | false | 3.66TB | 5 |
| 48 | | | | |
| node2 | node2_SSD_1 | false | 3.66TB | 5 |
| 48 | | | | |

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

| RAID Group In New | Disk | Usable Disk Or |
|------------------------------|------|---------------------|
| ---Count--- | | |
| Data Aggregate To Be Created | Type | Size Partition Data |
| Parity | | |
| ----- | | |

```

-----
/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 Type | Disk Or Usable | Remaining Size Partition | Spares |
|--------------|----------------|-------------------|-----------------------------|--------|
| node1 | SSD | 185.5GB | disk | 1 |
| node2 | SSD | 185.5GB | disk | 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 | 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 Type`

This parameter specifies the type for RAID groups on the aggregate. The values are *raid4* for RAID4, *raid_dp* for RAID Double Parity, and *raid_tec* for RAID Triple-Erasure-Code. The default setting is *raid_dp* unless the disks are HDDs with a capacity larger than 4 TB, in which case the default will be *raid_tec*. This parameter is not needed for array LUNs because they are always created with the *raid0* raidtype. *raid4* is 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 (privilege: advanced)`

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

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

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 (%) (privilege: advanced)`

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 (%) (privilege: advanced)`

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 (privilege: advanced)

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 | VMLUN | VMLUN-SSD}] - Disk Type

This parameter specifies either the HDD tier or the SSD tier when changing the RAID type of a Flash Pool. 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.

[-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.

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 Type`

This parameter specifies the RAID type for RAID groups on the aggregate. The possible values are `raid4` for RAID4, `raid_dp` for RAID-DP, and `raid_tec` for RAID-TEC. If you change the RAID type from RAID4 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 to RAID-TEC or vice-versa is not supported. To change the RAID type from RAID4 to RAID-TEC, first change from RAID4 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 (privilege: advanced)

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 (%) (privilege: advanced)

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 (%) (privilege: advanced)

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 nvlogged 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.

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 (privilege: advanced)

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

-nodename {<nodename>|local} - Node Name (privilege: advanced)

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
Aggregate                               Current          Target
Progress                               Usable Size      Usable Size      Provisioning
-----
node1_SSD_1                           0B              23.65GB Creating
node2_SSD_1                           0B              23.65GB Creating

cluster1::> storage aggregate show-auto-provision-progress
Aggregate                               Current          Target
Progress                               Usable Size      Usable Size      Provisioning
-----
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] (privilege: 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.

[-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 (privilege: advanced)

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 (privilege: advanced)

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 {<integer>[KB|MB|GB|TB|PB]}] - Total Data Reduction Logical Used (privilege: advanced)

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

[-total-data-reduction-physical-used {<integer>[KB|MB|GB|TB|PB]}] - Total Data Reduction Physical Used (privilege: advanced)

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

[-total-data-reduction-efficiency-ratio <text>] - Total Data Reduction Efficiency Ratio

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

[`-volume-logical-used` {<integer>[KB|MB|GB|TB|PB]}] - Logical Space Used for All volumes (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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
-----
      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] (privilege: 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 (privilege: advanced)

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 (privilege: advanced)

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 {<integer>[KB|MB|GB|TB|PB]}`] - Total Data Reduction Logical Used (privilege: advanced)

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

[`-total-data-reduction-physical-used {<integer>[KB|MB|GB|TB|PB]}`] - Total Data Reduction Physical Used (privilege: advanced)

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

[`-total-data-reduction-efficiency-ratio <text>`] - Total Data Reduction Efficiency Ratio

Displays the total storage efficiency ratio obtained by Deduplication, Compression, Data Compaction, Pattern Detection and FlexClone data reduction technologies 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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

Displays the logical size used by the aggregate.

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

Displays the physical size used by the aggregate.

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

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 (privilege: advanced)

Displays the storage efficiency ratio of the aggregate.

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

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 Percentage |
|-----------|----------------|---------------------|
| 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 accomodate 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.

`[-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 Referenced Capacity

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

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

If this parameter is specified, the command displays information only about aggregates whose 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]}] - Object Store Metadata

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>] - Object Store Metadata Percent

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]}] - Total Unreclaimed Space

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>] - Object Store Unreclaimed Space Percentage

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.

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`] (privilege: advanced)

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 | 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 | 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 | 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`
(privilege: advanced)**

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`
(privilege: advanced)**

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 (privilege: advanced)

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 auto-scaled units.

`[-local-usable-data2-size {<integer>[KB|MB|GB|TB|PB]}]` - Local Node Data2 Usable Size (privilege: advanced)

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 auto-scaled units.

Examples

Display spare disks owned by node *node-b*.

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

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 | zeroed | | | | |

| | | | | | |
|---------|--------|----------|------|-------|---------|
| 1.1.15 | BSAS | capacity | 7200 | block | 413.2GB |
| 414.0GB | zeroed | | | | |

Original Owner: node-b

Pool0

Partitioned Spares

Local

Local

Data

Root Physical

| Disk | Type | Class | RPM | Checksum | Usable |
|------|------|-------|-----|----------|--------|
|------|------|-------|-----|----------|--------|

| Usable | Size | Status | | | |
|--------|------|--------|--|--|--|
|--------|------|--------|--|--|--|

| | | | | | |
|---------|---------|-------------|-------|-------|---------|
| 1.0.8 | SAS | performance | 10000 | block | 472.9GB |
| 73.89GB | 547.1GB | zeroed | | | |

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 (privilege: advanced)

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, raid0, raid_dp, raid_tec, redirect, and wafl inconsistent. You can specify multiple values (for example, reconstruct and growing).

[-t, -raidtype {raid_tec|raid_dp|raid4}] - 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, *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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (%) (privilege: advanced)

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 (%) (privilege: advanced)

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.

`[-azcs-read-optimization-enabled {true|false}]` - azcs-read-optimization Enabled

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

`[-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.

Examples

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

```
cluster1::> storage aggregate show
```

| Aggregate Status | Size | Available | Used% | State | #Vols | Nodes | RAID |
|------------------------------------|---------|-----------|-------|--------|-------|-------------|--------|
| aggr0 raid_dp, | 6.21TB | 1.78TB | 71% | online | 49 | cluster1-01 | normal |
| aggr1 raid_dp, mirrored, | 56.04MB | 55.89MB | 0% | online | 0 | cluster1-02 | normal |
| aggr2 raid_dp, | 1.77TB | 1.63TB | 8% | online | 1 | cluster1-01 | normal |
| aggr3 raid_dp, | 1.77TB | 1.73TB | 2% | online | 2 | cluster1-02 | 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 |
|-----------|------|-----------|-------|-------|-------|-------|------|
|-----------|------|-----------|-------|-------|-------|-------|------|

| Status | | | | | | | |
|--------|--|--|--|--|--|--|--|
|--------|--|--|--|--|--|--|--|

| | | | | | | | |
|-------|---|---|---|----------------|--|--|--|
| 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 |
|-----------|------|-----------|-------|-------|-------|-------|------|
|-----------|------|-----------|-------|-------|-------|-------|------|

| Status | | | | | | | |
|--------|--|--|--|--|--|--|--|
|--------|--|--|--|--|--|--|--|

| | | | | | | | |
|-------|---|---|---|----------------|--|--|--|
| aggr2 | - | - | - | remote_cluster | | | |
|-------|---|---|---|----------------|--|--|--|

| | | |
|---|---|---|
| - | - | - |
|---|---|---|

| |
|---|
| - |
|---|

| | | | | | | | |
|-------|---|---|---|----------------|--|--|--|
| aggr3 | - | - | - | remote_cluster | | | |
|-------|---|---|---|----------------|--|--|--|

| | | |
|---|---|---|
| - | - | - |
|---|---|---|

| |
|---|
| - |
|---|

```
cluster1:
```

| Aggregate | Size | Available | Used% | State | #Vols | Nodes | RAID |
|-----------|------|-----------|-------|-------|-------|-------|------|
|-----------|------|-----------|-------|-------|-------|-------|------|

| Status | | | | | | | |
|--------|--|--|--|--|--|--|--|
|--------|--|--|--|--|--|--|--|

| | | | | | | | |
|-------|--------|--------|-----|--------|---|-------------|--|
| 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-vsimg2

Has Cross Volume Deduplication Savings:           false
Cross Volume Background Deduplication:             false
Cross Volume Inline Deduplication:                 false

Aggregate: aggr1
  Node: vivek6-vsimg2

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 cross-volume-dedupe revert-to` command is used to revert cross volume deduplication savings on an aggregate.

Parameters

-aggregate <aggregate name> - Aggregate (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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 (privilege: advanced)

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
(privilege: advanced)**

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 (privilege: advanced)

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
(privilege: advanced)**

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
(privilege: advanced)**

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 reassigned 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 (privilege: advanced)`

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-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 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 (privilege: advanced)

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 (privilege: advanced)`

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 (privilege: advanced)`

This parameter specifies the node on which the object store profiler is running.

-object-store-name <text> - Object Store Configuration Name (privilege: advanced)

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 (privilege: advanced)

This parameter specifies the node on which the profiler was started.

[-object-store-name <text>] - ONTAP Name for this Object Store Configuration (privilege: advanced)

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.

[-profiler-status <text>] - Profiler Status (privilege: advanced)

Current status of the profiler.

[-start-time <MM/DD/YYYY HH:MM:SS>] - Profiler Start Time (privilege: advanced)

Time at which profiler run started.

[`-op-name <text>`] - Operation Name - PUT/GET (privilege: advanced)

Name of the operation. Possible values are PUT or GET.

[`-op-size {<integer>[KB|MB|GB|TB|PB]}`] - Size of Operation (privilege: advanced)

Size of the PUT or GET operation.

[`-op-count <integer>`] - Number of Operations Performed (privilege: advanced)

Number of operations issued to the object store.

[`-op-failed <integer>`] - Number of Operations Failed (privilege: advanced)

Number of operations that failed.

[`-op-latency-minimum <integer>`] - Minimum Latency for Operation in Milliseconds (privilege: advanced)

Minimum latency of the operation in milliseconds, as measured from the filesystem layer.

[`-op-latency-maximum <integer>`] - Maximum Latency for Operation in Milliseconds (privilege: advanced)

Maximum latency of the operation in milliseconds, as measured from the filesystem layer.

[`-op-latency-average <integer>`] - Average Latency for Operation in Milliseconds (privilege: advanced)

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 (privilege: advanced)

Throughput per second for the operation.

[`-op-errors <text>,...`] - Error Reasons and Count (privilege: advanced)

Error reasons and count for failed operation.

[`-op-latency-histogram <text>,...`] - Latency Histogram (privilege: advanced)

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 (privilege: advanced)

This parameter specifies the node from which PUT/GET/DELETE operations are sent.

-object-store-name <text> - Object Store Configuration Name (privilege: advanced)

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 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 Online | Is Resyncing | Resyncing Percent | Status |
|-----------|-------|-----------|--------------|-------------------|---------------|
| aggr0 | plex0 | true | false | - | normal,active |
| aggr1 | plex0 | true | false | - | normal,active |
| aggr1 | plex1 | true | false | - | normal,active |
| aggr2 | plex0 | true | false | - | normal,active |
| aggr2 | plex2 | true | false | - | normal,active |

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
                -          -
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 (privilege: advanced)

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
```

storage array commands

storage array modify

Make changes to an array's profile.

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

Description

The `storage array modify` command lets the user change several array parameters.

Parameters

-name <text> - Name

Storage array name, either generated by Data ONTAP or assigned by the user.

[-prefix <text>] - Prefix

Abbreviation for the named array.

[-vendor <text>] - Vendor

Array manufacturer.

[-model <text>] - Model

Array model number.

[-options <text>] - options

Vendor specific array settings.

[-max-queue-depth <integer>] - Target Port Queue Depth (privilege: advanced)

The target port queue depth for all target ports on this array.

[-lun-queue-depth <integer>] - LUN Queue Depth (privilege: advanced)

The queue depth assigned to array LUNs from this array.

{ [-is-upgrade-pending {true|false}] - Upgrade Pending (privilege: advanced)

Set this parameter to *true* if the array requires additional Data ONTAP resilience for a pending firmware upgrade. Keep this parameter *false* during normal array operation. This value can not be set to *true* if `-path-failover-time` is greater than zero.

[`-path-failover-time` <integer>] - Path Failover Time (sec)

The time delay (in secs) before switching the I/O path when the path is deleted. The maximum time delay is 30 sec. The default is 0. This value can not be greater than zero if `-is-upgrade-pending` is *true*.

[`-all-path-fail-delay` <integer>] - Extend All Path Failure Event (secs)

Use this parameter to increase the delay before Data ONTAP declares an "all path failure" event for an array. Delaying the "all path failure" event allows Data ONTAP to suspend I/O operations for a longer period of time before declaring a data access disruption, allowing for I/O operations to resume if any path comes back online within the specified duration. A valid delay is any value between 30 and 90 seconds. A value of 0 will reset the delay, resulting in default actions being taken whenever an "all path failure" event is detected.

Examples

This command changes the model to FastT.

```
cluster1::> storage array modify -name IBM_1722_1 -model FastT
```

storage array remove

Remove a storage array record from the array profile database.

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

Description

The `storage array remove` command discards array profile records for a particular storage array from the cluster database. The command fails if a storage array is still connected to the cluster. Use the [storage array config show](#) command to view the array connectivity status. The array target port can be removed using the [storage array port remove](#) command.

Parameters

`-name` <text> - Name

Name of the storage array you want to remove from the database.

Examples

```
cluster1::> storage array remove IBM_1722_1
```

Related Links

- [storage array config show](#)
- [storage array port remove](#)

storage array rename

Change the name of a storage array in the array profile database.

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

Description

The `storage array rename` command permits substitution of the array profile name which Data ONTAP assigned during device discovery. By default, the name that Data ONTAP assigned to the storage array during discovery is shown in Data ONTAP displays and command output.

Parameters

-name <text> - Name

Storage array name either generated by Data ONTAP or assigned by the user.

-new-name <text> - The new name to assign to this array profile. (28 chars max)

New name to assign to the storage array.

Examples

```
cluster1::> storage array rename -name HITACHI_DF600F_1 -new-name MyArray
```

storage array show

Display information about SAN-attached storage arrays.

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

Description

The `storage array show` command displays information about arrays visible to the cluster. If no parameters are specified, the command displays the following information about all storage arrays:

- Prefix
- Name
- Vendor
- Model
- Options

To display detailed information about a single array, use the ``-name`` parameter. The detailed view adds the following information:

- Serial Number
- Optimization Policy
- Affinity
- Errors
- Path Failover Time

- Extend All Path Failure Event

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.

[-name <text>] - Name

Selects the arrays that match this parameter value.

[-prefix <text>] - Prefix

Abbreviation for the named array.

[-vendor <text>] - Vendor

Array manufacturer.

[-model <text>] - Model

Array model number.

[-options <text>] - options

Vendor specific array settings.

[-serial-number <text>] - Serial Number

Array product identifier.

[-max-queue-depth <integer>] - Target Port Queue Depth (privilege: advanced)

Selects the arrays that match this parameter value.

[-lun-queue-depth <integer>] - LUN Queue Depth (privilege: advanced)

Selects the arrays that match this parameter value.

[-optimization-policy {iALUA|eALUA|symmetric|proprietary|mixed|unknown}] - Optimization Policy

Selects the arrays that match this parameter value.

[-affinity {none|aaa|ap|mixed|unknown}] - Affinity

Selects the arrays that match this parameter value.

[-error-text <text>,...] - Error Text

Selects the arrays that match this parameter value.

[-is-upgrade-pending {true|false}] - Upgrade Pending (privilege: advanced)

Selects the arrays that match this parameter value.

[-path-failover-time <integer>] - Path Failover Time (sec)

Use this parameter to list arrays that have path failover time set to the value you specify.

[-all-path-fail-delay <integer>] - Extend All Path Failure Event (secs)

Use this parameter to list arrays that have the all path failure event delay set to the value you specify.

Examples

The following example displays information about all arrays.

```
cluster1::> storage array show
Prefix          Name      Vendor      Model Options
-----
                HITACHI_DF600F_1  HITACHI      DF600F
                IBM_1722_1        IBM          1722
2 entries were displayed.
```

The following example displays detailed information about a specific array:

```
cluster1::> storage array show -name HITACHI_DF600F_1
Name: HITACHI_DF600F_1
                Prefix: abc
                Vendor: HITACHI
                Model: DF600F
                options:
                Serial Number: 4291000000000000
                Optimization Policy: iALUA
                Affinity: aaa
                Error Text:
                Path Failover Timeout (sec): 30
                Extend All Path Failure Event (secs): 50
```

storage array config show

Display connectivity to back-end storage arrays.

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

Description

The `storage array config show` command displays information about how the storage arrays connect to the cluster, LUN groups, number of LUNS, and more. Use this command to validate the configuration and to assist in troubleshooting.

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.

| [-switch]

If you specify this parameter, switch port information is shown.

| [-instance] }

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

[-node {<nodename>|local}] - Controller Name

Selects the arrays that match this parameter value.

[-group <integer>] - LUN Group

Selects the arrays that match this parameter value. A LUN group is a set of LUNs that shares the same path set.

[-target-wwpn <text>] - Array Target Ports

Selects the arrays that match this parameter value (the World Wide Port Name of a storage array port).

[-initiator <text>] - Initiator

Selects the arrays that match this parameter value (the host bus adapter that the clustered node uses to connect to storage arrays).

[-array-name <array name>] - Array Name

Selects the arrays that match this parameter value.

[-target-side-switch-port <text>] - Target Side Switch Port

Selects the arrays that match this parameter value.

[-initiator-side-switch-port <text>] - Initiator Side Switch Port

Selects the arrays that match this parameter value.

[-lun-count <integer>] - Number of array LUNs

Selects the arrays that match this parameter value.

[-ownership {all|assigned|unassigned}] - Ownership

Selects the arrays that match this parameter value.

Examples

```
cluster1::> storage array config show
          LUN    LUN
Node      Group Count          Array Name      Array Target
Port Initiator
-----
-----
```

| | | | |
|------------------|---|----|----------------------|
| vnv3070f19a | 0 | 20 | DGC_RAID5_1 |
| 5006016030229f13 | | 0d | |
| 5006016130229f13 | | 0c | |
| 5006016830229f13 | | 0b | |
| 5006016930229f13 | | 0a | |
| | 1 | 21 | HITACHI_OPEN_1 |
| 50060e80034fe704 | | 0c | |
| 0d | | | |
| 50060e80034fe714 | | 0a | |
| 0b | | | |
| 50060e80034fe715 | | 0b | |
| 50060e80034fe716 | | 0c | |
| 0d | | | |
| | 2 | 8 | EMC_SYMMETRIX_1 |
| 50060482cb1bce1d | | 0a | |
| 0b | | | |
| 5006048acb1bce0c | | 0c | |
| 0d | | | |
| | 3 | 10 | IBM_UniversalXport_1 |
| 202600a0b8322d10 | | 0c | |
| 0d | | | |
| 204700a0b8322d10 | | 0a | |
| 0b | | | |
| vnv3070f19b | 0 | 20 | DGC_RAID5_1 |
| 5006016030229f13 | | 0d | |
| 5006016130229f13 | | 0c | |
| 5006016830229f13 | | 0b | |
| 5006016930229f13 | | 0a | |

```

1      21      HITACHI_OPEN_1
50060e80034fe704      0c
0d
50060e80034fe714      0a
0b
50060e80034fe715      0b
50060e80034fe716      0c
0d
2      8      EMC_SYMMETRIX_1
50060482cb1bce1d      0a
0b
5006048acb1bce0c      0c
0d
3      10      IBM_UniversalXport_1
202600a0b8322d10      0c
0d
204700a0b8322d10      0a
0b
38 entries were displayed.
```

Warning: Configuration errors were detected. Use 'storage errors show' for detailed information.

storage array disk paths show

Display a list of LUNs on the given array

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

Description

The `storage array disk paths show` command displays information about disks and array LUNs. Where it appears in the remainder of this document, "disk" may refer to either a disk or an array LUN. By default, the command displays the following information about all disks:

- Disk Unique Identifier
- Controller name
- Initiator Port
- LUN ID
- Failover optimization type
- The Use State of the LUN on this path
- Target Port
- Target IQN
- TPGN
- Port speeds
- Kbytes/sec on Disk (Rolling Average)
- Number IOPS per second on disk (Rolling Average)

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

Parameters

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

Displays the specified fields for all disks, in column style output.

| [-switch]

Displays the switch port information for all disks, in column style output.

| [-instance] }

Displays detailed disk information. If no disk path name is specified, this parameter displays the same detailed information for all disks as does the `-disk` parameter. If a disk path name is specified, then this parameter displays the same detailed information for the specified disks as does the `-disk` parameter.

[-uid <text>] - Disk Unique Identifier

Selects the disks whose unique id matches this parameter value. A disk unique identifier has the form:

```
`20000000:875D4C32:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000`
```

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

Displays detailed information about the specified disks.

[-array-name <array name>] - Array Name

Selects information about the LUNs presented by the specified storage array.

[-diskpathnames <disk path name>,...] - Path-Based Disk Names

Selects information about disks that have all of the specified path names.

[-nodelist {<nodename>|local}] - Controller name

Selects information about disks that are visible to all of the specified nodes .

[-initiator <text>,...] - Initiator Port

Selects information about disks that are visible to the initiator specified. Disks that are not currently in use by that initiator are included.

[-lun <integer>,...] - LUN ID

Selects information about the specified LUNs.

[-target-wwpn <text>,...] - Target Port

Selects information about disks that are visible on target ports identified by their World Wide Port Name.

[-initiator-side-switch-port <text>,...] - Initiator Side Switch Port

Selects information about disks visible to an initiator that is connected to the specified switch port.

[-lun-path-use-state <text>,...] - The Use State of the LUN on this path

Selects information about LUNs reporting the specified in-use state.

[-tpgn <integer>,...] - Target Port Group Number

Selects information about disks that belong to the specified Target Port Group Number.

[-port-speed <text>,...] - Port Speed

Selects information about disks served by a Host Bus Adapter that is running at the specified port speed.

[-lun-io-kbps <integer>,...] - Kbytes/sec on Disk (Rolling Average)

Selects information about the LUNs that have reached the specified I/O throughput.

[-lun-iops <integer>,...] - Number IOPS per second on disk (Rolling Average)

Selects information about the LUNs that have reached the specified number of IOPs.

[-target-side-switch-port <text>,...] - Target Side Switch Port

Selects information about disks that are visible on target ports identified by the switch port to which they are connected.

[-target-port-access-state <text>,...] - Failover optimization type

Selects information about disks visible on target ports that have the specified access state.

[-initiator-io-kbps <integer>,...] - Kbytes of I/O per second on Initiator (Rolling Average)

Selects information about disks visible to an initiator that has executed I/O at the specified throughput.

[-initiator-iops <integer>,...] - Number of IOPS on Initiator (Rolling Average)

Selects information about disks visible to an initiator that has executed the specified number of IOPs.

[-target-io-kbps <integer>,...] - Kbytes of I/O per second to Target (Rolling Average)

Selects information about disks visible on target ports that have reached the specified I/O throughput.

[-target-iops <integer>,...] - Number of IOPS to Target (Rolling Average)

Selects information about disks visible on target ports that have performed the specified number of IOPs.

[-path-link-errors <integer>,...] - Link Error count on path

Selects information about disks with paths that have incurred the specified number of FC link errors.

[-path-io-kbps <integer>,...] - Kbytes of I/O per second on Path (Rolling Average)

Selects information about disk with paths that have reached the specified I/O throughput.

[-path-iops <integer>,...] - Number of IOPS on Path (Rolling Average)

Selects information about disks on those paths that have reached the specified number of IOPs.

[-path-quality <integer>,...] - Percentage of weighted error threshold

Selects information about disks on paths that have incurred the specified number of errors. The value displayed is a measure of the health of a path expressed as a percentage of an error threshold. Once a path has reached or surpassed the error threshold, another path will be selected for I/O transfer, if there is one available.

[-path-lun-in-use-count <integer>,...] - Number of LUNs in the in-use state on this path

Selects information about disks with paths that have the specified in-use-count.

[-initiator-lun-in-use-count <integer>,...] - Number of LUNs in the in-use state on this initiator

Selects information about disks with a path through an initiator that has the specified in-use-count.

[-target-lun-in-use-count <integer>,...] - Number of LUNs in the in-use state on this target

Selects information about disks with a path through a target port that has the specified in-use-count.

[-preferred-target-port {true|false}] - Whether or not target port group is preferred

Selects information about disks that match the specified parameter value indicating whether the backing storage is ALUA (Assymetric Logical Unit Access) capable and has specified the array target port on this path to be a preferred target port for I/O.

[-vmdisk-device-id <integer>,...] - Virtual disk device ID

Selects information about disks that have the specified virtual disk device ID.

[-host-adapter <text>] - Primary Path Host Adapter

Selects information about disks that are currently using the specified Host Bus Adapter.

[-primary-port <text>] - Primary Path Disk Port

Selects information about disks that use the specified primary port.

[-secondary-name <disk path name>] - Secondary Path Name

Selects information about disks that use the specified secondary path name, for multipath configuration.

[-secondary-port <text>] - Secondary Path Disk Port

Selects information about disks that use the specified secondary port.

Examples

The following example displays information about all disks:

```
cluster1::> storage array disk paths show
Disk Name: 1.0.20
UID:
```

5000C500:0979E09F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000

LUN

| Link | Disk I/O | Initiator | ID | Acc | Use | Target Port |
|------------|----------|-----------|----|------|-----|------------------|
| Controller | | | | | | |
| TPGN | Speed | (KB/s) | | IOPS | | |
| node2 | | 3a | 0 | AO | INU | 5000c5000979e09d |
| 80 | 9 Gb/S | 0 | | 0 | | |
| node2 | | 3c | 0 | AO | RDY | 5000c5000979e09e |
| 12 | 9 Gb/S | 0 | | 0 | | |
| node1 | | 3a | 0 | AO | RDY | 5000c5000979e09e |
| 12 | 9 Gb/S | 0 | | 0 | | |
| node1 | | 3c | 0 | AO | INU | 5000c5000979e09d |
| 80 | 9 Gb/S | 0 | | 0 | | |

Disk Name: 1.0.22

UID:

5000C500:0979E3C3:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000

LUN

| Link | Disk I/O | Initiator | ID | Acc | Use | Target Port |
|------------|----------|-----------|----|------|-----|------------------|
| Controller | | | | | | |
| TPGN | Speed | (KB/s) | | IOPS | | |
| node2 | | 3a | 0 | AO | INU | 5000c5000979e3c1 |
| 83 | 9 Gb/S | 0 | | 0 | | |
| node2 | | 3c | 0 | AO | RDY | 5000c5000979e3c2 |
| 15 | 9 Gb/S | 0 | | 0 | | |
| node1 | | 3a | 0 | AO | RDY | 5000c5000979e3c2 |
| 15 | 9 Gb/S | 0 | | 0 | | |
| node1 | | 3c | 0 | AO | INU | 5000c5000979e3c1 |
| 83 | 9 Gb/S | 0 | | 0 | | |

Disk Name: 1.0.19

UID:

5000C500:0979E3F3:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000

LUN

| Link | Disk I/O | Initiator | ID | Acc | Use | Target Port |
|------------|----------|-----------|----|------|-----|------------------|
| Controller | | | | | | |
| TPGN | Speed | (KB/s) | | IOPS | | |
| node2 | | 3a | 0 | AO | RDY | 5000c5000979e3f1 |
| 86 | 9 Gb/S | 0 | | 0 | | |
| node2 | | 3c | 0 | AO | INU | 5000c5000979e3f2 |
| 18 | 9 Gb/S | 0 | | 0 | | |
| node1 | | 3a | 0 | AO | INU | 5000c5000979e3f2 |
| 18 | 9 Gb/S | 0 | | 0 | | |
| node1 | | 3c | 0 | AO | RDY | 5000c5000979e3f1 |
| 86 | 9 Gb/S | 0 | | 0 | | |


```

Disk Name: 1.0.16
UID:
5000C500:0979EBEB:00000000:00000000:00000000:00000000:00000000:00
000000:00000000

                                LUN
Link      Disk I/O
Controller Initiator      ID  Acc  Use  Target Port
TPGN      Speed      (KB/s)      IOPS
node2      3a          0  AO   INU   5000c5000979ebe9
71  9 Gb/s      283          3
node2      3c          0  AO   RDY   5000c5000979ebeb
3  9 Gb/s      0          0
node1      3a          0  AO   RDY   5000c5000979ebeb
3  9 Gb/s      0          0
node1      3c          0  AO   INU   5000c5000979ebe9
71  9 Gb/s      3          0
[...]
```

storage array port modify

Make changes to a target port record.

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

Description

The `storage array port modify` command lets the user change array target port parameters.

Parameters

-name <text> - Name

Selects the array ports that match this parameter value. The storage array name is either generated by Data ONTAP or assigned by the user.

-wwnn <text> - WWNN

Selects the array ports that match this parameter value.

-wwpn <text> - WWPN

Selects the array ports that match this parameter value.

[-max-queue-depth <integer>] - Target Port Queue Depth

The target port queue depth for this target port.

[-utilization-policy {normal|defer}] - Utilization Policy

The policy used in automatically adjusting the queue depth of the target port based on its utilization.

Examples

This command changes the maximum queue depth for this target port to 32.

```
cluster1::> storage array port modify -name HITACHI_DF600F_1 -wwnn  
50060e80004291c0 -wwpn 50060e80004291c0 -max-queue-depth 32
```

storage array port remove

Remove a port record from an array profile.

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

Description

The `storage array port remove` command removes a port from the array database. You might want to remove ports that are no longer connected to the clustered node. Port information can change after hardware replacement, rezoning, or similar configuration activities. The database retains the records about previous ports unless you remove the information.

Parameters

-name <text> - Name

Selects the array ports that match this parameter value. The storage array name is either generated by Data ONTAP or assigned by the user.

{ [-wwnn <text>] - WWNN

Selects the array ports that match this parameter value.

[-wwpn <text>] - WWPN

Selects the array ports that match this parameter value.

| [-target-iqn <text>] - Target IQN

Selects the array ports that match this parameter value.

[-tpgt <integer>] - TPGT }

Selects the array ports that match this parameter value.

Examples

This command removes a port record from the array profiles database.

```
cluster1::> storage array port remove -name HITACHI_DF600F_1 -wwnn  
50060e80004291c0 -wwpn 50060e80004291c0
```

storage array port show

Display information about a storage array's target ports.

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

Description

The `storage array port show` command displays all the target ports known to the cluster for a given storage array (if an array name is specified) or for all storage arrays if no storage array name is specified. Target ports remain in the database as part of an array profile unless you explicitly remove them from the database.

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.

[-name <text>] - Name

Selects the array ports that match this parameter value. The storage array name is either generated by Data ONTAP or assigned by the user.

[-wwnn <text>] - WWNN

Selects the array ports that match this parameter value.

[-wwpn <text>] - WWPN

Selects the array ports that match this parameter value.

[-max-queue-depth <integer>] - Target Port Queue Depth

Selects the array ports that match this parameter value.

[-node {<nodename>|local}] - Controller Name

Selects the array ports that match this parameter value.

[-initiator-port <text>,...] - Initiator Port

Selects the array ports that match this parameter value.

[-average-dynamic-queue-depth <integer>,...] - Average Dynamic Queue Depth (privilege: advanced)

The average value of the dynamic target port queue depth.

[-average-latency-per-iop <integer>,...] - Average Latency Per IOP

Selects the array ports that match this parameter value (average latency per I/O performed in microseconds).

[-average-pending <integer>,...] - Average Pending (privilege: advanced)

Selects the array ports that match this parameter value (average over time of how many commands are on the outstanding queue).

[`-average-waiting` <integer>,...] - Average Waiting (privilege: advanced)

Selects the array ports that match this parameter value (average over time of how many commands are on the waiting queue).

[`-connection-type` {`direct`|`fabric`|`ISCSI`}] - Connection Type

Selects the array ports that match this parameter value (type of connection between the controller and the back end storage).

[`-dynamic-queue-depth` <integer>,...] - Dynamic Queue Depth (privilege: advanced)

Current dynamic target port queue depth, the maximum number of commands allowed outstanding.

[`-max-pending` <integer>,...] - Max Pending (privilege: advanced)

Selects the array ports that match this parameter value (largest number of commands observed on the outstanding queue).

[`-max-waiting` <integer>,...] - Max Waiting (privilege: advanced)

Selects the array ports that match this parameter value (largest number of commands observed on the waiting queue).

[`-path-link-errors` <integer>,...] - Link Error count on path

Selects the array ports that match this parameter value.

[`-percent-busy` <integer>,...] - Percent Busy

Selects the array ports that match this parameter value (percentage of time I/Os are outstanding on the port).

[`-percent-waiting` <integer>,...] - Percent Waiting

Selects the array ports that match this parameter value (percentage of time there are I/Os waiting on the throttle list on the target port).

[`-switch-port` <text>] - Switch Port

Selects the array ports that match this parameter value (for fabric attached connections, the switch port the array target port is connected to; N/A for direct attached).

[`-target-io-kbps` <integer>,...] - Kbytes of I/O per second to Target (Rolling Average)

Selects the array ports that match this parameter value.

[`-target-iops` <integer>,...] - Number of IOPS to Target (Rolling Average)

Selects the array ports that match this parameter value.

[`-target-lun-in-use-count` <integer>,...] - Target LUN In Use Count

Selects the array ports that match this parameter value (number of IN-USE disks on this target port).

[`-target-port-speed` <text>] - Target Port Speed

Selects the array ports that match this parameter value (speed that the target port has negotiated with its connected switch port, or initiator port if direct attached).

[`-utilization-policy` {`normal`|`defer`}] - Utilization Policy

The policy used when sending I/O to an array target port when it reaches maximum queue depth. Possible values are:

- normal - This policy aggressively competes for target port resources, in effect competing with other hosts. (default)
- defer - This policy does not aggressively compete for target port resources, in effect deferring to other hosts.

Examples

The example below displays the port information for a single port.

```
cluster1::> storage array port show -wwpn 50060e80004291c0
Array Name: HITACHI_DF600F_1
WWNN: 50060e80004291c0
WWPN: 50060e80004291c0
Connection Type: fabric
Switch Port: vgbr300s89:9
Link Speed: 4 GB/s
Max Queue Depth: 1024
Utilization Policy: normal
```

| Link | Initiator | LUN | Count | IOPS | KB/s | %busy | %waiting |
|-------|-------------|-------|-------|-------|-------|-------|----------|
| Node | | | | | | | |
| Errs | | | | | | | |
| ----- | ----- | ----- | ----- | ----- | ----- | ----- | ----- |
| ---- | | | | | | | |
| 0 | vnv3070f20a | 0b | 2 | 0 | 0 | 0 | 0 |
| 0 | vnv3070f20b | 0b | 2 | 0 | 0 | 0 | 0 |

storage automated-working-set-analyzer commands

storage automated-working-set-analyzer show

Display running instances

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

Description

The automated-working-set-analyzer show command displays the Automated Working-set Analyzer running instances.

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 (privilege: advanced)

This parameter indicates the node name that the AWA instance runs on.

[-flash-cache {true|false}] - Flash Cache Node-wide Modeling (privilege: advanced)

This parameter indicates whether the AWA is modeling flash-cache.

[-aggregate-uuid <UUID>] - Uuid of the Aggregate (privilege: advanced)

This parameter indicates the aggregate uuid that the AWA instance runs on.

[-aggregate <aggregate name>] - Aggregate (privilege: advanced)

This parameter indicates the aggregate name that the AWA instance runs on.

[-working-set-size {true|false}] - Working Set Size (privilege: advanced)

This parameter indicates whether the AWA instance is configured to find the working set size.

[-start-time <Date>] - Starting Time (privilege: advanced)

This parameter indicates the time when the AWA instance was started.

[-total-intervals <integer>] - Total Interval Count (privilege: advanced)

This parameter indicates the total number of intervals that the AWA instance has covered.

[-read-throughput {<integer>[Bps|KBps|MBps|GBps]}] - Read Throughput (privilege: advanced)

This parameter indicates the maximum read throughput over an interval that AWA has observed from the storage disks.

[-write-throughput {<integer>[Bps|KBps|MBps|GBps]}] - Write Throughput (privilege: advanced)

This parameter indicates the maximum write throughput over an interval that AWA has observed to the storage disks

[-cacheable-read <percent>] - Cacheable Read (privilege: advanced)

This parameter indicates the maximum percent of cacheable read over an interval that AWA has observed. Cacheable reads are non-sequential reads, i.e., the percentage of data reads that could have been cached.

[-cacheable-write <percent>] - Cacheable Write (privilege: advanced)

This parameter indicates the maximum percent of cacheable write over an interval that AWA has observed. Cacheable writes are random overwrites, percentage of disk writes that could have been cached.

[-projected-cache-size {<integer>[KB|MB|GB|TB|PB]}] - Max Projected Cache Size (privilege: advanced)

This parameter indicates the projected Flash Pool cache usage.

[-projected-read-hit <percent>] - Projected Read Hit (privilege: advanced)

This parameter indicates the percentage of blocks that could be read from the Flash Pool cache instead of HDDs.

[`-projected-write-hit <percent>`] - Projected Write Hit (privilege: advanced)

This parameter indicates the percentage of block overwrites that could go to the Flash Pool cache instead of HDDs.

[`-referenced-interval-id <integer>`] - Referenced Interval ID (privilege: advanced)

This parameter indicates the interval in which the cache size effect information is derived from.

[`-referenced-interval-time <Date>`] - Referenced Interval Time (privilege: advanced)

This parameter indicates the time when the referenced interval for the cache size effect information is derived from.

[`-referenced-interval-cache-size {<integer>[KB|MB|GB|TB|PB]}`] - Referenced Interval Cache Size (privilege: advanced)

This parameter indicates the cache size at the end of the referenced interval from which the cache size effect information is based on.

[`-read-hit-20 <percent>`] - 20% Cache Read Hit (privilege: advanced)

This parameter indicates the predicted read hit rate when the cache size is 20% of the referenced cache size.

[`-read-hit-40 <percent>`] - 40% Cache Read Hit (privilege: advanced)

This parameter indicates the predicted read hit rate when the cache size is 40% of the referenced cache size.

[`-read-hit-60 <percent>`] - 60% Cache Read Hit (privilege: advanced)

This parameter indicates the predicted read hit rate when the cache size is 60% of the referenced cache size.

[`-read-hit-80 <percent>`] - 80% Cache Read Hit (privilege: advanced)

This parameter indicates the predicted read hit rate when the cache size is 80% of the referenced cache size.

[`-read-hit-100 <percent>`] - 100% Cache Read Hit (privilege: advanced)

This parameter indicates the predicted read hit rate when the cache size is 100% of the referenced cache size.

[`-write-hit-20 <percent>`] - 20% Cache Write Hit (privilege: advanced)

This parameter indicates the predicted write hit rate when the cache size is 20% of the referenced cache size.

[`-write-hit-40 <percent>`] - 40% Cache Write Hit (privilege: advanced)

This parameter indicates the predicted write hit rate when the cache size is 40% of the referenced cache size.

[`-write-hit-60 <percent>`] - 60% Cache Write Hit (privilege: advanced)

This parameter indicates the predicted write hit rate when the cache size is 60% of the referenced cache size.

[`-write-hit-80 <percent>`] - 80% Cache Write Hit (privilege: advanced)

This parameter indicates the predicted write hit rate when the cache size is 80% of the referenced cache size.

size.

[-write-hit-100 <percent>] - 100% Cache Write Hit (privilege: advanced)

This parameter indicates the predicted write hit rate when the cache size is 100% of the referenced cache size.

[-num-intervals-show <integer>] - Number of intervals to show (privilege: advanced)

This parameter indicates the number of intervals to the past this command is showing.

Examples

The following example shows a running instance of automated-working-set-analyzer on node *node1* for aggregate *aggr0*.

```
cluster1::> cluster-1::*> storage automated-working-set-analyzer show
Node          FC      Aggregate  wss      Intervals Start Time
-----
node1         false aggr0      false      125 Wed Jul 22 13:58:17
2015
```

storage automated-working-set-analyzer start

Command to start Automated Working Set Analyzer on node or aggregate

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

Description

The automated-working-set-analyzer start command enables the Automated Workload Analyzer that is capable of doing the following:

- Flash Pool modeling for an aggregate
- Flash Cache modeling for a node - can not specify an aggregate.
- Working set size estimation
- Workload monitoring

Parameters

-node <nodename> - Node Name (privilege: advanced)

This parameter indicates the node name that the AWA instance runs on.

[-flash-cache {true|false}] - Flash Cache Node-wide Modeling (privilege: advanced)

This parameter indicates whether the AWA is modeling flash-cache.

[-aggregate <aggregate name>] - Aggregate (privilege: advanced)

This parameter indicates the aggregate name that the AWA instance runs on.

[`-working-set-size {true|false}`] - Working Set Size (privilege: advanced)

This parameter indicates whether the AWA instance is configured to find the working set size.

Examples

```
cluster1::> storage automated-working-set-analyzer start -node vsim1
-aggregate aggr0
```

storage automated-working-set-analyzer stop

Command to stop Automated Working Set Analyzer on node or aggregate

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

Description

The storage automated-working-set-analyzer stop command terminates one or multiple Automated Workload Analyzer running instances.

Parameters

`-node <nodename>` - Node Name (privilege: advanced)

This parameter indicates the node name that the AWA instance runs on.

[`-flash-cache {true|false}`] - Flash cache node-wide modeling (privilege: advanced)

This parameter indicates whether the AWA is modeling flash-cache.

[`-aggregate <aggregate name>`] - Aggregate (privilege: advanced)

This parameter indicates the aggregate name that the AWA instance runs on.

Examples

```
cluster1::> storage automated-working-set-analyzer stop -node vsim1
-aggregate aggr1
```

storage automated-working-set-analyzer volume show

Displays the Automated Working Set Analyzer volume table

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

Description

The automated-working-set-analyzer volume show command displays the volume statistics reported by the corresponding Automated Working-set Analyzer running instances.

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 (privilege: advanced)

This parameter indicates the node name that the AWA instance runs on.

[-flash-cache {true|false}] - Flash Cache Node-wide Modeling (privilege: advanced)

This parameter indicates whether the AWA is modeling flash-cache.

[-vol-uuid <UUID>] - Uuid of the Volume (privilege: advanced)

This parameter indicates the volume uuid that this command is issued on.

[-aggregate <aggregate name>] - Aggregate (privilege: advanced)

This parameter indicates the aggregate name that the AWA instance runs on.

[-volume <volume name>] - Volume (privilege: advanced)

This parameter indicates the volume name that this command is issued on.

[-rank <integer>] - Cache Benefit Rank (privilege: advanced)

This parameter indicates the rank of this volume among all volumes that would be most benefited by the modeled cache technology based on the AWA prediction.

[-read-throughput {<integer>[Bps|KBps|MBps|GBps]}] - Read Throughput (privilege: advanced)

This parameter indicates the maximum read throughput over an interval that AWA has observed from the storage disks for this volume.

[-write-throughput {<integer>[Bps|KBps|MBps|GBps]}] - Write Throughput (privilege: advanced)

This parameter indicates the maximum write throughput over an interval that AWA has observed to the storage disks for this volume.

[-cacheable-read <percent>] - Cacheable Read (privilege: advanced)

This parameter indicates the maximum percent of cacheable read over an interval that AWA has observed for this volume. Cacheable reads are non-sequential reads, i.e., the percentage of data reads that could have been cached.

[-cacheable-write <percent>] - Cacheable Write (privilege: advanced)

This parameter indicates the maximum percent of cacheable write over an interval that AWA has observed. Cacheable writes are random overwrites, percentage of disk writes that could have been cached.

[-projected-cache-size {<integer>[KB|MB|GB|TB|PB]}] - Max Projected Cache Size (privilege: advanced)

This parameter indicates the projected Flash Pool cache usage by this volume.

[`-projected-read-hit <percent>`] - Projected Read Hit (privilege: advanced)

This parameter indicates the percentage of blocks that could be read from the Flash Pool cache instead of HDDs for this volume.

[`-projected-write-hit <percent>`] - Projected Write Hit (privilege: advanced)

This parameter indicates the percentage of block overwrites that could go to the Flash Pool cache instead of HDDs for this volume.

[`-num-intervals-show <integer>`] - Number of intervals to show (privilege: advanced)

This parameter indicates the number of intervals to the past this command is showing.

Examples

```
cluster1::> cluster-1::*> storage automated-working-set-analyzer volume
show
Node          FC      Aggregate  Volume      Rank  Read Thrupt  Write
Thrupt
-----
vsim1         false aggr0      vol0         1     230.47KBps
580.09KBps
```

storage disk commands

storage disk assign

Assign ownership of a disk to a system

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

Description

The `storage disk assign` command is used to assign ownership of an unowned disk or array LUN to a specific node. You can also use this command to change the ownership of a disk or an array LUN to another node. You can designate disk ownership by specifying disk names, array LUN names, wildcards, or all (for all disks or array LUNs visible to the node). For disks, you can also set up disk ownership autoassignment. You can also assign disks to a particular pool. You can also assign disks by copying ownership from another disk.

Parameters

{ [`-disk <disk path name>`] - Disk Path

This specifies the disk or array LUN that is to be assigned. Disk names take one of the following forms:

- Disks are named in the form `<stack-id>.<shelf>.<bay>`
- Disks on multi-disk carriers are named in the form `<stack-id>.<shelf>.<bay>.<lun>`
- Virtual disks are named in the form `<prefix>.<number>`, where prefix is the storage array's prefix and number is a unique ascending number.

Disk names take one of the following forms on clusters that are not yet fully upgraded to Data ONTAP 8.3:

- Disks that are not attached to a switch are named in the form `<node>:<host_adapter>.<loop_ID>`. For disks with a LUN, the form is `<node>:<host_adapter>.<loop_ID>L<LUN>`. For instance, disk number 16 on host adapter 1a on a node named node0a is named node0a:1a.16. The same disk on LUN lun0 is named node0a:1a.16Lun0.
- Disks that are attached to a switch are named in the form `<node>:<switch_name>:<switch_port>.<loop_ID>`. For disks with a LUN, the form is `<node>:<switch_name>:<switch_port>.<loop_ID>L<LUN>`. For instance, disk number 08 on port 11 of switch fc1 on a node named node0a is named node0a:fc1:11.08. The same disk on LUN lun1 is named node0a:fc1:11.08Lun1.

Before the cluster is upgraded to Data ONTAP 8.3, the same disk can have multiple disk names, depending on how the disk is connected. For example, a disk known to a node named alpha as alpha:1a.19 can be known to a node named beta as beta:0b.37. All names are listed in the output of queries and are equally valid. To determine a disk's unique identity, run a detailed query and look for the disk's universal unique identifier (UUID) or serial number.

A subset of disks or array LUNs can be assigned using the wildcard character (*) in the `-disk` parameter. Either the `-owner`, the `-sysid`, or the `-copy-ownership-from` parameter must be specified with the `-disk` parameter. Do not use the `-node` parameter with the `-disk` parameter.

| **-disklist <disk path name>,...** - Disk list

This specifies the List of disks to be assigned.

| **-all <true>** - Assign All Disks

This optional parameter causes assignment of all visible unowned disks or array LUNs to the node specified in the `-node` parameter. The `-node` parameter must be specified with the `-all` parameter. When the `-copy-ownership-from` parameter is specified with the `-node` parameter, it assigns disk ownership based on the `-copy-ownership-from` parameter; otherwise it assigns ownership of the disks based on the `-node` parameter. Do not use the `-owner` or the `-sysid` parameter with the `-all` parameter.

| **[-T, -type {ATA | BSAS | FCAL | FSAS | LUN | MSATA | SAS | SSD | VMDISK | SSD-NVM | SSD-CAP | VMLUN | VMLUN-SSD}]** - Storage Type

This optional parameter assigns ownership of a specific type of disk or array LUN (or a set of disks/array LUNs) to a node. The `-count` parameter must be specified with the `-type` parameter.

-n, -count <integer> - Disk Count

This optional parameter assigns ownership of a number of disks or array LUNs specified in the `-count` parameter, to a node.

| **-auto <true>** - Auto Assign }

This optional parameter causes all visible disks eligible for autoassignment to be immediately assigned to the node specified in the `-node` parameter, regardless of the setting of the `disk.auto_assign` option. Only unowned disks on loops or stacks owned wholly by that system and which have the same pool information will be assigned. The `-node` parameter must be specified with the `-auto` parameter. Do not use the `-owner`, the `-sysid`, or the `-copy-ownership-from` parameter with the `-auto` parameter. When possible, use the `-auto` parameter rather than the `-all` parameter to conform to disk ownership best practices. The `-auto` parameter is ignored for array LUNs.

[-p, -pool <integer>] - Pool

This optional parameter specifies the pool to which a disk must be assigned. It can take values of Pool0 or Pool1.

{ [-o, -owner <nodename>] - Owner Name

This optional parameter specifies the node to which the disk or array LUN has to be assigned.

[-s, -sysid <nvramid>] - New Owner ID

This optional parameter specifies the serial number (NVRAM ID) of the node to which the disk or array LUN has to be assigned.

| [-copy-ownership-from <disk path name>] - Disk Name to Copy Ownership }

This optional parameter specifies the disk name from where the node needs to copy disk ownership information. You can use this parameter for disks to have the same ownership as the provided input disk.

[-c, -checksum {block|zoned|advanced_zoned}] - Checksum Compatibility

This optional parameter is used to set the checksum type for a disk or an array LUN. The possible values are *block*, *zoned*, and *advanced_zoned*. This operation will fail if the specified disk is incompatible with the specified checksum type. A newly created aggregate with zoned checksum array LUNs is assigned advanced zoned checksum (AZCS) checksum type. AZCS checksum type provides more functionality than the "version 1" zoned checksum type which has been supported in previous Data ONTAP releases. Zoned checksum spare array LUNs added to an existing zoned checksum aggregate continue to be zoned checksum. Zoned checksum spare array LUNs added to an AZCS checksum type aggregate use the AZCS checksum scheme for managing checksums. For some disks (e.g. FCAL, SSD, SAS disks), the checksum type cannot be modified. For more information on modifying the checksum type, refer to the "Physical Storage Management Guide".

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

This optional parameter forces the assignment of ownership of an already owned disk to a node. This parameter could also be used to assign an array LUN with a redundancy error, for example, if the array LUN is available on only one path. For a disk which is part of a live aggregate, even specification of the *-force* parameter would not force the assignment, since it would be catastrophic.

[-N, -node <nodename>] - Node Name (For Auto Assign)

This optional parameter is used with either the *-auto* or the *-all* parameter. If used with the *-auto* parameter, all disks which are visible to the node specified in the *-node* parameter and which are eligible for autoassignment would be assigned to it. If used with the *-all* parameter, all unowned disks or array LUNs visible to the node would be assigned to it.

{ [-root <true>] - Root Partition of Root-Data or Root-Data1-Data2 Partitioned Disk (privilege: advanced)

This optional parameter assigns the root partition of a root-data/root-data1-data2 partitioned disk. You cannot use this parameter with disks that are part of a storage pool. The default value is *false*.

| [-data <true>] - Data Partition of Root-Data Partitioned Disk (privilege: advanced)

This optional parameter assigns the data partition of a root-data partitioned disk. You cannot use this parameter with disks that are part of a storage pool. The default value is *false*.

| [-data1 <true>] - Data1 Partition of Root-Data1-Data2 Partitioned Disk (privilege: advanced)

This optional parameter assigns the data1 partition of a root-data1-data2 partitioned disk. You cannot use this parameter with disks that are part of a storage pool. The default value is *false*.

| [-data2 <true>] - Data2 Partition of Root-Data1-Data2 Partitioned Disk (privilege: advanced) }

This optional parameter assigns the data2 partition of a root-data1-data2 partitioned disk. You cannot use this parameter with disks that are part of a storage pool. The default value is *false*.

Examples

The following example assigns ownership of an unowned disk named ``_1_``. ``_1_``. ``_16_`` to a node named ``_node1_`` :

```
cluster1::> storage disk assign -disk 1.1.16 -owner node1
```

The following example assigns all unowned disks or array LUNs visible to a node named *node1* to itself:

```
cluster1::> storage disk assign -all -node node1
```

The following example autoassigns all unowned disks (eligible for autoassignment) visible to a node named *node1* to itself:

```
cluster1::> storage disk assign -auto -node node1
```

The following two examples show the working of the *-force* parameter with a spare disk that is already owned by another system:

```
cluster1::> storage disk assign -disk 1.1.16 -owner node1
Error: command failed: Failed to assign disks. Reason: Disk 1.1.16 is
already owned.
```

```
cluster1::> storage disk assign -disk 1.1.16 -owner node1 -force
Success.
```

The following example assigns ownership of the set of unowned disks on *<stack>1*, to a node named *node1*:

```
cluster1::> storage disk assign -disk 1.* -owner node1
```

The following example assigns ownership of unowned disk *1.1.16* by copying ownership from disk *1.1.18*:

```
cluster1:> storage disk assign -disk 1.1.16  
-copy-ownership-from 1.1.18
```

The following example assigns all unowned disks visible to a node named ``_node1_`` by copying ownership from disk ``_1_`` . ``_1_`` . ``_18_`` :

```
cluster1:> storage disk assign -all -node node1  
-copy-ownership-from 1.1.18
```

The following example assigns the root partition of disk *1.1.16* to node1.

```
cluster1:> storage disk assign -disk 1.1.16 -owner node1 -root true  
-force true
```

The following example assigns the data partition of root-data partitioned disk *1.1.16* to node1.

```
cluster1:> storage disk assign -disk 1.1.16 -owner node1 -data true  
-force true
```

The following example assigns the data1 partition of root-data1-data2 partitioned disk *1.1.24* to node1.

```
cluster1:> storage disk assign -disk 1.1.24 -owner node1 -data1 true  
-force true
```

The following example assigns the data2 partition of root-data1-data2 partitioned disk *1.1.24* to node1.z33

```
cluster1:> storage disk assign -disk 1.1.24 -owner node1 -data2 true  
-force true
```

storage disk fail

Fail the file system disk

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

Description

The `storage disk fail` command can be used to manually force a file system disk to fail. It is used to remove a file system disk that may be logging excessive errors and requires replacement. To unfail a disk, use

the [storage disk unfail](#) command.

Parameters

-disk <disk path name> - Disk Name

This parameter specifies the disk to be failed.

[-i, -immediate <true>] - Fail immediately

This parameter optionally specifies whether the disk is to be failed immediately. It is used to avoid Rapid RAID Recovery and remove the disk from the RAID configuration immediately. Note that when a file system disk has been removed in this manner, the RAID group to which the disk belongs enters degraded mode (meaning a disk is missing from the RAID group). If a suitable spare disk is available, the contents of the disk being removed are reconstructed onto that spare disk.

Examples

The following example fails a disk named 1.1.16 immediately:

```
cluster1::> storage disk fail -disk 1.1.16 -i true
WARNING: The system will not prefail the disk and its contents will not be
copied to a replacement disk before being failed out. Do you want to
fail out the disk immediately? {y|n}: y
```

Related Links

- [storage disk unfail](#)

storage disk reassign

(DEPRECATED)-Change the default owner of all disks from one node to another

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

Description

The `storage disk reassign` is deprecated and may be removed in a future release of Data ONTAP. Disk reassignment is no longer required as part of a controller replacement procedure. For further information, see the latest controller or NVRAM FRU replacement flyer for your system. This command changes the ownership of all disks on a node to the ownership of another node. Use this command only when a node has a complete failure (for instance, a motherboard failure) and is replaced by another node. If the node's disks have already been taken over by its storage failover partner, use the `-force` parameter.

Parameters

-s, -homeid <nvr amid> - Current Home ID (privilege: advanced)

This specifies the serial number of the failed node.

-d, -newhomeid <nvr amid> - New Home ID (privilege: advanced)

This specifies the serial number of the node that is to take ownership of the failed node's disks.

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

This optionally specifies whether to force the reassignment operation. The default setting is `false`.

Examples

In the following example, a node named `node0` and having serial number 12345678 has failed. Its disks have not been taken over by its storage failover partner. A replacement node with serial number 23456789 was installed and connected to `node0`'s disk shelves. To assign `node0`'s disks to the new node, start the new node and run the following command:

```
cluster::*> storage disk reassign -homeid 12345678 -newhomeid 23456789
node0's disks 1.1.11, 1.1.12, 1.1.13, 1.1.14, 1.1.15, 1.1.16, 1.1.23 and
1.1.24
were reassigned to new owner with serial number 23456789.
```

In the following example, a similar failure has occurred, except that `node0`'s disks have been taken over by its storage failover partner, `node1`. A new node with serial number 23456789 has been installed and configured. To assign the disks that previously belonged to `node0` to this new node, run the following command:

```
cluster::*> storage disk reassign -homeid 12345678 -newhomeid 23456789
-force true
node0's disks 1.1.11, 1.1.12, 1.1.13, 1.1.14, 1.1.15, 1.1.16, 1.1.23 and
1.1.24
were reassigned to new owner with serial number 23456789.
```

storage disk refresh-ownership

Refresh the disk ownership information on a node

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

Description

This command updates the disk ownership information for all the disks attached to a node to the latest view for all the nodes in the cluster. During normal operations, disk ownership is kept up to date automatically. In certain circumstances, however, disk ownership must be updated manually. If this is required, EMS messages will indicate that this command should be run. If the `-node` parameter is provided, the disk ownership information is updated only on the node specified.

Parameters

[-node {<nodename>|local}] - Node (privilege: advanced)

If this parameter is provided, the disk ownership information is updated only on the node specified.

Examples

The following example refreshes the disk ownership information for all the nodes in the cluster:

```
cluster1::> storage disk refresh-ownership
```

storage disk remove-reservation

Removes reservation from an array LUN marked as foreign.

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

Description

The `storage disk remove-reservation` command removes persistent reservation from a specified foreign array LUN.

Parameters

-disk <disk path name> - Disk Name (privilege: advanced)

This specifies the disk from which persistent reservation is to be removed.

Examples

The following example removes the persistent reservation from a disk named node1:switch01:port.126L1.

```
cluster1::> storage disk remove-reservation -disk  
node1:switch01:port.126L1
```

storage disk remove

Remove a spare disk

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

Description

The `storage disk remove` command removes the specified spare disk from the RAID configuration, spinning the disk down when removal is complete.

This command does not remove disk ownership information from the disk. Therefore, if you plan to reuse the disk in a different storage system, you should use the [storage disk removeowner](#) command instead. See the "Physical Storage Management Guide" for the complete procedure.



For systems with multi-disk carriers, it is important to ensure that none of the disks in the carrier are filesystem disks before attempting removal. To convert a filesystem disk to a spare disk, see [storage disk replace](#).

Parameters

-disk <disk path name> - Disk Name

This parameter specifies the disk to be removed.

Examples

The following example removes a spare disk named 1.1.16:

```
cluster1::> storage disk remove -disk 1.1.16
```

Related Links

- [storage disk removeowner](#)
- [storage disk replace](#)

storage disk removeowner

Remove disk ownership

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

Description

The `storage disk removeowner` command removes ownership from a specified disk. Then disk can then be reassigned to a new owner.

Parameters

-disk <disk path name> - Disk Name

This specifies the disk whose ownership is to be removed.

{ [-root <true>] - Root Partition of Root-Data/Root-Data1-Data2 Partitioned Disk (privilege: advanced)

This optional parameter removes ownership of the root partition of a root-data/root-data1-data2 partitioned disk. You cannot use this parameter with disks that are part of a storage pool. The default value is *false*.

| [-data <true>] - Data Partition of Root-Data Partitioned Disk (privilege: advanced)

This optional parameter removes ownership of the data partition of a root-data partitioned disk. You cannot use this parameter with a root-data1-data2 partitioned disk or disks that are part of a storage pool. The default value is *false*.

| [-data1 <true>] - Data1 Partition of a Root-Data1-Data2 Partitioned Disk (privilege: advanced)

This optional parameter removes ownership of the data1 partition of a root-data1-data2 partitioned disk. You cannot use this parameter with a root-data partitioned disk or disks that are part of a storage pool. The default value is *false*.

| [-data2 <true>] - Data2 Partition of a Root-Data1-Data2 Partitioned Disk (privilege: advanced) }

This optional parameter removes ownership of the data2 partition of a root-data1-data2 partitioned disk. You cannot use this parameter with a root-data partitioned disk or disks that are part of a storage pool. The default value is *false*.

Examples

The following example removes the ownership from a disk named 1.1.27.

```
cluster1::> storage disk removeowner -disk 1.1.27
```

The following example removes ownership of the root partition on disk 1.1.16.

```
cluster1::> storage disk removeowner -disk 1.1.16 -root true
```

The following example removes ownership of the data partition on disk 1.1.16.

```
cluster1::> storage disk removeowner -disk 1.1.16 -data true
```

The following example removes ownership of the data1 partition on disk 1.1.23.

```
cluster1::> storage disk removeowner -disk 1.1.23 -data1 true
```

The following example removes ownership of the data2 partition on disk 1.1.23.

```
cluster1::> storage disk removeowner -disk 1.1.23 -data2 true
```

storage disk replace

Initiate or stop replacing a file-system disk

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

Description

The `storage disk replace` command starts or stops the replacement of a file system disk with spare disk. When you start a replacement, Rapid RAID Recovery begins copying data from the specified file system disk to a spare disk. When the process is complete, the spare disk becomes the active file system disk and the file system disk becomes a spare disk. If you stop a replacement, the data copy is halted, and the file system disk and spare disk retain their initial roles.

Parameters

-disk <disk path name> - Disk Name

This specifies the file system disk that is to be replaced. Disk names take one of the following forms:

- Disks are named in the form `<stack-id>.<shelf>.<bay>`
- Disks on multi-disk carriers are named in the form `<stack-id>.<shelf>.<bay>.<lun>`
- Virtual disks are named in the form `<prefix>.<number>`, where prefix is the storage array's prefix and

number is a unique ascending number.

Disk names take one of the following forms on clusters that are not yet fully upgraded to Data ONTAP 8.3:

- Disks that are not attached to a switch are named in the form `<node>:<host_adapter>.<loop_ID>`. For disks with a LUN, the form is `<node>:<host_adapter>.<loop_ID>L<LUN>`. For instance, disk number 16 on host adapter 1a on a node named node0a is named node0a:1a.16. The same disk on LUN lun0 is named node0a:1a.16Llun0.
- Disks that are attached to a switch are named in the form `<node>:<switch_name>:<switch_port>.<loop_ID>`. For disks with a LUN, the form is `<node>:<switch_name>:<switch_port>.<loop_ID>L<LUN>`. For instance, disk number 08 on port 11 of switch fc1 on a node named node0a is named node0a:fc1:11.08. The same disk on LUN lun1 is named node0a:fc1:11.08Llun1.

Before the cluster is upgraded to Data ONTAP 8.3, the same disk can have multiple disk names, depending on how the disk is connected. For example, a disk known to a node named alpha as alpha:1a.19 can be known to a node named beta as beta:0b.37. All names are listed in the output of queries and are equally valid. To determine a disk's unique identity, run a detailed query and look for the disk's universal unique identifier (UUID) or serial number.

-action {start | stop} - Action

This specifies whether to start or stop the replacement process.

[-replacement <disk path name>] - Replacement

This specifies the spare disk that is to replace the file system disk.

[-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 replace a disk in 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 replace a disk in an aggregate with a disk 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, ensure that disks housed in the same carrier are in different RAID groups.

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

[-m, -allow-mixing <true>] - Allow Mixing of Disks of Different RPM or Pool

This optional parameter specifies whether the disk can be replaced with another disk of different RPM or from different Pool. This parameter affects only the current disk replacement operation.

Examples

The following example begins replacing a file system disk named `1.0.16` with a spare disk named `1.1.14`.

```
cluster1::> storage disk replace -disk 1.0.16 -replacement 1.1.14 -action
start
```

storage disk set-foreign-lun

Sets or Unsets an array LUN as foreign

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

Description

The `storage disk set-foreign-lun` command sets or unsets a specified array LUN as foreign. This command will enable/disable the feature of importing the data from foreign LUN.

Parameters

-disk <disk path name> - Disk Name

This parameter specifies the array LUN which is to be set or unset as foreign.

-is-foreign-lun <true> - Is Foreign LUN

If the parameter value specified is *true* then array LUN is set as foreign. If the parameter value specified is *false* then array LUN foreignness is cleared.

Examples

The following example shows how to set an array LUN as foreign:

```
cluster1::> storage disk set-foreign-lun -disk EMC-1.1 -is-foreign-lun
true
```

The following example shows how to mark an array LUN as not foreign:

```
cluster1::> storage disk set-foreign-lun -disk EMC-1.1 -is-foreign-lun
false
```

storage disk set-led

Identify disks by turning on their LEDs

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

Description

The `storage disk set-led` command controls the LED of a specified disk.

You can turn an LED on or off, cause it to blink or stop blinking, or test it.

This command is useful for locating a disk in its shelf.

Parameters

-action {on|off|blink|blinkoff|testall|resetall} - Action

This parameter specifies the state to which the LED is to be set. Possible values include the following:

- on - The LED is lit steadily
- off - The LED is not lit
- blink - The LED blinks
- blinkoff - The LED stops blinking and is not lit
- testall - This tests the operation of every disk enclosure's hardware and drivers per node. Do not use this value in normal operation.
- resetall - This resets the LED of every disk on the node and lights up the LED of disks with faults.

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

This specifies the disk whose LED is to be set. Disk names take one of the following forms:

- Disks are named in the form *<stack-id>.<shelf>.<bay>*
- Disks on multi-disk carriers are named in the form *<stack-id>.<shelf>.<bay>.<lun>*
- Virtual disks are named in the form *<prefix>.<number>*, where prefix is the storage array's prefix and number is a unique ascending number.

Disk names take one of the following forms on clusters that are not yet fully upgraded to Data ONTAP 8.3:

- Disks that are not attached to a switch are named in the form *<node>:<host_adapter>.<loop_ID>*. For disks with a LUN, the form is *<node>:<host_adapter>.<loop_ID>L<LUN>*. For instance, disk number 16 on host adapter 1a on a node named node0a is named node0a:1a.16. The same disk on LUN lun0 is named node0a:1a.16Llun0.
- Disks that are attached to a switch are named in the form *<node>:<switch_name>:<switch_port>.<loop_ID>*. For disks with a LUN, the form is *<node>:<switch_name>:<switch_port>.<loop_ID>L<LUN>*. For instance, disk number 08 on port 11 of switch fc1 on a node named node0a is named node0a:fc1:11.08. The same disk on LUN lun1 is named node0a:fc1:11.08Llun1.

Before the cluster is upgraded to Data ONTAP 8.3, the same disk can have multiple disk names, depending on how the disk is connected. For example, a disk known to a node named alpha as alpha:1a.19 can be known to a node named beta as beta:0b.37. All names are listed in the output of queries and are equally valid. To determine a disk's unique identity, run a detailed query and look for the disk's universal unique identifier (UUID) or serial number.

| [-adapter <text>] - Adapter Name

The name of the adapter to which the shelves of disks of interest are attached to.

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

The node for which action is to be taken.

[-duration <integer>] - Duration (minutes)

This specifies the duration, in minutes, that the LED is to remain in the specified state. Only actions "on" and "blink" are supported.

[`-iteration <integer>`] - Test iterations

This specifies the number of iterations to run the action for. Only action "test-all" is supported.

Examples

The following example causes the LEDs on all disks whose names match the pattern Cluster1* to turn on for 5 minutes:

```
Cluster1::> storage disk set-led -disk Cluster1* -action on -duration 5
```

The following example causes the LEDs on all disks attached to adapter 0b on Node2 to turn on for 1 minute:

```
Cluster1::> storage disk set-led -node Node2 -adapter 0b -action on  
-duration 1
```

The following example resets the LEDs on all disks on the local node and causes the LEDs of disks with faults to turn on:

```
Cluster1::> storage disk set-led -action resetall
```

The following example causes the LEDs on all disks whose names match the pattern Cluster1* to turn on for 2 minutes:

```
Cluster1::> storage disk set-led -disk Cluster1* -action on -duration 2
```

The following example tests the LEDs on all disks owned by the local node for 3 iterations:

```
Cluster1::> storage disk set-led -action testall -iteration 3
```

storage disk show

Display a list of disk drives and array LUNs

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

Description

The `storage disk show` command displays information about disks and array LUNs. Where it appears in the remainder of this document "disk" may refer to either a disk or an array LUN. By default, the command displays the following information about all disks in column style output:

- Disk name
- Usable space on the disk, in human readable units

- Shelf number
- Bay number
- Container type (aggregate, broken, foreign, labelmaint, maintenance, mediator, remote, shared, spare, unassigned, unknown, volume, or unsupported)
- Position (copy, data, dparity, orphan, parity, pending, present, shared or tparity)
- Container name
- Owning node name

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

Parameters

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

Displays the specified fields for all disks, in column style output.

| [-broken]

Displays the following RAID-related information about broken disks:

- Original owning node name
- Checksum compatibility
- Disk name
- Outage reason
- Host bus adapter
- Shelf number
- Bay number
- Primary port / Channel
- Pool
- Disk type
- RPM (Revolutions per minute)
- Usable size in human readable units
- Physical size in human readable units
- Current owner node

| [-errors]

Displays the following disk information about the disks which have errors.

- Disk Name
- Error Type
- Error Description and corresponding corrective action

| [-longop]

Displays the following information about long-running disk operations, in column style output:

- Disk name

- Whether the disk is marked as prefailed
- Whether the disk is being replaced
- Whether the disk is zeroed
- Copy destination
- Percentage of copy complete
- Percentage of zeroing complete
- Percentage of reconstruction complete

| [**-maintenance**]

Displays the following RAID-related information about disks in the maintenance center:

- Original owning node name
- Checksum compatibility
- Disk name
- Outage Reason
- Host bus adapter
- Shelf number
- Bay number
- Primary port / Channel
- Pool
- Disk type
- RPM (Revolutions per minute)
- Usable size in human readable units
- Physical size in human readable units
- Current owner node

| [**-ownership**]

Displays the following ownership-related information:

- Disk name
- Aggregate name
- Home node name
- Owning node name
- Disaster recovery home node name
- Home node system id
- Owning node system id
- Disaster recovery home node system id
- Reservation node system id
- SyncMirror pool

| **[-partition-ownership]**

Displays the following ownership-related information for partitioned disks:

- Disk name
- Aggregate name
- Owner of root partition on a partitioned disk
- Owner system id of root partition on a partitioned disk
- Owner of data or data1 partition on a root-data or a root-data1-data2 partitioned disk respectively
- Owner system id of data or data1 partition on a root-data or a root-data1-data2 partitioned disk respectively
- Owner of data2 partition on a root-data1-data2 partitioned disk
- Owner system id of data2 partition on a root-data1-data2 partitioned disk
- Owner of the disk which is partitioned
- Owner system id of the disk which is partitioned

| **[-physical]**

Displays the following information about the disk's physical attributes, in column style output:

- Disk name
- Disk type
- Disk vendor
- Disk model
- Firmware revision level
- RPM (Revolutions per minute)
- BPS (Bytes per sector)

| **[-port]**

Displays the following path-related information:

- Disk name and disk port associated with disk primary path
- Disk name and disk port associated with the disk secondary path, for a multipath configuration
- Type, shelf, and bay information for the disks

| **[-raid]**

Displays the following RAID-related information:

- Disk name
- Container type (aggregate, broken, labelmaint, maintenance, mediator, remote, shared, spare, unassigned, unknown, or volume)
- Outage reason
- Position (copy, data, dparity, orphan, parity, pending, present, shared or tparity)
- RAID group name
- Aggregate name

| **`-raid-info-for-aggregate`** |

Displays the following RAID-related information about the disks used in an aggregate:

- Owning node name
- Aggregate name
- Plex name
- RAID group name
- Position (copy, data, dparity, orphan, parity, pending, present, shared or tparity)
- Disk name
- Host bus adapter
- Shelf number
- Bay number
- Primary port / Channel
- Pool
- Disk type
- RPM (Revolutions per minute)
- Usable size in human readable units
- Physical size in human readable units

When this parameter is specified, RAID groups that use shared disks are not included. Use [storage aggregate show-status](#) to show information for all RAID groups and aggregates.

| **`-spare`** |

Displays the following RAID-related information about available spare disks:

- Original owning node name
- Checksum compatibility
- Disk name
- Host bus adapter
- Shelf number
- Bay number
- Primary port / Channel
- Pool
- Disk type
- Disk class
- RPM (Revolutions per minute)
- Usable size in human readable units
- Physical size in human readable units
- Current owner node

| [-ssd-wear]

Displays the following wear life related information about solid state disks:

- Rated Life Used : An estimate of the percentage of device life that has been used, based on the actual device usage and the manufacturer's prediction of device life. A value greater than 99 indicates that the estimated endurance has been used, but this does not necessarily indicate a device failure. Omitted if value is unknown.
- Spare Blocks Consumed Limit : Spare blocks consumed percentage limit reported by the device. When the Spare Blocks Consumed percentage for the device reaches this read-only value, Data ONTAP initiates a disk copy operation to prepare to remove the device from service. Omitted if value is unknown.
- Spare Blocks Consumed : Percentage of device spare blocks that have been used. Each device has a number of spare blocks that will be used when a data block can no longer be used to store data. This value reports what percentage of the spares have already been consumed. Omitted if value is unknown.

| [-virtual-machine-disk-info]

Displays information about Data ONTAP virtual disks, their mapped datastores and their specific backing device attributes, such as: disk or LUN, adapter and initiator details (if applicable).

- Disk name.
- Name of the node.
- Data ONTAP-supplied serial number of the system disk.
- Size of the system disk.
- Name of the disk backing store. A backing store represents a storage location for virtual machine files. It can be a VMFS volume, a directory on network-attached storage, or a local file system path.
- File name of the virtual disk used by the hypervisor. Each Data ONTAP disk is mapped to a unique VM disk file.
- Type of the disk backing store. It can be a VMFS volume, a directory on network-attached storage, or a local file system path.
- Size of the disk backing store.
- Full path to the backing store for network-attached storage. This field is valid only for NAS connections.
- Backing adapter PCI device ID for the virtual disk, for example "50:00.0".
- Backing adapter device name, for example "vmhba32".
- Backing adapter model type, for example "LSI1064E".
- Backing adapter driver name of the initiator.
- The iSCSI name of the disk backing target. This field is valid only for iSCSI connections.
- The iSCSI IP address of the disk backing target. This field is valid only for iSCSI connections.
- SCSI device name for the backing disk. It takes the form target-id:lun-id, for example "2:1".
- Hypervisor-assigned unique ID of the backing device (disk or LUN).
- Backing disk partition number where the corresponding VM disk file resides.
- Size of the backing device (disk or LUN).
- Backing device manufacturer, for example "FUJITSU" or "IBM".
- Backing device model, for example "MBE2073RC" or "LUN".

- Storage account associated with the VM Disk.
- Container associated with the VM Disk.
- Page blob associated with the VM Disk.
- Error (if any) while retrieving virtual disk details.

| **[-vmdisk-backing-info]**

Displays information about the backing disks on certain Data ONTAP-v models:

- Disk name
- Backing disk vendor
- Backing disk model
- Backing disk serial number
- Backing disk device id

| **[-foreign] (privilege: advanced)**

Displays the following foreign LUN import related information about foreign disks:

- Disk name
- Array name
- Capacity in sectors
- Capacity in mb
- Serial Number

| **[-physical-location] (privilege: advanced)**

Displays the following information about disks:

- Disk name
- Container type
- Primary path
- Location
- Home node name
- Physical size in human readable units

| **[-primary-paths] (privilege: advanced)**

Displays the following information about disks:

- Disk Name
- Shelf
- Bay
- Container Type
- Primary Path

[`-instance`] }

Displays detailed disk information. If no disk path name is specified, this parameter displays the same detailed information for all disks as does the `-disk` parameter. If a disk path name is specified, then this parameter displays the same detailed information for the specified disks as does the `-disk` parameter.

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

Displays detailed information about the specified disks. Disk names take one of the following forms:

- Disks are named in the form `<stack-id>.<shelf>.<bay>`
- Disks on multi-disk carriers are named in the form `<stack-id>.<shelf>.<bay>.<lun>`
- Virtual disks are named in the form `<prefix>.<number>`, where prefix is the storage array's prefix and number is a unique ascending number.

Disk names take one of the following forms on clusters that are not yet fully upgraded to Data ONTAP 8.3:

- Disks that are not attached to a switch are named in the form `<node>:<host_adapter>.<loop_ID>`. For disks with a LUN, the form is `<node>:<host_adapter>.<loop_ID>L<LUN>`. For instance, disk number 16 on host adapter 1a on a node named node0a is named node0a:1a.16. The same disk on LUN lun0 is named node0a:1a.16Llun0.
- Disks that are attached to a switch are named in the form `<node>:<switch_name>:<switch_port>.<loop_ID>`. For disks with a LUN, the form is `<node>:<switch_name>:<switch_port>.<loop_ID>L<LUN>`. For instance, disk number 08 on port 11 of switch fc1 on a node named node0a is named node0a:fc1:11.08. The same disk on LUN lun1 is named node0a:fc1:11.08Llun1.

Before the cluster is upgraded to Data ONTAP 8.3, the same disk can have multiple disk names, depending on how the disk is connected. For example, a disk known to a node named alpha as alpha:1a.19 can be known to a node named beta as beta:0b.37. All names are listed in the output of queries and are equally valid. To determine a disk's unique identity, run a detailed query and look for the disk's universal unique identifier (UUID) or serial number.

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

Selects information about disks that are owned by the specified node.

[`-owner-id <nvrmaid>`] - Owner System ID

Selects the disks that are owned by the node with the specified system ID.

[`-is-foreign {true|false}`] - Foreign LUN (privilege: advanced)

Selects information about array LUNs that have been declared to be foreign LUNs.

[`-uid <text>`] - Disk Unique ID

Selects the disks whose unique id matches this parameter value. A disk unique identifier has the form:

``20000000:875D4C32:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000``

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

Selects information about disks that belong to the specified aggregate.

[`-array-name <array name>`] - Array Name

Selects information about the LUNs presented by the specified storage array.

[-average-latency <integer>] - Average I/O Latency Across All Active Paths

Selects information about disks that have the specified average latency.

[-bay <integer>] - Bay

Selects information about disks that are located in the carrier within the specified shelf bay.

[-bps <integer>] - Bytes Per Sector

Selects information about disks that have the specified number of bytes per sector. Possible settings are 512, 520, 4096, and 4160.

[-carrier-id <text>] - Carrier ID

Selects information about disks that are located within the specified multi-disk carrier.

[-checksum-compatibility {advanced_zoned | block | none}] - Checksum Compatibility

Selects information about disks that have the specified checksum compatibility.

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

Selects information about disks that have the specified disk class.

- 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.

[-container-type {aggregate | broken | foreign | labelmaint | maintenance | mediator | remote | shared | spare | unassigned | unknown | unsupported}] - Container Type

Selects information about disks that have the specified container type.

- Aggregate = Disk is used as a physical disk in an aggregate.
- Broken = Disk is in broken pool.
- Foreign = Array LUN has been marked foreign.
- Labelmaint = Disk is in online label maintenance list.
- Maintenance = Disk is in maintenance center.
- Mediator = A mediator disk is a disk used on non-shared HA systems hosted by an external node which is used to communicate the viability of the storage failover between non-shared HA nodes.
- Remote = Disk belongs to the remote cluster.
- Shared = Disk is partitioned or in a storage pool.
- Spare = Disk is a spare disk.
- Unassigned = Disk ownership has not been assigned.
- Unknown = Container is currently unknown. This is the default setting.

- Unsupported = Disk is not supported.

[-container-name <text>] - Container Name

Selects information about disks that have the specified container name. + If a disk is in an aggregate or storage pool, the container name is the name of the aggregate or storage pool. + Spare disks show the SyncMirror Pool to which they belong. + Partitioned disks could return multiple aggregate names.

[-copy-destination <disk path name>] - Copy Destination Name

Selects information about disks whose contents are being copied (due to either Rapid RAID Recovery or disk replacement) to the specified spare disk.

[-copy-percent <integer>] - Percentage of Copy Complete

Selects information about disks that are involved as either a source or destination of a copy operation, (due to either disk replacement or Rapid RAID Recovery) and that have the specified percentage of the copy operation completed.

[-data-owner {<nodename>|local}] - Owner of Data Partition of Root-Data Partitioned Disk

Selects information about disks that have the specified data partition owner name. Used with root-data partitioned disks.

[-data1-owner {<nodename>|local}] - Owner of Data1 Partition of Root-Data1-Data2 Partitioned Disk

Selects information about disks that have the specified data1 partition owner name. Used with root-data1-data2 partitioned disks.

[-data2-owner {<nodename>|local}] - Owner of Data2 Partition of Root-Data1-Data2 Partitioned Disk

Selects information about disks that have the specified data2 partition owner name. Used with root-data1-data2 partitioned disks.

[-data-home {<nodename>|local}] - Home Owner of Data Partition of Root-Data Partitioned Disk

Selects information about disks that have the specified data partition home owner name. Used with root-data partitioned disks.

[-data1-home {<nodename>|local}] - Home Owner of Data1 Partition of Root-Data1-Data2 Partitioned Disk

Selects information about disks that have the specified data1 partition home owner name. Used with root-data1-data2 partitioned disks.

[-data2-home {<nodename>|local}] - Home Owner of Data2 Partition of Root-Data1-Data2 Partitioned Disk

Selects information about disks that have the specified data2 partition home owner name. Used with root-data1-data2 partitioned disks.

[-data-owner-id <nvrleid>] - Owner System ID of Data Partition of Root-Data Partitioned Disk

Selects information about disks that have the specified data partition owner system ID. Used with root-data partitioned disks.

[-data1-owner-id <nvrleid>] - Owner System ID of Data1 Partition of Root-Data1-Data2 Partitioned Disk

Selects information about disks that have the specified data1 partition owner system ID. Used with root-

data1-data2 partitioned disks.

[-data2-owner-id <nvramid>] - Owner System ID of Data2 Partition of Root-Data1-Data2 Partitioned Disk

Selects information about disks that have the specified data2 partition owner system ID. Used with root-data1-data2 partitioned disks.

[-data-home-id <nvramid>] - Home Owner System ID of Data Partition of Root-Data Partitioned Disk

Selects information about disks that have the specified data partition home owner system ID. Used with root-data partitioned disks.

[-data1-home-id <nvramid>] - Home Owner System ID of Data1 Partition of Root-Data1-Data2 Partitioned Disk

Selects information about disks that have the specified data1 partition home owner system ID. Used with root-data1-data2 partitioned disks.

[-data2-home-id <nvramid>] - Home Owner System ID of Data2 Partition of Root-Data1-Data2 Partitioned Disk

Selects information about disks that have the specified data2 partition home owner system ID. Used with root-data1-data2 partitioned disks.

[-disk-io-kbps-total <integer>] - Total Disk Throughput in KBPS Across All Active Paths

Selects information about disks that have attained the specified I/O throughput on all connected paths.

[-disk-iops-total <integer>] - Total Disk IOPs Across All Active Paths

Selects information about disks that have achieved the specified number of IOPs per second on all connected paths.

[-diskpathnames <disk path name>,...] - List of Path-Based Disk Names

Selects information about disks that have all of the specified path names.

[-effective-rpm <integer>] - Effective RPM

Selects information about disks with the specified effective rotational speed.

[-dr-home {<nodename>|local}] - Disaster Recovery Home

Selects information about disks that have the specified Disaster home node.

[-dr-home-id <nvramid>] - Disaster Recovery Home System ID

Selects information about disks whose Disaster home node has the specified system id.

[-drawer <integer>] - Drawer

Selects information about disks that are located in the specified drawer.

[-error-type

{onepath|onedomain|control|foreign|toobig|toosmall|invalidblocksize|targetasymmap|deviceasymmap|failovermisconfig|unknown|netapp|fwdownrev|qualfail|diskfail|notallflashdisk}] - Error Type

Selects information about disks that have the specified error types.

- onepath = The array LUN is accessible only via a single path.

- onedomain = The array LUN is accessible only via a single fault domain.
- control = The array LUN cannot be used because it is a control device.
- foreign = The array LUN is marked as foreign and has some external SCSI reservations other than those from Data ONTAP.
- toobig = The array LUN exceeds the maximum array LUN size that Data ONTAP supports.
- toosmall = The array LUN is less than the minimum array LUN size that Data ONTAP supports.
- invalidblocksize = The array LUN is not a valid block size.
- targetasymmap = The array LUN is presented more than once on a single target port.
- deviceassymmap = The array LUN is presented with multiple IDs.
- failovermisconfig = The array LUN is configured with inconsistent failover methods.
- unknown = The array LUN from a storage array that is not supported by this version of Data ONTAP.
- netapp = A SAN front-end LUN from one Data ONTAP system that is presented as external storage to another Data ONTAP system.
- fwdownrev = The disk firmware is a down version.
- qualfail = The disk is not supported.
- diskfail = The disk is in a failed state.
- notallflashdisk = The disk does not match the All-Flash Optimized personality of the system.

`[-firmware-revision <text>] - Firmware Revision`

Selects information about disks that have the specified firmware revision level.

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

Selects information about disks that have the specified home node.

`[-home-id <nvrmaid>] - Home System ID`

Selects information about disks whose home node has the specified system ID.

`[-host-adapter <text>] - Primary Path Host Adapter`

Selects information about disks that are currently using the specified Host Bus Adapter.

`[-hw-minimum-os <text>] - Hardware Minimum Supported Data ONTAP Version`

Selects information about disks that have the specified hardware minimum supported Data ONTAP version.

`[-import-in-progress {true|false}] - Foreign LUN import in progress`

Selects information about the array LUNs that are currently being imported

`[-initiator <text>,...] - Initiator Port`

Selects information about disks that are visible to the initiator specified. Disks that are not currently in use by that initiator are included.

`[-initiator-iops <integer>,...] - Number of IOPS on Initiator (Rolling Average)`

Selects information about disks that are visible to an initiator that has executed the specified number of IOPs.

[-initiator-io-kbps <integer>,...] - Kbytes of I/O per second on Initiator (Rolling Average)

Selects information about disks visible to an initiator that has executed I/O at the specified throughput.

[-initiator-lun-in-use-count <integer>,...] - Number of LUNs in the in-use state on this initiator

Selects information about disks with a path through an initiator that has the specified in-use-count.

[-initiator-side-switch-port <text>,...] - Initiator Side Switch Port

Selects information about disks that are visible to an initiator connected to the specified switch port.

[-is-multidisk-carrier {true|false}] - Multi Disk Carrier?

Selects information about disks that are located within a multi-disk carrier.

[-is-local-attach {true|false}] - Indicates If the Disk Is Local to This Cluster

Selects information about disks attached to the local(true) or remote(false) MetroCluster site.

[-location {<nodename>|local}] - Physical Location

Selects information about disks attached to the specified node.

[-location-id <nvramid>] - The system ID of the node where the disk is attached

Selects information about disks attached to the node with the specified system ID.

[-lun <integer>,...] - LUN ID

Selects information about the specified LUNs.

[-lun-iops <integer>,...] - Number IOPS per second on disk (Rolling Average)

Selects information about the LUNs that have reached the specified number of IOPs.

[-lun-io-kbps <integer>,...] - Kbytes/sec on Disk (Rolling Average)

Selects information about the LUNs that have reached the specified I/O throughput.

[-lun-path-use-state <text>,...] - The Use State of the LUN on this path

Selects information about LUNs reporting the specified in-use state.

[-model <text>] - Model

Selects information about disks of the specified model.

[-nodelist {<nodename>|local}] - Controller name

Selects information about disks that are visible to all of the specified nodes .

[-outage-reason <text>] - Outage Reason

Selects information about disks that are not in service for the specified reason. Possible values are: admin failed, admin removed, admin testing, evacuated, bad label, bypassed, failed, init failed, label version, labeled broken, labelmaint, LUN resized, missing, not responding, predict failure, rawsize shrank, recovering, sanitizing, sanitized, SnapLock Disk, testing, unassigned, unknown.

[-path-error-count <integer>] - Path Error Count

Selects information about disks that are visible on a path that has incurred the specified number of errors.

[-path-iops <integer>,...] - Number of IOPS on Path (Rolling Average)

Selects information about disks on those paths that have reached the specified number of IOPs.

[-path-io-kbps <integer>,...] - Kbytes of I/O per second on Path (Rolling Average)

Selects information about disk with paths that have reached the specified I/O throughput

[-path-link-errors <integer>,...] - Link Error count on path

Selects information about disks with paths that have incurred the specified number of FC link errors.

[-path-lun-in-use-count <integer>,...] - Number of LUNs in the in-use state on this path

Selects information about disks with paths that have the specified in-use-count.

[-path-quality <integer>,...] - Percentage of weighted error threshold

Selects information about disks on paths that have incurred the specified number of errors. The value displayed is a measure of the health of a path expressed as a percentage of an error threshold. Once a path has reached or surpassed the error threshold, another path will be selected for I/O transfer, if there is one available.

[-physical-size-mb <integer>] - Physical Size (MB)

Selects information about disks that have the specified physical capacity, in megabytes.

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

Selects information about disks that have the specified physical capacity, in human readable units.

[-physical-size-512b <integer>] - Physical Size in Units of 512 Bytes

Selects information about disks that have the specified physical capacity, in 512-byte chunks. This parameter is present only for backwards compatibility with Data ONTAP 8.0.

[-plex <text>] - Plex Name

Selects information about disks that belong to the specified RAID plex.

[-pool <text>] - Assigned Pool

Selects information about disks that belong to the specified SyncMirror pool (pool0 or pool1).

[-port-speed <text>,...] - Port Speed

Selects information about disks that are served by a Host Bus Adapter that is running at the specified port speed.

[-position <diskPositionType>] - Disk Position

Selects information about disks that have the specified position within their disk container.

[-power-on-hours <integer>] - Hours Powered On

Selects information about disks that have the specified number of hours being powered up.

[-prefailed {true|false}] - Marked for Rapid RAID Recovery?

Selects information about disks that match the specified parameter value indicating whether the disk is either awaiting or is in process of Rapid RAID Recovery.

[-preferred-target-port {true|false}] - Whether or not target port group is preferred (privilege: advanced)

Selects information about disks that match the specified parameter value indicating whether the backing storage is ALUA (Asymmetric Logical Unit Access) capable and has specified the array target port on this path to be a preferred target port for I/O.

[-primary-port <text>] - Primary Path Disk Port

Selects information about disks that use the specified primary port.

[-raid-group <text>] - Raid Group Name

Selects information about disks that belong to the specified RAID group.

[-reconstruction-percent <integer>] - Percentage of Reconstruction Complete

Selects information about disks that are being reconstructed and that have the specified percentage of the reconstruction operation completed.

[-replacing {true|false}] - Being Replaced?

Selects information about disks that match the specified boolean value indicating whether the disk is either awaiting or in process of disk replacement.

[-reservation-key <text>] - Reservation Key

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

[-reservation-type {rs|we|re|ea|sa|wero|earo|wear|ear|none}] - Reservation Type

If this parameter is specified, the command displays information only about the disk or disks that have the specified persistent reservation type. Possible values are: rs, we, re, ea, sa, wero, earo, wear, ear, or none.

[-reserver-id <integer>] - Reservation System ID

Selects information about disks that are reserved by the node with the specified system ID.

[-root-owner {<nodename>|local}] - Owner of Root Partition of Root-Data/Root-Data1-Data2 Partitioned Disk

Selects information about disks that have the specified root partition owner name. Used with root-data/root-data1-data2 partitioned disks.

[-root-owner-id <nvramid>] - Owner System ID of Root Partition of Root-Data/Root-Data1-Data2 Partitioned Disk

Selects information about disks that have the specified root partition owner system ID. Used with root-data/root-data1-data2 partitioned disks.

[-root-home {<nodename>|local}] - Home Owner of Root Partition of Root-Data/Root-Data1-Data2 Partitioned Disk

Selects information about disks that have the specified root partition home owner name. Used with root-data/root-data1-data2 partitioned disks.

[-root-home-id <nvramid>] - Home Owner System ID of Root Partition of Root-Data/Root-Data1-Data2 Partitioned Disk

Selects information about disks that have the specified root partition home owner system ID. Used with root-data/root-data1-data2 partitioned disks.

[-rpm <integer>] - Revolutions Per Minute

Selects information about disks that have the specified rotational speed.

[-secondary-name <disk path name>] - Secondary Path Name

Selects information about disks that use the specified secondary path name, for multipath configuration.

[-secondary-port <text>] - Secondary Path Disk Port

Selects information about disks that use the specified secondary port.

[-serial-number <text>] - Serial Number

Selects information about the disk that has the specified serial number.

[-storage-pool <text>] - Storage Pool Name

Selects information about disks that belong to the specified SSD storage pool.

[-shelf <integer>] - Shelf

Selects information about disks that are located within the specified shelf.

[-shelf-uid <text>] - Shelf UID

Selects information about disks that are located within a shelf with the specified Shelf UID.

[-slot <integer>] - Slot

Selects information about disks that are located in a drawer with the specified slot.

[-stack-id <integer>] - Stack ID

A cluster unique id for a collection of one or more interconnected shelves.

[-target-iops <integer>,...] - Number of IOPS to Target (Rolling Average)

Selects information about disks that are visible on target ports that have performed the specified number of IOPs.

[-target-io-kbps <integer>,...] - Kbytes of I/O per second to Target (Rolling Average)

Selects information about disks that are visible on target ports that have reached the specified I/O throughput.

[-target-lun-in-use-count <integer>,...] - Number of LUNs in the in-use state on this target

Selects information about disks with a path through a target port that has the specified in-use-count.

[-target-port-access-state <text>,...] - Failover optimization type

Selects information about disks that are visible on target ports that have the specified access state.

[-target-side-switch-port <text>,...] - Target Side Switch Port

Selects information about disks that are visible on target ports identified by the switch port to which they are connected.

[-target-wwpn <text>,...] - Target Port

Selects information about disks that are visible on target ports identified by their World Wide Port Name.

[-tpgn <integer>,...] - Target Port Group Number

Selects information about disks that belong to the specified Target Port Group Number.

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

Selects information about disks that have the specified disk type.

[-usable-size-mb <integer>] - Usable Size (MB)

Selects information about disks that have the specified usable space, in megabytes.

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

Selects information about disks that have the specified usable space, in human readable units.

[-vendor <text>] - Vendor Name

Selects information about disks that have the specified vendor.

[-vmdisk-device-id <integer>,...] - Virtual Disk Device ID

Selects information about disks that have the specified virtual disk device ID.

[-zeroed {true|false}] - Zeroed?

Selects information about disks that have (true) or have not (false) been fully pre-zeroed.

[-zeroing-percent <integer>] - Percentage of Zeroing Complete

Selects information about disks that are zeroing and have the specified percentage complete.

[-carrier-serialno <text>] - Carrier Serial Number

Selects information about disks that are located within the multi-disk carrier specified by the serial number.

[-vmdisk-target-address <text>] - Target Address of the VM Disk

Displays the VM Disk's target address either in the form of bus target lun or bus unit.

[-path-count <integer>] - Number of Paths to Disk

Selects information about disks that have the specified number of paths.

Examples

The following example displays information about all disks:


```
cluster1::> storage disk show
```

| Disk Owner | Usable Size | Shelf | Bay | Container Type | Position | Aggregate |
|---------------|----------------|-------|-----|-------------------|----------|-------------|
| ----- | ----- | ----- | --- | ----- | ----- | ----- |
| 1.1.1 | 10GB | 1 | 1 | spare | present | - |
| node1 | | | | | | |
| 1.1.4 | 78.59GB | 1 | 4 | spare | present | - |
| node1 | | | | | | |
| 1.1.12 | 10GB | 1 | 12 | spare | present | - |
| node1 | | | | | | |
| 1.2.12 | 10GB | 2 | 12 | broken | present | - |
| node1 | | | | | | |
| 1.3.7 | 78.59GB | 3 | 7 | aggregate | parity | aggr0_u23 |
| node1 | | | | | | |
| 1.1.6 | 78.59GB | 1 | 6 | broken | present | - |
| node1 | | | | | | |
| 1.2.10 | 78.59GB | 2 | 10 | aggregate | dparity | aggr0_u23 |
| node1 | | | | | | |
| 1.4.9 | 78.59GB | 4 | 9 | aggregate | data | aggr0_u23 |
| node1 | | | | | | |
| 1.1.0 | 10GB | 1 | 0 | aggregate | dparity | aggr0_u22 |
| node2 | | | | | | |
| 1.4.1 | 10GB | 4 | 1 | aggregate | data | dp_degraded |
| node2 | | | | | | |
| 1.1.2 | 10GB | 1 | 2 | spare | present | - |
| node2 | | | | | | |
| 1.1.3 | 20GB | 1 | 3 | spare | present | - |
| node2 | | | | | | |
| 1.4.4 | 20GB | 4 | 4 | spare | present | - |
| node2 | | | | | | |
| 1.4.6 | 10GB | 4 | 6 | aggregate | data | dp_sdc |
| node2 | | | | | | |
| 1.1.5 | 268.0GB | 1 | 5 | maintenance | present | - |
| node2 | | | | | | |
| 1.3.0 | 10GB | 3 | 0 | aggregate | parity | aggr0_u22 |
| node2 | | | | | | |
| 1.4.11 | 10GB | 4 | 11 | spare | present | - |
| node2 | | | | | | |
| 1.4.13 | 20GB | 4 | 13 | broken | present | - |
| node2 | | | | | | |
| [...] | | | | | | |

The following example displays detailed information about a disk named 1.0.75

```

cluster1::> storage disk show -disk 1.0.75
          Disk: 1.0.75
    Container Type: spare
      Owner/Home: node2 / node2
        DR Home: -
Stack ID/Shelf/Bay: 1 / 0 / 75
      LUN: 0
      Array: N/A
      Vendor: NETAPP
      Model: X267_HKURO500SSX
      Serial Number: ZAKAS0GH
      UID:
1FF17846:0A419201:9325845A:3ABD5075:00000000:00000000:00000000:00000000:00
000000:00000000
          BPS: 512
    Physical Size: 10.15GB
      Position: present
Checksum Compatibility: block
      Aggregate: -
      Plex: -

Paths:

          LUN  Initiator Side          Target Side
Link
Controller      Initiator      ID  Switch Port          Switch Port
Acc Use  Target Port          TPGN  Speed          I/O KB/s
IOPS
-----
node1          0d          0  N/A          N/A
AO  INU  220a000a3384e4d2          21  2 Gb/S          0
0
node1          0c          0  N/A          N/A
AO  RDY  2209000a3384e4d2          62  2 Gb/S          0
0
node2          0d          0  N/A          N/A
AO  INU  2209000a3384e4d2          62  2 Gb/S          3
0

Errors:
-

```

The following example displays RAID-related information about disks used in an aggregate:

```
cluster1::> storage disk show -raid-info-for-aggregate
```

```

Owner Node: node1
  Aggregate: aggr0_node1_0
    Plex: plex0
      RAID Group: rg0

```

```

Usable Physical
  Position Disk
  HA Shelf Bay Chan Pool  Type
RPM   Size   Size
-----
-----
15000  data    2.11.2    2d    11    2 B    Pool0  SAS
      9.77GB  9.93GB
      dparity 2.11.0    2d    11    0 B    Pool0  SAS
15000  9.77GB  9.93GB
      parity  2.11.1    2d    11    1 B    Pool0  SAS
15000  9.77GB  9.93GB

```

```

Owner Node: node2
  Aggregate: a1
    Plex: plex0
      RAID Group: rg0

```

```

Usable Physical
  Position Disk
  HA Shelf Bay Chan Pool  Type
RPM   Size   Size
-----
-----
7200   data    2.1.8     2a     1    8 B    Pool0  BSAS
      9.77GB  9.91GB
      dparity 2.1.6     2a     1    6 B    Pool0  BSAS
7200   9.77GB  9.91GB
      parity  2.1.7     2a     1    7 B    Pool0  BSAS
7200   9.77GB  9.91GB

```

```

Owner Node: node2
  Aggregate: a1
    Plex: plex0
      RAID Group: rg1

```

```

Usable Physical
  Position Disk
  HA Shelf Bay Chan Pool  Type
RPM   Size   Size
-----
-----
7200   data    2.1.11    2a     1   11 B    Pool0  BSAS
      9.77GB  9.91GB
      dparity 2.1.9     2a     1    9 B    Pool0  BSAS
7200   9.77GB  9.91GB

```

```

        parity    2.1.10                2a    1  10  B    Pool0  BSAS
7200    9.77GB    9.91GB
Owner Node: node2
  Aggregate: aggr0
    Plex: plex0
      RAID Group: rg0

Usable Physical
      Position Disk
RPM    Size    Size                HA Shelf Bay Chan Pool  Type
-----
-----
        data      2.1.5                2a    1   5  B    Pool0  BSAS
7200    9.71GB   10.03GB
        dparity   2.1.2                2a    1   2  B    Pool0  BSAS
7200    9.71GB   10.03GB
        parity    2.1.4                2a    1   4  B    Pool0  BSAS
7200    9.71GB   10.03GB
12 entries were displayed.

```

The following example displays RAID-related information about spares:

```

cluster1::> storage disk show -spare
Original Owner: node1
Checksum Compatibility: block

Physical
Disk          HA Shelf Bay Chan  Pool  Type  RPM  Size  Usable
Size Owner
-----
-----
1.1.23        0b      1  23    A  Pool0  FCAL  10000  132.8GB
134.2GB node1
1.1.25        0b      1  25    A  Pool0  FCAL  10000  132.8GB
133.9GB node1
1.1.26        0b      1  26    A  Pool1  FCAL  10000  132.8GB
133.9GB node1
1.1.27        0b      1  27    A  Pool1  FCAL  10000  132.8GB
134.2GB node1
Home Owner: node2
Checksum Compatibility: block

Physical
Disk          HA Shelf Bay Chan  Pool  Type  RPM  Size  Usable
Size Owner
-----
-----
1.1.19        0a      1  19    B  Pool1  FCAL  10000  132.8GB
133.9GB node2
1.1.20        0a      1  20    B  Pool0  FCAL  10000  132.8GB
133.9GB node2
1.1.21        0a      1  21    B  Pool0  FCAL  10000  132.8GB
133.9GB node2
[...]
```

The following example displays RAID-related information about broken disks:

```
cluster1::> storage disk show -broken
```

```
Original Owner: node1
```

```
Checksum Compatibility: block
```

```
Usable Physical
```

| Disk Size | Size | Outage Reason | HA | Shelf | Bay | Chan | Pool | Type | RPM |
|--------------|------|---------------|----|-------|-----|------|------|------|-----|
|--------------|------|---------------|----|-------|-----|------|------|------|-----|

| | | | | | | | | | |
|-------|---------|---------------|----|---|---|---|-------|------|-------|
| 1.1.0 | 132.8GB | admin failed | 0b | 1 | 0 | A | Pool0 | FCAL | 10000 |
| 1.2.6 | 133.9GB | admin removed | 0b | 2 | 6 | A | Pool1 | FCAL | 10000 |

| | | | | | | | | | |
|-------|---------|--|--|--|--|--|--|--|--|
| 1.2.6 | 134.2GB | | | | | | | | |
|-------|---------|--|--|--|--|--|--|--|--|

```
Original Owner: node2
```

```
Checksum Compatibility: block
```

```
Usable Physical
```

| Disk Size | Size | Outage Reason | HA | Shelf | Bay | Chan | Pool | Type | RPM |
|--------------|------|---------------|----|-------|-----|------|------|------|-----|
|--------------|------|---------------|----|-------|-----|------|------|------|-----|

| | | | | | | | | | |
|-------|---------|--------------|----|---|---|---|-------|------|-------|
| 1.1.0 | 132.8GB | admin failed | 0a | 1 | 0 | B | Pool0 | FCAL | 10000 |
|-------|---------|--------------|----|---|---|---|-------|------|-------|

| | | | | | | | | | |
|--------|---------|---------------|----|---|----|---|-------|------|-------|
| 1.1.13 | 133.9GB | admin removed | 0a | 1 | 13 | B | Pool0 | FCAL | 10000 |
|--------|---------|---------------|----|---|----|---|-------|------|-------|

```
4 entries were displayed.
```

The following example displays RAID-related information about disks in maintenance center:

```
cluster1::> storage disk show -maintenance
```

```
Original Owner: node1
```

```
Checksum Compatibility: block
```

```
Usable Physical
```

| Disk | Outage Reason | HA Shelf | Bay | Chan | Pool | Type | RPM |
|------|---------------|----------|-----|------|------|------|-----|
| Size | Size | | | | | | |

| | | | | | | | |
|-----------------|------------------|---|---|---|-------|------|-------|
| 1.1.8 | admin testing 0b | 1 | 8 | A | Pool0 | FCAL | 10000 |
| 132.8GB 133.9GB | | | | | | | |

| | | | | | | | |
|-----------------|------------------|---|----|---|-------|------|-------|
| 1.2.11 | admin testing 0b | 2 | 11 | A | Pool1 | FCAL | 10000 |
| 132.8GB 134.2GB | | | | | | | |

```
Original Owner: node2
```

```
Checksum Compatibility: block
```

```
Usable Physical
```

| Disk | Outage Reason | HA Shelf | Bay | Chan | Pool | Type | RPM |
|------|---------------|----------|-----|------|------|------|-----|
| Size | Size | | | | | | |

| | | | | | | | |
|-----------------|------------------|---|----|---|-------|------|-------|
| 1.2.10 | admin testing 0a | 2 | 10 | B | Pool1 | FCAL | 10000 |
| 132.8GB 133.9GB | | | | | | | |

| | | | | | | | |
|-----------------|------------------|---|----|---|-------|------|-------|
| 1.2.13 | admin testing 0a | 2 | 13 | B | Pool1 | FCAL | 10000 |
| 132.8GB 134.2GB | | | | | | | |

```
4 entries were displayed.
```

The following example displays partition-related information about disks:

```
cluster1::> storage disk show -partition-ownership
```

| Disk ID | Partition | Home | Owner | Home ID | Owner |
|----------|-----------|--------------------|--------------------|------------|------------|
| VMw-1.13 | Container | pvaruncluster-2-01 | pvaruncluster-2-01 | 4087518786 | 4087518786 |
| VMw-1.14 | Container | pvaruncluster-2-01 | pvaruncluster-2-01 | 4087518786 | 4087518786 |
| VMw-1.15 | Container | pvaruncluster-2-01 | pvaruncluster-2-01 | 4087518786 | 4087518786 |
| | Root | pvaruncluster-2-01 | pvaruncluster-2-01 | 4087518786 | 4087518786 |
| | Data | pvaruncluster-2-01 | pvaruncluster-2-01 | 4087518786 | 4087518786 |
| VMw-1.16 | Container | pvaruncluster-2-01 | pvaruncluster-2-01 | 4087518786 | 4087518786 |
| | Root | pvaruncluster-2-01 | pvaruncluster-2-01 | 4087518786 | 4087518786 |
| | Data1 | pvaruncluster-2-01 | pvaruncluster-2-01 | 4087518786 | 4087518786 |
| | Data2 | pvaruncluster-2-01 | pvaruncluster-2-01 | 4087518786 | 4087518786 |
| VMw-1.17 | Container | pvaruncluster-2-01 | pvaruncluster-2-01 | 4087518786 | 4087518786 |
| VMw-1.18 | Container | pvaruncluster-2-01 | pvaruncluster-2-01 | 4087518786 | 4087518786 |
| | Root | - | - | - | - |
| | Data | pvaruncluster-2-01 | pvaruncluster-2-01 | 4087518786 | 4087518786 |
| VMw-1.19 | Container | pvaruncluster-2-01 | pvaruncluster-2-01 | 4087518786 | 4087518786 |
| | Root | pvaruncluster-2-01 | pvaruncluster-2-01 | 4087518786 | 4087518786 |
| | Data1 | - | - | - | - |
| | Data2 | pvaruncluster-2-01 | pvaruncluster-2-01 | 4087518786 | 4087518786 |

Related Links

- [storage aggregate show-status](#)

storage disk unfail

Unfail a broken disk

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

Description

The `storage disk unfail` command can be used to unfail a broken disk.

If the attempt to unfail the disk is unsuccessful, the disk remains in the broken state.

The disk unfail command prompts for confirmation unless you specify the `"-quiet"` parameter.

Parameters

`-disk <disk path name>` - Disk Name (privilege: advanced)

This parameter specifies the disk to be unfailed.

`[-s, -spare <true>]` - Make the Disk Spare (privilege: advanced)

This parameter specifies whether the unfailed disk will be made a spare disk. The disk is forced to become a spare disk if this parameter is specified.

If this parameter is not specified, the disk is brought back into its parent aggregate. Setting this parameter might result in the aggregate coming back online if it is not complete or online. The default value is `false`.

`[-q, -quiet <true>]` - Confirmations off (privilege: advanced)

You can set this parameter to `true` to suppress the confirmation messages. However, before proceeding with the command, you should be aware that the confirmation message contains important information about the effect of unfailling a disk. This command cannot be reversed once it is invoked. The default value is `false`.

Examples

The following example unfails a disk named `1.1.16` to become a spare disk:

```
cluster1::*> storage disk unfail -disk 1.1.16 -spare
```

Warning: Failed disk "1.1.16" may have aggregate labels and file system data present. In that case use of this command to release this disk as a spare disk, will lead to permanent loss of any file system data that may be on this disk, and may prevent recovery of a failed aggregate. Are you sure you want to continue with disk unfail? {y|n}:

storage disk updatefirmware

(DEPRECATED) - Update disk firmware

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

Description



This command is deprecated and may be removed in a future release of Data ONTAP. Use the `"storage disk firmware update"` command.

The `storage disk updatefirmware` command updates the firmware on one or more disks.

You can download the latest firmware by using the [storage firmware download](#) command.

You can specify a list of one or more disks whose firmware is to be updated by using the `-disk` parameter, or you can update the firmware on all local disks by omitting the `-disk` parameter.

Parameters

[`-disk <disk path name>,...`] - Disk

This specifies the disk or disks whose firmware is to be updated.

If you do not specify this option, all local disks' firmware is updated.

Examples

The following example updates the firmware on all disks:

```
cluster1::> storage disk updatefirmware
```

Related Links

- [storage disk firmware update](#)
- [storage firmware download](#)

storage disk zerospares

Zero non-zeroed spare disks

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

Description

The `storage disk zerospares` command zeroes all non-zeroed spare disks in all nodes or a specified node in the cluster. A node must be online to zero disks. This operation must be done before a disk can be reused in another aggregate. This version of ONTAP uses fast zeroing, which converts a spare disk from non-zeroed to zeroed without the long wait times required when physically zeroing a disk.

Parameters

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

If this parameter is specified, only non-zeroed spares assigned to the specified node will be zeroed.

Otherwise, all non-zeroed spares in the cluster will be zeroed.

Examples

The following example zeroes all non-zeroed spares owned by a node named `node4`, using fast zeroing:

```
cluster1::> storage disk zerospares -owner node4
```

storage disk error show

Display disk component and array LUN configuration errors.

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

Description

The `storage disk error show` command displays disk component and array LUN configuration errors.

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.

[-uid <text>] - UID

Displays the error information of the disk whose unique ID matches the value you specify. A disk unique identifier has the form:

```
`20000000:875D4C32:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000`
```

[-array-name <array name>] - Array Name

Displays the errors of the storage array whose name you specified.

[-node {<nodename>|local}] - Controller Name

Displays the error information for the disks on the clustered node whose name you specified.

[-disk <disk path name>] - Disk

Displays detailed error information about the disk you specified.

[-serial-number <text>] - Serial Number

Displays the error information for the disk whose serial number you specified.

[-error-id <integer>,...] - Error ID

Displays the error information for the disks whose Error IDs match IDs you specified.

[-error-type

{onepath|onedomain|control|foreign|toobig|toosmall|invalidblocksize|targetasymmap|deviceasymmap|failovermisconfig|unknown|netapp|fwdownrev|qualfail|diskfail|notallflashdisk}] - Error Type

Displays all disk errors of the error types you specified, grouped by type.

- onepath = The array LUN is accessible only via a single path.
- onedomain = The array LUN is accessible only via a single fault domain.
- control = The array LUN cannot be used because it is a control device.

- **foreign** = The array LUN is marked as foreign and has some external SCSI reservations other than those from Data ONTAP.
- **toobig** = The array LUN exceeds the maximum array LUN size that Data ONTAP supports.
- **toosmall** = The array LUN is less than the minimum array LUN size that Data ONTAP supports.
- **invalidblocksize** = The array LUN is not a valid block size.
- **targetasymmap** = The array LUN is presented more than once on a single target port.
- **deviceasymmap** = The array LUN is presented with multiple IDs.
- **failovermisconfig** = The array LUN is configured with inconsistent failover methods.
- **unknown** = The array LUN from a storage array that is not supported by this version of Data ONTAP.
- **netapp** = A SAN front-end LUN from one Data ONTAP system that is presented as external storage to another Data ONTAP system.
- **fwdownrev** = The disk firmware is a down version.
- **qualfail** = The disk is not supported.
- **diskfail** = The disk is in a failed state.
- **notallflashdisk** = The disk does not match the All-Flash Optimized personality of the system.

Examples

The following example displays configuration errors seen in the system:

```
cluster1::> storage disk error show
Disk           Error Type           Error Text
-----
1.02.0         qualfail           This disk failed dynamic disk
qualification. Update the Disk Qualification Package.
```

storage disk firmware revert

Revert disk firmware

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

Description

The `storage disk firmware revert` command reverts firmware on all disks or a specified list of disks on a node.

You can specify a list of one or more disks whose firmware is to be reverted by using the `-disk` parameter.

You can revert the firmware on all the disks owned by a node by using the `-node` parameter.

This command can make the disks inaccessible for up to five minutes after the start of its execution. Therefore, the network sessions that use the concerned node must be terminated before running the `storage disk firmware revert` command. This is particularly true for CIFS sessions that might be terminated when this

command is executed.

If you need to view the current firmware versions, use the [storage disk show](#)-fields`*firmware-revision* command. The following example displays partial output from the [storage disk show](#)-fields`*firmware-revision* command, where the firmware version for the disks is NA02:

```
cluster1::> storage disk show -fields firmware-revision
disk      firmware-revision
-----
1.0.0     NA02
1.0.1     NA02
1.0.2     NA02
1.0.3     NA02
1.0.4     NA02
1.0.5     NA02
```

The firmware files are stored in the `/mroot/etc/disk_fw` directory on the node. The firmware file name is in the form of "product-ID.revision.LOD". For example, if the firmware file is for Seagate disks with product ID X225_ST336704FC and the firmware version is NA01, the file name is X225_ST336704FC.NA01.LOD. If the node in this example contains disks with firmware version NA02, the `/mroot/etc/disk_fw/X225_ST336704FC.NA01.LOD` file is downloaded to every disk when you execute this command.

How to Revert the Firmware for an HA Pair in a Cluster

Use the following procedure to perform a revert on the disks in an HA environment:

- Make sure that the nodes are not in takeover or giveback mode.
- Download the latest firmware on both nodes by using the [storage firmware download](#) command.
- Revert the disk firmware on Node A's disks by entering the `storage disk firmware revert-node`node-A`` command.
- Wait until the `storage disk firmware revert` command completes on Node A, and then revert the firmware on Node B's disks by entering the `storage disk firmware revert-node`node-B`` command.

Parameters

{ -disk <disk path name>, ... - Disk Name (privilege: advanced) }

Specifies the disk or disks whose firmware is to be reverted.

| -node {<nodename>|local} - Node Name (privilege: advanced) }

Specifies the node name. The disk firmware will be reverted on all the disks owned by the node specified by this parameter.

Examples

- The following example reverts the firmware on all disks owned by cluster-node-01:

```
cluster1::*> storage disk firmware revert -node cluster-node-01
```

Warning: Disk firmware reverts can be disruptive to the system. Reverts involve

power cycling all of the affected disks, as well as suspending disk

I/O to the disks being reverted. This delay can cause client disruption. Takeover/giveback operations on a high-availability (HA)

group will be delayed until the firmware revert process is complete.

Disk firmware reverts should only be done one node at a time. Disk firmware reverts can only be performed when the HA group is healthy;

they cannot be performed if the group is in takeover mode.

Do you want to continue with disk firmware reverts? {y|n}: y

Info: Reverting disk firmware for disks on cluster-node-01.

- The following example reverts the firmware on disk 1.5.0 which is owned by node cluster-node-04:

```
cluster1::*> storage disk firmware revert -disk 1.5.0
```

Warning: Disk firmware reverts can be disruptive to the system. Reverts involve

power cycling all of the affected disks, as well as suspending disk

I/O to the disks being reverted. This delay can cause client disruption. Takeover/giveback operations on a high-availability (HA)

group will be delayed until the firmware revert process is complete.

Disk firmware reverts should only be done one node at a time. Disk firmware reverts can only be performed when the HA group is healthy;

they cannot be performed if the group is in takeover mode.

Do you want to continue with disk firmware reverts? {y|n}: y

Info: Reverting disk firmware for disks on cluster-node-04.

Related Links

- [storage disk show](#)
- [storage firmware download](#)

storage disk firmware show-update-status

Display disk firmware update status.

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

Description

The `storage disk firmware show-update-status` command displays the state of the background disk firmware update process.

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 (privilege: advanced)

Selects the node that matches this parameter value.

[-num-waiting-download <integer>] - The Number of Disks Waiting to Download (privilege: advanced)

Selects the nodes whose number of disks waiting to download by the BDFU process matches this parameter value.

[-total-completion-estimate <integer>] - Estimated Duration to Completion (mins) (privilege: advanced)

Selects the nodes whose Background Disk Firmware Update (BDFU) completion time estimate matches this parameter value. This indicates the amount of estimated time required for BDFU to complete the firmware update cycle.

[-average-duration-per-disk <integer>] - Average Firmware Update Duration per Disk (secs) (privilege: advanced)

Selects the nodes whose BDFU reports the average time required to update a single disk matches this parameter value. This indicates the average amount of time required by each disk drive.

[-unable-to-update <disk path name>,...] - List of Disks with a Failed Update (privilege: advanced)

Selects the nodes whose unable to update disk list matches this parameter value. This is a list of disks that failed to update the firmware.

[`-update-status {off|running|idle}`] - Background Disk Firmware Update Status (privilege: advanced)

Selects the nodes whose BDFU process status matches this parameter value. Possible values are:

- off - The BDFU process is off.
- running - The BDFU process is on and currently running.
- idle - The BDFU process is on and is currently idle.

Examples

```
cluster1::*> storage disk firmware show-update-status
```

| | | Number | Average | Total | |
|--------|---------|----------|-------------|------------|-----------|
| | Update | Waiting | Duration | Completion | |
| Node | State | Download | /Disk (Sec) | Est. (Min) | Unable to |
| Update | | | | | |
| ----- | ----- | ----- | ----- | ----- | |
| node1 | running | 2 | 120 | 4 | 1.3.3 |
| node2 | idle | 0 | 120 | 0 | - |
| node3 | off | 0 | 120 | 0 | - |

storage disk firmware update

Update disk firmware

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

Description

Use the `storage disk firmware update` command to manually update firmware on all disks or a specified list of disks on a node. However, the recommended way to update disk firmware in a cluster is to enable automatic background firmware update by enabling the `-bkg-firmware-update` parameter for all of the nodes in the cluster. You can do this by entering the [storage disk option modify](#) `-node`*-bkg-firmware-update`on` command.

You can download the latest firmware on the node by using the [storage firmware download](#) command.

You can specify a list of one or more disks whose firmware is to be updated by using the `-disk` parameter.

You can update the firmware on all the disks owned by a node by using the `-node` parameter.

This command can make the disks inaccessible for up to five minutes after the start of its execution. Therefore, the network sessions that use the concerned node must be terminated before running the `storage disk firmware update` command. This is particularly true for CIFS sessions that might be terminated when this command is executed.

The firmware is automatically downloaded to disks, which report previous versions of the firmware. For information on automatic firmware update downloads, see "Automatic versus Manual Firmware Download".

If you need to view the current firmware versions, use the `storage disk show -fields `firmware-revision`` command. The following example displays partial output from the `storage disk show -fields `firmware-revision`` command, where the firmware version for the disks is NA01:

```
cluster1:> storage disk show -fields firmware-revision
disk      firmware-revision
-----
1.0.0     NA01
1.0.1     NA01
1.0.2     NA01
1.0.3     NA01
1.0.4     NA01
1.0.5     NA01
```

The firmware files are stored in the `/mroot/etc/disk_fw` directory on the node. The firmware file name is in the form of "product-ID.revision.LOD". For example, if the firmware file is for Seagate disks with product ID X225_ST336704FC and the firmware version is NA02, the filename is X225_ST336704FC.NA02.LOD. The revision part of the file name is the number against which the node compares each disk's current firmware version. If the node in this example contains disks with firmware version NA01, the `/mroot/etc/disk_fw/X225_ST336704FC.NA02.LOD` file is used to update every eligible disk when you execute this command.

Automatic versus Manual Firmware Download

The firmware is automatically downloaded to those disks that report previous versions of firmware following a system boot or disk insertion. Note that:

- A manual download is a disruptive operation that makes disks inaccessible for up to five minutes after the download is started. Network sessions that use the node must be terminated before running the `storage disk firmware update` command.
- The firmware is not automatically downloaded to the node's partner node in an HA pair.
- The firmware is not automatically downloaded to unowned disks on nodes configured to use software-based disk ownership.
- The `bkg-firmware-update` parameter controls how the automatic firmware download feature works:
 - If the `bkg-firmware-update` parameter is set to off, then the `storage disk firmware update` will update the firmware on the drives in parallel.
 - If the `bkg-firmware-update` parameter is set to on, then the `storage disk firmware update` will update spares and filesystem disks in a nondisruptive manner in the background after boot. Firmware downloads for these disks will be done sequentially by temporarily taking them offline one at a time for the duration of the download. After the firmware is updated, the disk will be brought back online and restored to its normal operation.

During an automatic download to an HA environment, the firmware is not downloaded to the disks owned by the HA partner.

When you use the `storage disk firmware update` command, the firmware is:

- Updated on every disk regardless of whether it is on the A-loop, the B-loop, or in an HA environment.

- If the node is configured in a software-based disk ownership system, only disks owned by this node are updated.

During an automatic firmware download in a MetroCluster™ environment, the firmware is not downloaded to the disks owned by the partner cluster. During both manual and automatic firmware download in a MetroCluster-over-IP environment, the firmware is not downloaded to any remote disks located at the partner cluster while Disaster Recovery is in progress.

Follow the instructions in "How to Update the Firmware for an HA Pair in a Cluster" to ensure that the updating process is successful. Data ONTAP supports redundant path configurations for disks in a non-HA configuration. The firmware is automatically downloaded to disks on the A-loop or B-loop of redundant configurations that are not configured in an HA pair and are not configured to use software-based disk ownership.

Automatic Background Firmware Update

The firmware can be updated in the background so that the firmware update process does not impact the clients. This functionality is controlled with the `bkg-firmware-update` parameter. You can modify the parameter by using the CLI `storage disk option modify -node `node_name` bkg-firmware-update `on|off` command. The default value for this parameter is "on".

When disabled or set to "off", `storage disk firmware update` will update the firmware in automated mode. This means that all disks which had older firmware revision will be updated regardless of whether they are spare or filesystem disks.

When enabled or set to "on", the background `storage disk firmware update` will update firmware in automated mode only on disks that can be successfully taken offline from active filesystem RAID groups and from the spare pool. To ensure a faster boot process, the firmware is not downloaded to spares and filesystem disks at boot time.

This provides the highest degree of safety available, without the cost of copying data from each disk in the system twice. Disks are taken offline one at a time and then the firmware is updated on them. The disk is brought online after the firmware update and a mini/optimized reconstruct happens for any writes, which occurred while the disk was offline. Background disk firmware update will not occur for a disk if its containing RAID group or the volume is not in a normal state (for example, if the volume/plex is offline or the RAID group is degraded). However, due to the continuous polling nature of background disk firmware update, firmware updates will resume after the RAID group/plex/volume is restored to a normal mode. Similarly, background disk firmware updates are suspended for the duration of any reconstruction within the system.

How to Update the Firmware for an HA Pair in a Cluster

The best way to update the firmware in a cluster with HA pairs is to use automatic background firmware update by enabling the option `bkg-firmware-update` parameter for each node. Enable the `-bkg-firmware-update` parameter on all the nodes by entering the `storage disk option modify -node `node_name` bkg-firmware -update `on` command. Alternatively, use the following procedure to successfully perform a manual update on the disks in an HA environment:

- Make sure that the nodes are not in takeover or giveback mode.
- Download the latest firmware on both the nodes by using the `storage firmware download` command.
- Install the new disk firmware on Node A's disks by entering the `storage disk firmware update -node `node-A`` command.
- Wait until the `storage disk firmware update` command completes on Node A, and then install the new disk firmware on Node B's disks by entering the `storage disk firmware update -node `node-`

B` command.

Parameters

{ -disk <disk path name>, ... - Disk (privilege: advanced)

Specifies the disk or disks whose firmware is to be updated.

| -node {<nodename>|local} - node (privilege: advanced) }

Specifies the node name. The disk firmware will be updated on all the disks owned by the node specified by this parameter.

Examples

- The following example updates the firmware on all disks owned by cluster-node-01:

```
cluster1::*> storage disk firmware update -node cluster-node-01
```

Warning: Disk firmware updates can be disruptive to the system. Updates involve

power cycling all of the affected disks, as well as suspending disk

I/O to the disks being updated. This delay can cause client disruption. Takeover/giveback operations on a high-availability (HA)

group will be delayed until the firmware update process is complete.

Disk firmware updates should only be done one node at a time.

Disk

firmware updates can only be performed when the HA group is healthy;

they cannot be performed if the group is in takeover mode.

Do you want to continue with disk firmware updates? {y|n}: y

Info: Updating disk firmware for disks on cluster-node-01.

- The following example updates the firmware on disk 1.5.0 which is owned by node cluster-node-04:

```
cluster1::*> storage disk firmware update -disk 1.5.0
```

Warning: Disk firmware updates can be disruptive to the system. Updates involve

power cycling all of the affected disks, as well as suspending disk

I/O to the disks being updated. This delay can cause client disruption. Takeover/giveback operations on a high-availability

(HA) group will be delayed until the firmware update process is complete.

Disk firmware updates should only be done one node at a time.

Disk firmware updates can only be performed when the HA group is healthy;

they cannot be performed if the group is in takeover mode.

Do you want to continue with disk firmware updates? {y|n}: y

Info: Updating disk firmware for disks on cluster-node-04.

Related Links

- [storage disk option modify](#)
- [storage firmware download](#)
- [storage disk show](#)

storage disk option modify

Modify disk options

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

Description

The `storage disk option modify` command modifies the background firmware update setting, automatic copy setting, controls automatic disk assignment of all disks assigned to a specified node, or modifies the policy of automatic disk assignment of unowned disks.

Parameters

-node {<nodename>|local} - Node

This parameter specifies the node that owns the disks whose options are to be modified.

[-bkg-firmware-update {on|off}] - Background Firmware Update

This parameter specifies whether firmware updates run as a background process. The default setting is `on`, which specifies that firmware updates to spare disks and file system disks is performed nondisruptively via

a background process. If the option is turned off, automatic firmware updates occur at system startup or during disk insertion.

[`-autocopy {on|off}`] - Auto Copy

This parameter specifies whether data is to be automatically copied from a failing disk to a spare disk in the event of a predictive failure. The default setting is `on`. It is sometimes possible to predict a disk failure based on a pattern of recovered errors that have occurred. In such cases, the disk reports a predictive failure. If this option is set to `on`, the system initiates Rapid RAID Recovery to copy data from the failing disk to an available spare disk. When data is copied, the disk is marked as failed and placed in the pool of broken disks. If a spare is not available, the node continues to use the disk until it fails. If the option is set to `off`, the disk is immediately marked as failed and placed in the pool of broken disks. A spare is selected and data from the missing disk is reconstructed from other disks in the RAID group. The disk does not fail if the RAID group is already degraded or is being reconstructed. This ensures that a disk failure does not lead to the failure of the entire RAID group.

[`-autoassign {on|off}`] - Auto Assign

This parameter specifies whether automatic assignment of unowned disks is enabled or disabled. The default setting is `on`. This parameter is used to set both a node-specific and a cluster-wide disk option.

[`-autoassign-policy {default|bay|shelf|stack}`] - Auto Assignment Policy

This parameter defines the granularity at which auto assign should work. This option is ignored if the `-autoassign` option is off. Auto assignment can be done at the stack/loop, shelf, or bay level. The possible values for the option are `default`, `stack`, `shelf`, and `bay`. The default value is platform dependent. It is `stack` for all non-entry platforms and single-node systems, whereas it is `bay` for entry-level platforms.

Examples

The following example sets the background firmware update setting to `on` for all disks belonging to a node named `node0`:

```
cluster1::> storage disk option modify -node node0 -bkg-firmware-update on
```

The following example shows how to enable auto assignment for the disks on `node1`:

```
cluster1::> storage disk option modify -node node1 -autoassign on
cluster1::> storage disk option show
```

| Node | BKg. FW. Upd. | Auto Copy | Auto Assign | Auto Assign Policy |
|-------|---------------|-----------|-------------|--------------------|
| node1 | on | on | on | default |
| node2 | on | on | off | default |

2 entries were displayed.

The following example shows how to modify the auto assignment policy on `node1`:

```
cluster1::> storage disk option modify -node node1 -autoassign-policy bay
cluster1::> storage disk option show
```

| Node | BKg. FW. Upd. | Auto Copy | Auto Assign | Auto Assign Policy |
|-------|---------------|-----------|-------------|--------------------|
| node1 | on | on | on | bay |
| node2 | on | on | off | default |

2 entries were displayed.

storage disk option show

Display a list of disk options

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

Description

The `storage disk option show` command displays the settings of the following disk options:

- Background firmware update
- Automatic copying of data to a spare disk in the event of a predictive failure
- Automatic assignment of disks
- Policy that governs automatic assignment of unowned disks

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

Selects the node that owns the disks. If this parameter is not specified, the command displays information about the disk options on all the nodes.

[-bkg-firmware-update {on|off}] - Background Firmware Update

Selects the disks that match this parameter value.

[-autocopy {on|off}] - Auto Copy

Selects the disks that match this parameter value.

[-autoassign {on|off}] - Auto Assign

Displays the auto assignment status of unowned disks. The default value is `on`.

[`-autoassign-policy {default|bay|shelf|stack}`]] - Auto Assignment Policy

Selects the disks that match the automatic assignment policy value:

- Default
- Stack/loop
- Shelf
- Bay

Examples

The following example displays disk-option settings for disks owned by all nodes in the cluster:

```
cluster1::> storage disk option show
Node          BKg. FW. Upd.  Auto Copy      Auto Assign  Auto Assign
Policy
-----
node0          on              on              on           default
node1          on              on              on           stack
node2          on              on              on           bay
node3          on              on              on           bay
4 entries were displayed.
```

storage dqp commands

storage dqp show

Display Disk Qualification Package details

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

Description

The `storage dqp show` command displays information about the Disk Qualification Package (DQP). The command displays the following information:

- Node Name
- Package Date
- File Version
- File Name
- Drive Record Count
- Drive Alias Record Count
- Device Class Record Count
- System Class Record Count

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

Displays DQP information on the specified node.

[-package-date <text>] - Package Date

Selects DQP information with the specified package date.

[-file-version <text>] - File Version

Selects DQP information with the specified file version.

[-filename <text>] - File Name

Selects DQP information with the specified file name.

[-drive-records <integer>] - Drive Record Count

Selects DQP information with the specified drive record count.

[-alias-records <integer>] - Alias Record Count

Selects DQP information with the specified drive alias record count.

[-device-records <integer>] - Device Record Count

Selects DQP information with the specified device class record count.

[-system-records <integer>] - System Record Count

Selects DQP information with the specified system class record count.

Examples

The following example displays the brief version of Disk Qualification Package details:

```
cluster1::> storage dqp show
Node                Package Date      Version
-----
cluster1-01        20180907         3.17
cluster1-02        20180907         3.17
2 entries were displayed.
```

The following example displays the `-instance` version of Disk Qualification Package details:


```
cluster1::> storage dqp show -instance
Node: cluster1-01
    Package Date: 20180907
    Version: 3.17
    File Name: qual_devices_v3
    Drive Record Count: 626
    Alias Record Count: 196
    Device Record Count: 28
    System Record Count: 3
Node: cluster1-02
    Package Date: 20180907
    Version: 3.17
    File Name: qual_devices_v3
    Drive Record Count: 626
    Alias Record Count: 196
    Device Record Count: 28
    System Record Count: 3
2 entries were displayed.
```

storage encryption commands

storage encryption disk destroy

Cryptographically destroy a self-encrypting disk

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

Description

The `storage encryption disk destroy` command cryptographically destroys a self-encrypting disk (SED), making it incapable of performing I/O operations. This command performs the following operations:

- Employs the inherent erase capability of SEDs to cryptographically sanitize the disk
- Permanently locks the disk to prevent further data access
- Changes the data and FIPS authentication keys to random values that are not recorded except within the SED.

Use this command with extreme care. The only mechanism to restore the disk to usability (albeit without the data) is the [storage encryption disk revert-to-original-state](#) operation that is available only on disks that have the physical secure ID (PSID) printed on the disk label.

The destroy command requires you to enter a confirmation phrase before proceeding with the operation.

The command releases the cluster shell after launching the operation. Monitor the output of the [storage encryption disk show-status](#) command for command completion.

Upon command completion, remove the destroyed SED from the system.

Parameters

-disk <disk path name> - Disk Name

This parameter specifies the name of the disk you want to cryptographically destroy. See the man page for the `storage disk modify` command for information about disk-naming conventions.

[-force-all-states <true>] - Destroy All Matching Disks

When this parameter is *false* or not specified, the operation defaults to spare and broken disks only, as reported in the output of the `storage disk show` command. When you specify this parameter as *true*, it allows you to cryptographically destroy all matching disk names regardless of their state, including those in active use in aggregates. This allows a quick destroy of all system disks if you use the `-disk` parameter with the asterisk wildcard (*). If you destroy active disks, the nodes might not be able to continue operation, and might halt or panic.

Examples

The following command cryptographically destroys the disk 1.10.20:

```
cluster1::> storage encryption disk destroy 1.10.20
```

```
Warning: This operation will cryptographically destroy 1 spare or broken  
self-encrypting disks on 1 node.
```

```
You cannot reuse destroyed disks unless you revert  
them to their original state using the PSID value.
```

```
To continue, enter
```

```
destroy disk
```

```
:destroy disk
```

```
Info: Starting destroy on 1 disk.
```

```
View the status of the operation by using the  
"storage encryption disk show-status" command.
```

```
cluster1::>
```

If you do not enter the correct confirmation phrase, the operation is aborted:

```
cluster1::> storage encryption disk destroy 1.10.2*
```

Warning: This operation will cryptographically destroy 5 spare or broken self-encrypting disks on 1 node.

You cannot reuse destroyed disks unless you revert them to their original state using the PSID value.

To continue, enter

destroy disk

:yes

No disks destroyed.

```
cluster1::>
```

The following command quickly cryptographically destroys all system disks, including those in active use in aggregates and shared devices:

```
cluster1::> storage encryption disk destroy -force-all-states -disk *
```

Warning: This operation will cryptographically destroy 96 self-encrypting disks on 4 nodes.

To continue, enter

destroy disk

:destroy disk

Info: Starting destroy on 96 disks.

View the status of the operation by using the

xref:{relative_path}storage-encryption-disk-show-status.html[storage encryption disk show-status] command.

```
cluster1::>
```

Related Links

- [storage encryption disk revert-to-original-state](#)
- [storage encryption disk show-status](#)
- [storage disk show](#)

storage encryption disk modify

Modify self-encrypting disk parameters

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

Description

The `storage encryption disk modify` command changes the data protection parameters of self-encrypting disks (SEDs) and FIPS-certified SEDs (FIPS SEDs); it also modifies the FIPS-compliance AK (FIPS AK) of FIPS SEDs. The current data AK and FIPS AK of the device are required to effect changes to the respective AKs and FIPS compliance. The current and new AKs must be available from the key servers or onboard key management.

The command releases the cluster shell after launching the operation. Monitor the output of the [storage encryption disk show-status](#) command for command completion.



To properly protect data at rest on a FIPS SED and place it into compliance with its FIPS certification requirements, set both the Data and FIPS AKs to a value other than the device's default key; depending on the device type, the default may be manufacture secure ID (MSID), indicated by a key ID with the special value `0x0`, or a null key represented by a blank AK. Verify the key IDs by using the [storage encryption disk show](#) and `storage encryption disk show-fips` commands.

Parameters

`-disk <disk path name>` - Disk Name

This parameter specifies the name of the SED or FIPS SED that you want to modify.

{ `[-data-key-id <text>]` - Key ID of the New Data Authentication Key

This parameter specifies the key ID associated with the data AK that you want the SED to use for future authentications. When the provided key ID is the MSID, data at rest on the SED is not protected from unauthorized access. Setting this parameter to a non-MSID value automatically engages the power-on-lock protections of the device, so that when the device is power-cycled, the system must authenticate with the device using the AK to reenale I/O operations. You cannot specify the null default key; use MSID instead.

| `[-fips-key-id <text>]` - Key ID of the New Authentication Key for FIPS Compliance }

This parameter specifies the key ID associated with the FIPS AK that you want the FIPS SED to apply to SED credentials other than the one that protects the data. When the value is not the MSID, these credentials are changed to the indicated AK, and other security-related items are set to conform to the FIPS certification requirements ("FIPS compliance mode") of the device. You may set the `-fips-key-id` to any one of the key IDs known to the system. The FIPS key ID may, but does not have to, be the same as the data key ID parameter. Setting `-fips-key-id` to the MSID key ID value disables FIPS compliance mode and restores the FIPS-related authorities and other components as required (other than data) to their default settings. A nonMSID FIPS-compliance key may be applied only to a FIPS SED.

Examples

The following command changes both the AK and the power-cycle protection to values that protect the data at rest on the disk. Note that the `-data-key-id` and `-fips-key-id` parameters require one of the key IDs that appear in the output of the [security key-manager query](#) command.

```
cluster1::> storage encryption disk modify -data-key-id
6A1E21D800000000010000000000000F5A1EB48EF26FD6A8E76549C019F2350 -disk
2.10.*
```

Info: Starting modify on 14 disks.
View the status of the operation by using the
storage encryption disk show-status command.

The following command changes the FIPS AK and sets the device into FIPS-compliance mode. Note that the `-fips-key-id` parameter requires one of the key IDs that appear in the output of the [security key-manager query](#) command.

```
cluster1::> storage encryption disk modify -fips-key-id
6A1E21D8000000000100000000000005A1FB4EE8F62FD6D8AE6754C9019F35A 2.10.*
```

Info: Starting modify on 14 disks.
View the status of the operation by using the
storage encryption disk show-status command.

Related Links

- [storage encryption disk show-status](#)
- [storage encryption disk show](#)
- [security key-manager query](#)

storage encryption disk revert-to-original-state

Revert a self-encrypting disk to its original, as-manufactured state

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

Description

Some self-encrypting disks (SEDs) are capable of an operation that restores them as much as possible to their as-manufactured state. The `storage encryption disk revert-to-original-state` command invokes this special operation that is available only in SEDs that have the physical secure ID (PSID) printed on their labels.

The PSID is unique to each SED, meaning the command can revert only one SED at a time. The disk must be in a "broken" or "spare" state as shown by the output of the [storage disk show](#) command.

The operation in the SED accomplishes the following changes:

- Sanitizes all data by changing the disk encryption key to a new random value
- Sets the data authentication key (AK) and FIPS AK to the default values
- Resets the data locking controls

- Resets the power-on lock state to *false*
- Initializes other vendor-unique encryption-related parameters

The command releases the cluster shell after launching the operation. Monitor the output of the [storage encryption disk show-status](#) command for command completion.

When the operation is complete, it is possible to return the SED to service using the [storage disk unfail](#) command in *advanced* privilege mode. To do so, you might also need to reestablish ownership of the SED using the [storage disk assign](#) command.

Parameters

-disk <disk path name> - Disk Name

The name of the SED to be reverted to its as-manufactured state. See the man page for the `storage disk modify` command for information about disk-naming conventions.

-psid <text> - Physical Secure ID

The PSID printed on the SED label.

Examples

The following command shows a SED being returned to its as-manufactured state:

```
cluster1::> storage encryption disk revert-to-original-state -disk 01.10.0
-psid AC65PYF8CG45YZABUQJKM98WV2VZGRLD
```

Related Links

- [storage disk show](#)
- [storage encryption disk show-status](#)
- [storage disk unfail](#)
- [storage disk assign](#)

storage encryption disk sanitize

Cryptographically sanitize a self-encrypting disk

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

Description

The `storage encryption disk sanitize` command cryptographically sanitizes one or more self-encrypting disks (SEDs), making the existing data on the SED impossible to retrieve. This operation employs the inherent erase capability of SEDs to perform all of the following changes:

- Sanitizes all data by changing the disk encryption key to a new random value
- Sets the data authentication key (AK) to the default AK (manufacture secure ID/MSID or null, depending on the device type)

- Unlocks the data band
- Resets the power-on lock state to *false*

There is no method to restore the disk encryption key to its previous value, meaning that you cannot recover the data on the SED. Use this command with extreme care.

The sanitize command requires you to enter a confirmation phrase before proceeding with the operation.

The command releases the cluster shell after launching the operation. Monitor the output of the [storage encryption disk show-status](#) command for command completion.

When the operation is complete, it is possible to return the SED to service using the [storage disk unfail](#) command in *advanced* privilege mode. To do so, you might also need to reestablish ownership of the SED using the [storage disk assign](#) command.

Parameters

-disk <disk path name> - Disk Name

This parameter specifies the name of the SEDs you want to cryptographically sanitize. See the man page for the `storage disk modify` command for information about disk-naming conventions.

[-force-all-states <true>] - Sanitize All Matching Disks

When this parameter is *false* or not specified, the operation defaults to spare and broken disks only, as reported in the output of the [storage disk show](#) command. When you specify this parameter as *true*, it allows you to cryptographically sanitize all matching disk names regardless of their state, including those in active use in aggregates. This allows a quick erasure of all system data if you use the `-disk` parameter with the asterisk wildcard (*). If you sanitize active disks, the nodes might not be able to continue operation, and might halt or panic.

Examples

The following command sanitizes the disk 1.10.20:

```
cluster1::> storage encryption disk sanitize 1.10.20
```

```
Warning: This operation will cryptographically sanitize 1 spare or broken
self-encrypting disk on 1 node.
```

```
To continue, enter
```

```
sanitize disk
```

```
:sanitize disk
```

```
Info: Starting sanitize on 1 disk.
```

```
View the status of the operation using the
```

```
xref:{relative_path}storage-encryption-disk-show-status.html[storage
encryption disk show-status] command.
```

```
cluster1::>
```

If you do not enter the correct confirmation phrase, the operation is aborted:

```
cluster1::> storage encryption disk sanitize 1.10.2*
```

```
Warning: This operation will cryptographically sanitize 5 spare or broken  
self-encrypting disks on 1 node.
```

```
To continue, enter
```

```
sanitize disk
```

```
:yes
```

```
No disks sanitized.
```

```
cluster1::>
```

The following command quickly cryptographically sanitizes all system disks, including those in active use in aggregates and shared devices:

```
cluster1::> storage encryption disk sanitize -force-all-states -disk *
```

```
Warning: This operation will cryptographically sanitize 96  
self-encrypting disks on 4 nodes.
```

```
To continue, enter
```

```
sanitize disk
```

```
:sanitize disk
```

```
Info: Starting sanitize on 96 disks.
```

```
View the status of the operation by using the
```

```
xref:{relative_path}storage-encryption-disk-show-status.html[storage  
encryption disk show-status] command.
```

```
cluster1::>
```

Related Links

- [storage encryption disk show-status](#)
- [storage disk unfail](#)
- [storage disk assign](#)
- [storage disk show](#)

storage encryption disk show-status

Display status of disk encryption operation

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

Description

The `storage encryption disk show-status` command displays the results of the latest destroy ,

modify, or sanitize operation of the storage encryption disk command family. Use this command to view the progress of these operations on self-encrypting disks (SEDs).

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

If you specify this parameter, the command displays disk encryption status for the nodes that match this parameter.

[-is-fips-support {true|false}] - Node Supports FIPS Disks

If you specify this parameter, the command displays disk encryption status for the nodes that match this parameter (*true* means the node supports FIPS-certified self-encrypting drives).

[-latest-op <Storage Disk Encryption Operation>] - Latest Operation Requested

If you specify this parameter, the command displays disk encryption status for the nodes with a most recent storage encryption disk operation that matches this parameter (one of `destroy`, `modify`, `revert-to-original-state`, `sanitize`, or `unknown`).

[-op-start-time <MM/DD/YYYY HH:MM:SS>] - Operation Start Time

Selects the nodes with operation start times that match this parameter.

[-op-execute-time <integer>] - Execution Time in Seconds

If you specify this parameter, the command displays disk encryption status for the nodes with operation execution time that matches this parameter. The operation may be partial or completed.

[-disk-start-count <integer>] - Number of Disks Started

If you specify this parameter, the command displays disk encryption status for the nodes that started this number of SEDs in their latest operation.

[-disk-done-count <integer>] - Number of Disks Done

Selects the nodes that report this number of SEDs having completed the latest operation, successfully or not.

[-disk-success-count <integer>] - Number of Disks Successful

If you specify this parameter, the command displays disk encryption status for the nodes that report this number of SEDs that successfully completed the latest operation. When the operation is finished, if the success count is not the same as the started count, some additional detail is available using the `-instance` or `-node` parameters.

[-disk-no-key-id-count <integer>] - Number of Disks with Key ID Not Found

If you specify this parameter, the command displays disk encryption status for the nodes that report this number of SEDs that failed the latest operation because Data ONTAP could not find the Key IDs associated with the required authentication key of the SED.

`[-disk-no-authent-count <integer>]` - Number of Disks Not Authenticated

If you specify this parameter, the command displays disk encryption status for the nodes that report this number of SEDs that failed the latest operation because the identified Authentication Key could not authenticate with the SED.

`[-op-sequence-count <integer>]` - Sequence Count of Latest Operation

If you specify this parameter, the command displays disk encryption status for that nodes that match the value list.

Examples

When no operation has been requested since node boot, the status for that node is empty. If you enter a node name, the output is in the same format as for the `-instance` parameter.

```
cluster1::> storage encryption disk show-status -node node
Node Name: node
  Node Supports FIPS-certified Self-Encrypting Disks: true
    Latest Operation Requested: unknown
      Operation Start Time: -
        Execution Time in Seconds: -
          Number of Disks Started: -
            Number of Disks Done: -
              Number of Disks Successful: -
                Number of Disks with Key ID Not Found: -
                  Number of Disks Not Authenticated: -
```

Once an operation begins, the status is dynamic until all devices have completed. When disks are modified, sanitized, or destroyed, sequential executions of `storage encryption disk show-status` appear as in this example that shows the progress of a modify operation on three SEDs on each node of a two-node cluster:

```
cluster1::> storage encryption disk show-status
```

| | SED | Latest | Start | Execution | Disks | Disks |
|------------|---------|---------|--------------------|------------|-------|-------|
| Disk | | | | | | |
| Node | Support | Request | Timestamp | Time (sec) | Begun | Done |
| Successful | | | | | | |
| node0 | true | modify | 9/22/2014 13:58:53 | 4 | 3 | 0 |
| node10 | true | modify | 9/22/2014 13:58:53 | 4 | 3 | 0 |


```
cluster1::> storage encryption disk show-status
```

| | SED | Latest | Start | Execution | Disks | Disks |
|------------|---------|---------|--------------------|------------|-------|-------|
| Disk | | | | | | |
| Node | Support | Request | Timestamp | Time (sec) | Begun | Done |
| Successful | | | | | | |
| node3 | true | modify | 9/22/2014 13:58:53 | 7 | 3 | 3 |
| node13 | true | modify | 9/22/2014 13:58:53 | 7 | 3 | 3 |

storage encryption disk show

Display self-encrypting disk attributes

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

Description

The `storage encryption disk show` command displays information about encrypting drives. When no parameters are specified, the command displays the following information about all encrypting drives:

- Disk name
- The protection mode of the device
- The key ID associated with the data authentication key ("data AK")

In MetroCluster systems, the information is valid from the cluster that owns the drive, or from the DR cluster when in switchover mode. If information is not available, perform the `show` command from the cluster partner.

You can use the following parameters together with the `-disk` parameter to narrow the selection of displayed drives or the information displayed about them.

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.

| [-fips]

If you specify this parameter, the command displays the key ID associated with the FIPS-compliance authentication key ("FIPS AK") instead of the data key ID.

| [-instance] }

If you specify this parameter, the command displays detailed disk information about all disks, or only those specified by a `-disk` parameter.

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

If you specify this parameter, the command displays information about the specified disks. If you specify a single disk path name, the output is the same as when you use the `-instance` parameter. See the man page for the `storage disk modify` command for information about disk-naming conventions. Default is all self-encrypting disks.

[-container-name <text>] - Container Name

This parameter specifies the container name associated with an encrypting drive. If you specify an aggregate name or other container name, only the encrypting drives in that container are displayed. See the man page for the [storage disk show](#) command for a description of the container name. Use the [storage aggregate show-status](#) and [storage disk show](#) commands to determine which aggregates the drives are in.

[-container-type {aggregate | broken | foreign | labelmaint | maintenance | mediator | remote | shared | spare | unassigned | unknown | unsupported}] - Container Type

This parameter specifies the container type associated with an encrypting drive. If you specify a container type, only the drives with that container type are displayed. See the man page for the [storage disk show](#) command for a description of the container type.

[-data-key-id <text>] - Key ID of the Current Data Authentication Key

This parameter specifies the key ID associated with the data AK that the encrypting drive requires for authentication with its data-protection authorities. The special key ID `0x0` indicates that the current data AK of the drive is the default manufacture secure ID (MSID) that is not secret. Some devices employ an initial null default AK that appears as a blank data-key-id; you cannot specify a null data-key-id value. To properly protect data at rest on the device, modify the data AK using a key ID that is not a default value (MSID or null). When you modify the data AK with a non-MSID key ID, the system automatically sets the device's power-on lock enable control so that authentication with the data AK is required after a device power-cycle. Use [storage encryption disk modify-data-key-id`key-id](#) to protect the data. Use [storage encryption disk modify-fips-key-id`key-id](#) to place the drives into FIPS-compliance mode.

[-fips-key-id <text>] - Key ID of the Current FIPS Authentication Key

This parameter specifies the key ID associated with the FIPS authentication key ("FIPS AK") that the system must use to authenticate with FIPS-compliance authorities in FIPS-certified drives. This parameter may not be set to a non-MSID value in drives that are not FIPS-certified.

[-is-power-on-lock-enabled {true|false}] - Is Power-On Lock Protection Enabled?

This parameter specifies the state of the control that determines whether the encrypting drive requires

authentication with the data AK after a power-cycle. The system enables this control parameter automatically when you use the [storage encryption disk modify](#)-data-key-id command to set the data AK to a value other than the default AK. Data is protected only when this parameter is *true* and the data AK is not a default. Compare with the values of the -protection-mode parameter below.

[-protection-mode <text>] - Mode of SED Data and FIPS-Compliance Protection

The protection mode that the drive is in:

- open - data is unprotected; drive is not in FIPS-compliance mode
- data - data is protected; drive is not in FIPS-compliance mode
- part - data is unprotected; drive is otherwise in FIPS-compliance mode
- full - data is protected; drive is in FIPS-compliance mode
- miss - protection mode information is not available

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

This parameter selects the drive type to include in the output.

Examples

The following command displays information about all encrypting drives:

```
cluster1:> storage encryption disk show
Disk      Mode Data Key ID
-----
-----
0.0.0     open 0x0
0.0.1     part 0x0
0.0.2     data
0A9C9CFC00000000010000000000000345CFD1BAD310CA8EDB377D439FB5C9A
1.10.0    open
0A53ED2A00000000010000000000000BEDC1B27AD3F0DB8891375AED2F34D0B
1.10.1    part
0A9C9CFC00000000010000000000000345CFD1BAD310CA8EDB377D439FB5C9A
1.10.2    full
0A9C9CFC00000000010000000000000345CFD1BAD310CA8EDB377D439FB5C9A
[...]
```

Note in the example that only disk 1.10.2 is fully protected with FIPS mode, power-on-lock enable, and an AK that is not the default MSID or a null key.

The following command displays information about the protection mode and FIPS key ID for all encrypting drives:

```
cluster1::> storage encryption disk show -fips
Disk      Mode FIPS-Compliance Key ID
-----
-----
0.0.0    open 0x0
0.0.1    part
0A53ED2A000000000100000000000000C1B27AD3F0DB8891375AED2F34D0BBED
0.0.2    data 0x0
1.10.0   open
0A53ED2A000000000100000000000000BEDC1B27AD3F0DB8891375AED2F34D0B
1.10.1   part
0A9C9CFC000000000100000000000000345CFD1BAD310CA8EDB377D439FB5C9A
1.10.2   full
0A9C9CFC000000000100000000000000BEDC1B27AD3F0DB8891375AED2F34D0B
[...]
```

Note again that only disk 1.10.2 is fully protected with FIPS-compliance mode set, power-on-lock enabled, and a data AK that is not the default MSID or a null key.

The following command displays the individual fields for disk 1.10.2:

```
cluster1::> storage encryption disk show -disk 1.10.2
Disk Name: 1.10.2
                                Container Name: aggr0
                                Container Type: shared
                                Is Drive FIPS-certified?: true
Key ID of the Current Data Authentication Key:
0A9C9CFC000000000100000000000000345CFD1BAD310CA8EDB377D439FB5C9A
Key ID of the Current FIPS Authentication Key:
0A9C9CFC000000000100000000000000BEDC1B27AD3F0DB8891375AED2F34D0B
                                Is Power-On Lock Protection Enabled?: true
                                Mode of Data and FIPS-Compliance Protection: open
```

Related Links

- [storage disk show](#)
- [storage aggregate show-status](#)
- [storage encryption disk modify](#)

storage errors commands

storage errors show

Display storage configuration errors.

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

Description

The `storage errors show` command displays configuration errors with back end storage arrays.

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.

[-uid <text>] - UID

Selects the disks that match this parameter value.

[-array-name <array name>] - Array Name

Selects the disks that have the specified name for the storage array that is connected to the cluster.

[-node {<nodename>|local}] - Controller Name

Selects the disks that match this parameter value.

[-disk <disk path name>] - Disk

Selects the disks that match this parameter value.

[-serial-number <text>] - Serial Number

Selects the disks that match this parameter value.

[-error-id <integer>,...] - Error ID

Selects the disks with error-id values that match this parameter value.

[-error-type

{onepath|onedomain|control|foreign|toobig|toosmall|invalidblocksize|targetasymmap|deviceasymmap|failovermisconfig|unknown|netapp|fwdownrev|qualfail|diskfail|notallflashdisk}] - Error Type

Selects the disks with error types values that match this parameter value.

Examples

The following example displays configuration errors seen in the system:

```
cluster1::> storage errors show
Disk: vnv3070f20b:vnci9124s54:1-24.126L23
-----
vnci9124s54:1-24.126L23 (600a0b800019e999000036b24bac3983): This array LUN
reports an invalid block size and is not usable. Only a block size of 512
is supported.
```

storage failover commands

storage failover giveback

Return failed-over storage to its home node

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

Description

The `storage failover giveback` command returns storage that has failed over to a node's partner back to the home node. This operation fails if other resource-intensive operations (for instance, system dumps) are running and make the giveback operation potentially dangerous or disruptive. Some options are available only at the advanced privilege level and higher. Run the `storage failover show-giveback` command to check the status of giveback operations.

NOTE:

- If the system ID of the partner has changed while the node is in takeover mode, the `storage failover giveback` command updates the ownership of the partner's disks to the new system ID while giving back.
- If the giveback operation fails due to the operation being vetoed by a subsystem, check the syslog or EMS output for a subsystem-specific reason for the abort. The corrective action is subsystem-specific and is detailed in the corrective action portion of the message. Follow the corrective action specified by the subsystem and then reissue the `storage failover giveback` command. If you cannot perform the corrective action, then use the `override-vetoes` option in the `storage failover giveback` command to force the giveback.
- If the giveback operation fails because the node cannot communicate with its partner, check the EMS output for the corrective action. Follow the corrective action and then reissue the `storage failover giveback` command. If you cannot perform the corrective action, then use the `-require-partner -waiting false` option in the `storage failover giveback` command to force the giveback. This parameter is available only at the advanced privilege level and higher.
- If the node does not receive notification that the partner has brought online the given-back aggregate and its volumes, the `storage failover show-giveback` command displays the giveback status for the aggregate as failed. A possible reason for this failure is that the partner is overloaded and slow in bringing the aggregate online. Run the `storage aggregate show` command to verify that the aggregate and its volumes are online on the partner node. The node will not attempt the giveback operation for remaining aggregates. To force the giveback, use the `-require-partner-waiting false` option in the `storage failover giveback` command. This parameter is available only at the advanced privilege level and higher.

Parameters

{ -ofnode {<nodename>|local} - Node to which Control is Givenback

Specifies the node whose storage is currently taken over by its partner and will be given back by the giveback operation.

| -fromnode {<nodename>|local} - Node Initiating Giveback }

Specifies the node that currently holds the storage that is to be returned to the partner node.

[-require-partner-waiting {true|false}] - Require Partner in Waiting (privilege: advanced)

If this optional parameter is used and set to false, the storage is given back regardless of whether the partner node is available to take back the storage or not. If this parameter is used and set to true, the storage will not be given back if the partner node is not available to take back the storage. If this parameter is not used, the behavior defaults to the setting of the `-check-partner` option set with the [storage failover modify](#) command.

[-override-vetoes <true>] - Override All Vetoes

If this optional parameter is used, the system overrides veto votes during a giveback operation. If this parameter is not used, the system does not proceed with a giveback if it is vetoed. This parameter, if used, can only be set to true.

[-only-cfo-aggregates <true>] - Giveback Only CFO Aggregates

If this optional parameter is used, giveback of only the CFO aggregates (root aggregate and CFO style data aggregates) will be attempted. If this parameter is not used, giveback of all the aggregates (CFO and SFO aggregates) will be attempted. This parameter, if used, can only be set to true.

Examples

The following example gives back storage that is currently held by a node named node1. The partner must be available for the giveback operation to occur.

```
node::> storage failover giveback -fromnode node1
```

The following example gives back only the CFO aggregates to a node named node2 (the aggregates are currently held by a node named node1). The partner must be available for the giveback operation to occur, and the veto-giveback process can be overridden.

```
node::> storage failover giveback -ofnode node2  
-override-vetoes true -only-cfo-aggregates true
```

Related Links

- [storage failover show-giveback](#)
- [storage aggregate show](#)
- [storage failover modify](#)

storage failover modify

Modify storage failover attributes

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

Description

The `storage failover modify` command changes the storage-failover options for a node. Some options are available only at the advanced privilege level and higher.

Parameters

-node {<nodename>|local} - Node

This specifies the node whose storage-failover options are to be modified.

{ [-enabled {true|false}] - Takeover Enabled

This optionally specifies whether storage failover is enabled. The default setting is `true`.

[[-mode {ha|non_ha}] - HA Mode }

This specifies whether the node is set up in high-availability mode or stand-alone mode. If the node is a member of a high-availability configuration, set the value to `ha`. If the node is stand-alone, set the value to `non_ha`. Before setting the HA mode, you must complete the platform dependent steps to set up the system in a stand-alone or HA configuration as shown in the documentation for your platform.

[-auto-giveback {true|false}] - Auto Giveback Enabled

This optionally specifies whether automatic giveback operations are enabled. An automatic giveback operation is invoked when one node of a failover pair is in takeover mode and the failed node is repaired and restarts. When the repaired node boots, the node in takeover mode detects this and initiates a giveback operation. The default setting is `true`. This parameter is not applicable if takeover was because of a disruption in the partner's operation. For those cases, use the `-auto-giveback-after-panic` parameter, instead.

[-detection-time <integer>] - Takeover Detection Time (secs)

This optionally specifies the amount of time, in seconds, that a node remains unresponsive before its partner initiates a takeover operation. Possible values range from 10 to 180 seconds. The default setting is 15 seconds.

[-onfailure {true|false}] - Takeover on Failure Enabled (privilege: advanced)

This optionally specifies whether the node automatically takes over for its partner node if the partner node fails. The default setting is `true`. This parameter is available only at the advanced privilege level and higher.

[-onpanic {true|false}] - Takeover on Panic Enabled

This optionally specifies whether the node automatically takes over for its partner node if the partner node panics. The default setting is `true`. Changing this parameter on one node automatically makes the same change on its partner node.

[-onshort-uptime {true|false}] - Takeover on Short Uptime Enabled (privilege: advanced)

This optionally specifies whether the node takes over for its partner node if the partner node fails within 60 seconds of starting up; the time period is modifiable by using the `-short-uptime` parameter. The default

setting is `true` . This parameter is available only at the advanced privilege level and higher.

`[-short-uptime <integer>]` - Short Uptime (secs) (privilege: advanced)

This optionally specifies the time period used by the `-onshort-uptime` parameter. The default setting is 60 seconds. This parameter is available only at the advanced privilege level and higher.

`[-attempts <integer>]` - Number of Giveback Attempts (privilege: advanced)

This optionally specifies the number of times the node attempts an automatic giveback operation within 60 minutes; the time period is modifiable by using the `-attempts-time` parameter. The default setting is 2 attempts. This parameter is available only at the advanced privilege level and higher.

`[-attempts-time <integer>]` - Giveback Attempts Period (minutes) (privilege: advanced)

This optionally specifies the time period used by the `-attempts` parameter. The default setting is 60 minutes. This parameter is available only at the advanced privilege level and higher.

`[-propagate {true|false}]` - Propagate Status via Mailbox (privilege: advanced)

This optionally specifies whether storage-failover status is communicated via mailbox disks. The default setting is `true` . This parameter is available only at the advanced privilege level and higher.

`[-read-interval <integer>]` - Node Status Read Interval (secs) (privilege: advanced)

This optionally specifies, in seconds, how frequently the node reads its partner node's status from the mailbox disks. The default setting is 5 seconds. This parameter is available only at the advanced privilege level and higher.

`[-write-interval <integer>]` - Node Status Write Interval (secs) (privilege: advanced)

This optionally specifies, in seconds, how frequently the node writes its status to the mailbox disks. The default setting is 5 seconds. This parameter is available only at the advanced privilege level and higher.

`[-onreboot {true|false}]` - Takeover on Reboot Enabled

This optionally specifies whether the node automatically takes over for its partner if the partner reboots. The default setting is `true` . Takeover can occur if the partner exceeds the expected time to reboot even when this option is set to `false`. The expected time to reboot is different for different platforms. The minimum expected time to reboot is 180 seconds. The `-inhibit-takeover` option of the [system node reboot](#) command overrides this option: if a node is rebooted with `-inhibit-takeover` set to `true` then takeover does not occur, even if the `takeover on reboot` option is `true`. If a node does takeover due to the partner rebooting, then it will automatically giveback after the reboot, even if the `-auto-giveback` option is set to `false` . This is non-persistent behavior: if the node does takeover due to partner reboot and then itself reboots (prior to giveback) then it will not automatically giveback if the `-auto-giveback` option is set to `false` .

`[-delay-seconds <integer>]` - Delay Before Auto Giveback (secs)

This optionally specifies the minimum time that a node will stay in takeover state prior to performing an automatic giveback. If the taken over node recovers quickly (for example, if the takeover was due to a reboot), by delaying the giveback for a few minutes the outage during the takeover and giveback can be reduced to two short outages instead of one longer one. The allowed range is 0 to 600, inclusive. The default setting is 600 seconds. This option affects all types of auto-giveback. This parameter is available only at the advanced privilege level and higher.



This delay does not affect manual giveback.

`[-hwassist {true|false}] - Hardware Assist Enabled`

This optionally specifies whether the hardware assist feature is enabled. If set to `true` this feature helps in fast takeover detection times in certain cases.

`[-hwassist-partner-ip <IP Address>] - Partner's Hwassist IP`

This optionally specifies the Ip address on which the partner node receives hardware assist alerts. For the hardware assist feature to be active, the value of this option should be equal to partner's node management Ip address.

`[-hwassist-partner-port <integer>] - Partner's Hwassist Port`

This optionally specifies the port number on which partner node listens to hardware assist alerts. It is recommended to have this value to be between 4000-4500. The default value is 4444.

`[-hwassist-health-check-interval <integer>] - Hwassist Health Check Interval (secs)`

This optionally specifies, in seconds, how frequently the hardware assist hardware on a node sends a heartbeat to its partner. The default value is 180.

`[-hwassist-retry-count <integer>] - Hwassist Retry Count`

This optionally specifies the number of times we repeat sending an hardware assist alert. The default value is 2.

`[-auto-giveback-after-panic {true|false}] - Auto Giveback After Takeover On Panic`

This optionally specifies whether a node should attempt automatic giveback operations if takeover was because of a disruption in the partner's operation. An automatic giveback operation is invoked when one node of a failover pair is in takeover mode and the failed node is repaired and restarts. When the repaired node boots, the node in takeover mode detects this and initiates a giveback operation automatically. The default setting is `true`.



This parameter is independent of the `-auto-giveback` parameter. If this parameter is enabled and the takeover is due to disruption in the partner's operation, giveback will be initiated even if `-auto-giveback` parameter is `false`.

`[-aggregate-migration-timeout <integer>] - Aggregate Migration Timeout (secs) (privilege: advanced)`

This optionally specifies the amount of time, in seconds, the source node has to wait for the destination node to complete the aggregate migration before declaring the migration as failed. The default setting is 120 seconds.

Examples

The following example enables the storage-failover service on a node named node0:

```
node::> storage failover modify -node node0 -enabled true
```

The following examples enable storage-failover takeover on a short uptime of 30 seconds on a node named node0:

```
node::*> storage failover modify -node node0 -onshort-uptime true -short
-uptime 30
```

Related Links

- [system node reboot](#)

storage failover show-giveback

Display giveback status

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

Description

The `storage failover show-giveback` command displays information about the giveback status of high-availability (HA) partner aggregates. The command displays the following information when no parameters are specified:

- Node name
- Partner aggregate name
- Giveback Status

You can specify additional parameters to display only the information that matches those parameters. For example, to display information only about a particular aggregate, run the command with the ``-aggregate aggregate_name`` 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.

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

If this parameter is used, the command displays information about the giveback status of the aggregates belonging to the HA partner of the specified node.

[-aggregate <text>] - Aggregate

If this parameter is used, the command displays information about the giveback status of the specified aggregate.

[`-giveback-status <text>,...`] - Aggregates Giveback State

If this parameter is used, the command displays information about the aggregates with the specified giveback status.

[`-destination <text>`] - Destination for Giveback

If this parameter is used, the command displays information about the giveback status of the aggregates whose destination after the giveback is the specified node.

Examples

The following example displays information about giveback status on all nodes:

```
node::> storage failover show-giveback
      Partner
Node      Aggregate      Giveback Status
-----
node0
      -                  No aggregates to give back
node1
      -                  No aggregates to give back
node2
      -                  No aggregates to give back
node3
      -                  No aggregates to give back
4 entries were displayed.
```

storage failover show-takeover

Display takeover status

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

Description

The `storage failover show-takeover` command displays information about the takeover status of nodes in a cluster. The command also displays the takeover status of aggregates being taken over. During each phase of takeover, the takeover node and the node being taken over display their takeover status and the status of the aggregates being taken over. The command displays the following information when no parameters are specified:

- Node name
- Node takeover status - This contains a descriptive information about the phase of takeover.
- Aggregate
- Aggregate takeover status - This contains the following information:
- Takeover status of the aggregate, such as "Done", "Failed", "In progress" and "Not attempted yet".

- Reason for an aggregate takeover failure.
- Corrective action, in case of an aggregate takeover failure.

You can specify additional parameters to display only the information that matches those parameters. For example, to display information only about a particular node, run the command with the `-node`_node_name_`` parameter.

Parameters

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

If this parameter is specified, the command displays the specified fields for all nodes, in column style output.

| [-instance] }

If this parameter is specified, the command displays the same detailed information as for the `-node` parameter, but for all nodes.

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

If this parameter is specified, the command displays information about the takeover status of the specified node, and the takeover status of the aggregates being taken over.

[-node-takeover-status <text>] - Node's Takeover Status

If this parameter is specified, the command displays information about the takeover status of the nodes with the specified node-takeover-status. The command also displays the takeover status of the aggregates belonging to the node being taken over.

[-aggregate <text>] - Aggregate Being Taken Over

If this parameter is specified, the command displays information about the takeover status of the specified aggregate, and the takeover status of the nodes containing the specified aggregate.

[-aggregate-takeover-status <text>] - Aggregate's Takeover Status

If this parameter is specified, the command displays information about the takeover status of the aggregates with the specified aggregate takeover status, and the takeover status of the nodes containing those aggregates.

Examples

The following example shows the takeover status of two nodes, nodeA and nodeB, in an High Availability (HA) pair, when both are in normal mode; neither node has taken over its HA partner. In this case, there is no takeover status for the aggregates.

```
cluster1::> storage failover show-takeover
```

| Node | Node Status | Aggregate | Takeover Status |
|-------|-------------------------|-----------|-----------------|
| nodeA | Takeover not attempted. | - | - |
| nodeB | Takeover not attempted. | - | - |

The following example shows the takeover status of two nodes, nodeA and nodeB, in an HA pair, when nodeA is in the SFO phase of an optimized takeover of nodeB. In this case, nodeA does not have information about the takeover status of nodeB's aggregates.

```
cluster1::> storage failover show-takeover
```

| Node | Node Status | Aggregate | Takeover Status |
|-------|--|----------------------------------|---|
| nodeA | Optimized takeover of partner in progress. Partner, ("nodeB"), is relocating its SFO aggregates. Run the command "storage failover show-takeover -node nodeB" to display the relocation status of the partner. | - | - |
| nodeB | Being taken over. | aggr1 aggr2 CFO aggregates | In progress, Module: backup. Not attempted yet Not attempted yet. |

The following example shows the takeover status of two nodes, nodeA and nodeB, in an HA pair, when nodeA has completed the SFO phase of an optimized takeover of nodeB (but has not completed the CFO phase of the optimized takeover). In this case, nodeA has information about the takeover status of nodeB's aggregates.


```

cluster1::> storage failover show-takeover
ode      Node Status      Aggregate      Takeover Status
-----
odeA      Partner has
relocated its
aggregates. Takeover
in progress.
          aggr1          Done
          aggr2          Done
          CFO aggregates In progress.
odeB      Relocated aggregates
to partner. Waiting
for partner to
takeover.
          aggr1          Done
          aggr2          Done
          CFO aggregates Not attempted yet.

```

The following example shows the takeover status of two nodes, nodeA and nodeB, in an HA pair, when nodeA has completed the SFO and CFO phases of an optimized takeover of nodeB. In this case, nodeA has information about the takeover status of nodeB's aggregates. Since nodeB is not operational, an Remote Procedure Call(RPC) error is indicated in the command output.

```

cluster1::> storage failover show-takeover
ode      Node Status      Aggregate      Takeover Status
-----
odeA      Partner has
relocated its
aggregates. In
takeover.
          aggr1          Done
          aggr2          Done
          CFO aggregates Done.
arning: Unable to list entries on node nodeB. RPC: Port mapper failure -
RPC:
          Timed out

```

The following example shows the takeover status of two nodes, nodeA and nodeB, in an HA pair, when nodeA has aborted the SFO phase of an optimized takeover of nodeB. In this case, nodeA does not have information about the takeover status of nodeB's aggregates.

```

cluster1::> storage failover show-takeover
ode          Node Status          Aggregate          Takeover Status
-----
odeA          Optimized takeover
              of partner aborted.
              Run the command
              "storage failover
              show-takeover -node
              nodeB" to display the
              relocation status of
              the partner.
              -
odeB          Optimized takeover
              by partner aborted.
              aggr1          Failed: Destination node did
                              not online the aggregate on
                              time. To takeover the
                              remaining aggregates, run the
                              "storage failover takeover
                              -ofnode nodeB
                              -bypass-optimization true"
                              command. To giveback the
                              relocated aggregates, run the
                              "storage failover giveback
                              -ofnode nodeB" command.
              aggr2          Not attempted yet
              CFO aggregates Not attempted yet.

```

storage failover show

Display storage failover status

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

Description

The `storage failover show` command displays information about storage-failover configurations. By default, the command displays the following information:

- Node name.
- Partner node name.
- Whether storage failover is possible.
- The current state of storage failover. If the takeover is disabled the appropriate reason would be displayed.

To display detailed information about storage failover on a specific node, run the command with the `-node`

parameter. The detailed view adds the following information:

- Node NVRAM ID.
- Partner NVRAM ID.
- Whether storage failover is enabled.
- Whether the storage-failover interconnect is available.
- Status of individual storage-failover interconnect links.
- Type and vendor of the storage-failover interconnect.
- Partner State
- Status codes from the takeover-by-partner process. Possible values include:
 - NVRAM_DOWN
 - OPERATOR_DISABLE_NVRAM
 - PARTNER_RESET
 - FM_TAKEOVER
 - NVRAM_MISMATCH
 - OPERATOR_DENY
 - CLUSTER_DISABLE
 - VERSION
 - SHELF_HOT
 - REVERT_IN_PROGRESS
 - HALT_NOTKOVER
 - TAKEOVER_ON_PANIC
- Reasons why takeover is not possible, if applicable. Possible values include:
 - NOT_INIT
 - DISABLED
 - DEGRADED
 - MBX_UNKNOWN
 - FM_VERSION
 - PARTNER_DISABLED
 - OPERATOR_DENY
 - NVRAM_MISMATCH
 - VERSION
 - IC_ERROR
 - BOOTING
 - SHELF_HOT
 - PARTNER_REVERT_IN_PROGRESS
 - LOCAL_REVERT_IN_PROGRESS
 - PARTNER_TAKEOVER

- LOCAL_TAKEOVER
- HALT_NOTKOVER
- LOG_UNSYNC
- UNKNOWN
- WAITING_FOR_PARTNER
- LOW_MEMORY
- HALTING
- MBX_UNCERTAIN
- NO_AUTO_TKOVER
- Time until takeover, in seconds.
- Time until auto giveback, in seconds.
- Delay for auto giveback, in seconds.
- List of local mailbox disks.
- List of partner mailbox disks.
- Whether operator-initiated planned takeover will be optimized for performance by relocating SFO (non-root) aggregates serially to the partner prior to takeover.

You can specify additional parameters to select the displayed information. For example, to display information only about storage-failover configurations whose interconnect is down, run the command with ``-interconnect-up false``.

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.

| [-options]

Displays the following information:

- Node name
- Whether automatic giveback operations are enabled
- Whether long-running operations are terminated when an automatic giveback operation is initiated
- Whether the node checks its partner's readiness before initiating a giveback operation
- The time, in seconds, that the node remains unresponsive before its partner initiates a takeover operation
- Whether the node automatically takes over for its partner if the partner fails
- Whether the node automatically takes over for its partner if the partner panics
- Whether the node automatically takes over for its partner if the partner reboots
- whether Hardware Assisted Takeover is enabled
- Ip address on which the partner node listens to the Hardware Assist alerts
- Port number on which the partner node listens to the Hardware Assist alerts

- Whether operator-initiated planned takeover will be optimized for performance by relocating SFO (non-root) aggregates serially to the partner prior to takeover

If this parameter is specified when the privilege level is set to advanced or higher, the command displays the information in the previous list and the following additional information:

- Whether the node takes over for its partner if its partner fails after a period of time, which is listed in the following field
- The number of seconds before the node takes over for its partner
- The number of times the node attempts an automatic giveback operation within a period of time
- The number of minutes in which the automatic giveback attempts can occur
- Whether storage-failover status is communicated via mailbox disks
- The interval at which the node reads its partner node's status from the mailbox disks
- The interval at which the node writes its status to the mailbox disks
- '
 - The interval at which Hardware assist h/w sends a heartbeat
 - The number of times the Hardware assist alert is sent

| **[-takeover-status]**

Displays the following information:

- Node name
- Partner name
- Takeover enabled
- Takeover possible
- Interconnect up
- State
- Node NVRAM ID
- Partner NVRAM ID
- Reason Takeover Not Possible By Partner
- Reason Takeover Not Possible
- Time Until Takeover

| **[-advanced] (privilege: advanced)**

Displays the following information:

- Node name
- Whether kill messages are issued during a takeover operation
- Whether the node controls its partner's storage aggregates
- The time when firmware notification was received
- The time when booting notification was received
- The time at which the last takeover or giveback operation occurred, in microseconds

- The number of times the failover log was unsynchronized

| **[-iotime] (privilege: advanced)**

Displays the following information:

- Node name
- Primary normal I/O time
- Primary transition I/O time
- Backup normal I/O time
- Backup transition I/O time

| **[-mailbox-status] (privilege: advanced)**

Displays the following information:

- Node name
- Primary mailbox status
- Backup mailbox status

| **[-more-options] (privilege: advanced)**

Displays the following information:

- Node name
- Whether takeover on short uptime is enabled
- Short uptime, in seconds
- Number of giveback attempts
- Interval of giveback attempts, in minutes
- Whether the primary mailbox is online
- Mailbox status read interval, in seconds
- Mailbox status write interval, in seconds

| **[-progress] (privilege: advanced)**

Displays the following information:

- Node name
- Maximum resource-table index number
- Current resource-table index number
- Current resource-table entry

| **[-timeout] (privilege: advanced)**

Displays the following information:

- Node name
- Fast timeout
- Slow timeout

- Mailbox timeout
- Connection timeout
- Operator timeout
- Firmware timeout
- Dump-core timeout
- Booting timeout
- Reboot timeout

| [-transit] (privilege: advanced)

Displays the following information:

- Node name
- Transit Timer Enabled
- Transit Timeout

| [-instance] }

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

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

Selects the nodes whose name matches this parameter value.

[-partner-name <text>] - Partner Name

Selects the nodes that have the specified partner-name setting.

[-nvramid <integer>] - Node NVRAM ID

Selects the nodes that have the specified NVRAM ID setting.

[-partner-nvramid <integer>] - Partner NVRAM ID

Selects the nodes that have the specified partner NVRAM ID setting.

[-enabled {true|false}] - Takeover Enabled

Selects the nodes that have the specified takeover-enabled setting.

[-mode {ha|non_ha}] - HA Mode

Selects the nodes that have the specified HA-mode setting. If the value is set to *ha* then the node is a member of a storage-failover configuration. If it is set to *non-ha* then it is in a stand alone configuration.

[-possible {true|false}] - Takeover Possible

Selects the nodes that have the specified failover-possible setting.

[-reason <text>, ...] - Reason Takeover not Possible

Selects the nodes that have the specified reason-not-possible setting. Possible values include:

- NOT_INIT
- DISABLED
- DEGRADED

- MBX_UNKNOWN
- FM_VERSION
- PARTNER_DISABLED
- OPERATOR_DENY
- NVRAM_MISMATCH
- VERSION
- IC_ERROR
- BOOTING
- SHELF_HOT
- PARTNER_REVERT_IN_PROGRESS
- LOCAL_REVERT_IN_PROGRESS
- PARTNER_TAKEOVER
- LOCAL_TAKEOVER
- HALT_NOTKOVER
- LOG_UNSYNC
- UNKNOWN
- WAITING_FOR_PARTNER
- LOW_MEMORY
- HALTING
- MBX_UNCERTAIN
- NO_AUTO_TKOVER

`[-interconnect-up {true|false}] - Interconnect Up`

Selects the nodes that have the specified interconnect-up setting.

`[-interconnect-links <text>] - Interconnect Links`

Selects the nodes that have the specified interconnect-links setting.

`[-interconnect-type <text>] - Interconnect Type`

Selects the nodes that have the specified interconnect-type setting.

`[-state-description <text>] - State Description`

Selects the nodes that have the specified state-description setting.

`[-partner-state <text>] - Partner State`

Selects the nodes that have the specified partner-state setting. Possible values include:

- OPERATOR COMPLETED
- DEBUGGUER COMPLETED
- PROGRESS COUNTER
- I/O ERROR

- BAD CHECKSUM
- RESERVED
- UNKNOWN
- INITIALIZING
- IN POWER-ON SELF TEST
- BOOTING
- BOOT FAILED
- WAITING
- KERNEL LOADED
- UP
- IN DEBUGGER
- WAITING FOR OPERATOR INPUT
- DUMPING CORE
- HALTED
- REBOOTING
- WAITING FOR GIVEBACK (DISK RESERVATIONS)
- WAITING FOR GIVEBACK (HA MAILBOXES)
- DUMPING SPARECORE
- MULTI-DISK PANIC
- IN TAKEOVER

`[-time-until-takeover <integer>]` - Time Until Takeover

Selects the nodes that have the specified time-until-takeover setting.

`[-partner-reason <text>,...]` - Reason Takeover not Possible by Partner

Selects the nodes that have the specified partner-reason text setting.

`[-killpackets {true|false}]` - Issue Kill Packets (privilege: advanced)

Selects the nodes that have the specified kill packets setting.

`[-partner-aggregates {true|false}]` - Control Partner Aggregates (privilege: advanced)

Selects the nodes that have the specified partner aggregates setting.

`[-current-index <integer>]` - Current Progress Index (privilege: advanced)

Selects the nodes that have the specified current-progress index setting.

`[-current-entry <text>]` - Current Progress Entry (privilege: advanced)

Selects the nodes that have the specified current-progress entry setting.

`[-maximum-index <integer>]` - Maximum Progress Index (privilege: advanced)

Selects the nodes that have the specified maximum-progress index setting.

[-pmbox-status <text>,...] - Primary Mailbox Status (privilege: advanced)

Selects the nodes that have the specified primary mailbox status setting. Possible values include:

- MBX_STATUS_NODISKS
- MBX_STATUS_UNCERTAIN
- MBX_STATUS_STALE
- MBX_STATUS_CONFLICTED
- MBX_STATUS_OLD_VERSION
- MBX_STATUS_NOT_FOUND
- MBX_STATUS_WRONG_STATE
- MBX_STATUS_BACKUP

[-bmbox-status <text>,...] - Backup Mailbox Status (privilege: advanced)

Selects the nodes that have the specified backup-mailbox status setting. See the description of the -pmbox -status parameter for a list of possible values.

[-major-seq-num-local <integer>] - Local Major Sequence Number (privilege: advanced)

Selects the nodes that have the specified mailbox heartbeat major sequence number on the local node.

[-minor-seq-num-local <integer>] - Local Minor Sequence Number (privilege: advanced)

Selects the nodes that have the specified mailbox heartbeat minor sequence number on the local node.

[-major-seq-num-partner <integer>] - Partner Major Sequence Number (privilege: advanced)

Selects the nodes that have the specified mailbox heartbeat major sequence number on the partner node.

[-minor-seq-num-partner <integer>] - Partner Minor Sequence Number (privilege: advanced)

Selects the nodes that have the specified mailbox heartbeat minor sequence number on the partner node.

[-local-mbx-node-status <Mailbox Status>] - Local Mailbox Node Status (privilege: advanced)

Selects the nodes that have the specified local mailbox node status. Possible values include:

- MBX_UNKNOWN - Local node is up, mailbox uninitialized
- MBX_TAKEOVER_DISABLED - Local node is up but takeover is disallowed
- MBX_TAKEOVER_ENABLED - Local node is up and takeover is allowed
- MBX_TAKEOVER_ACTIVE - Partner node has taken over
- MBX_GIVEBACK_DONE - Giveback completed, but local node has not yet restarted

[-mbx-abs-time-local <integer>] - Local Mailbox Absolute Time (privilege: advanced)

Selects the nodes that have the specified local mailbox channel absolute time. This time is measured in msec since 1/1/1970 (epoch).

[-mbx-sk-time-local <integer>] - Local Mailbox Kernel Time (privilege: advanced)

Selects the nodes that have the specified local mailbox channel Kernel Time.

[-mbx-sk-cycles-local <integer>] - Local Mailbox CPU Cycles (privilege: advanced)

Selects the nodes that have the specified local mailbox channel CPU Cycle count.

[-ic-abs-time-local <integer>] - Local IC Absolute Time (privilege: advanced)

Selects the nodes that have the specified local Interconnect channel absolute time. This time is measured in msec since 1/1/1970 (epoch).

[-ic-sk-time-local <integer>] - Local IC Kernel Time (privilege: advanced)

Selects the nodes that have the specified local Interconnect channel Kernel Time.

[-ic-sk-cycles-local <integer>] - Local IC CPU Cycles (privilege: advanced)

Selects the nodes that have the specified local Interconnect channel CPU Cycle count.

[-partner-mbx-node-status <Mailbox Status>] - Partner Mailbox Node Status (privilege: advanced)

Selects the nodes that have the specified partner mailbox node status. Possible values include:

- MBX_UNKNOWN
- MBX_TAKEOVER_DISABLED
- MBX_TAKEOVER_ENABLED
- MBX_TAKEOVER_ACTIVE
- MBX_GIVEBACK_DONE

[-mbx-abs-time-partner <integer>] - Partner Mailbox Absolute Time (privilege: advanced)

Selects the nodes that have the specified partner mailbox channel absolute time. This time is measured in msec since 1/1/1970 (epoch).

[-mbx-sk-time-partner <integer>] - Partner Mailbox Kernel Time (privilege: advanced)

Selects the nodes that have the specified partner mailbox channel Kernel Time.

[-mbx-sk-cycles-partner <integer>] - Partner Mailbox CPU Cycles (privilege: advanced)

Selects the nodes that have the specified partner mailbox channel CPU Cycle count.

[-mbx-major-seq-num-partner <integer>] - Partner Mailbox Major Sequence Number (privilege: advanced)

Selects the nodes that have the specified partner mailbox channel major sequence number.

[-mbx-minor-seq-num-partner <integer>] - Partner Mailbox Minor Sequence Number (privilege: advanced)

Selects the nodes that have the specified partner mailbox channel minor sequence number.

[-ic-abs-time-partner <integer>] - Partner IC Absolute Time (privilege: advanced)

Selects the nodes that have the specified partner Interconnect channel absolute time. This time is measured in msec since 1/1/1970 (epoch).

[-ic-sk-time-partner <integer>] - Partner IC Kernel Time (privilege: advanced)

Selects the nodes that have the specified partner Interconnect channel Kernel Time.

`[-ic-sk-cycles-partner <integer>]` - Partner IC CPU Cycles (privilege: advanced)

Selects the nodes that have the specified partner Interconnect channel CPU Cycle count.

`[-ic-major-seq-num-partner <integer>]` - Partner IC Major Sequence Number (privilege: advanced)

Selects the nodes that have the specified partner Interconnect channel major sequence number.

`[-ic-minor-seq-num-partner <integer>]` - Partner IC Minor Sequence Number (privilege: advanced)

Selects the nodes that have the specified partner Interconnect channel minor sequence number.

`[-local-takeover-info <text>]` - Local Takeover Info (privilege: advanced)

Selects the nodes that have the specified local node takeover information. This includes the type of negotiated failover request, or if takeover is not possible, the reason why takeover is disabled. Possible values include:

- NOTKOVER_NVRAM_DOWN - NVRAM mirror is down
- NOTKOVER_OPERATOR_DISABLE_NVRAM - Operator disabled
- NOTKOVER_PARTNER_RESET - A link reset is in progress
- NOTKOVER_FM_TAKEOVER - The failover monitor has declared takeover
- NOTKOVER_NVRAM_MISMATCH - NVRAM sizes mismatch
- NOTKOVER_OPERATOR_DENY - Operator denies takeover
- NOTKOVER_CLUSTER_DISABLE - Cluster is disabled
- NOTKOVER_VERSION - Version mismatch
- NOTKOVER_SHELF_HOT - Disk shelf is too hot
- NOTKOVER_REVERT_IN_PROGRESS - Revert is in progress
- NOTKOVER_HALT_NOTKOVER - Node halted in no-takeover mode
- TKOVER_ON_REBOOT - Enable takeover on reboot
- TKOVER_ON_PANIC - Enabled takeover on panic
- TKOVER_ON_STUTTER_DISABLED - Disable takeover on short uptime
- NFO_DISK_SHELF_ENABLED - Negotiated failover for disk shelf module is enabled
- NFO_ISCSI_ENABLED - Negotiated failover for network interfaces module is enabled
- NFO_FCP_TARGET_ENABLED - Negotiated failover for fcp target module is enabled

`[-partner-takeover-info <text>]` - Partner Takeover Info (privilege: advanced)

Selects the nodes that have the specified partner node takeover information. This includes the type of negotiated failover request, or if takeover is not possible, the reason why takeover is disabled. Possible values include:

- NOTKOVER_NVRAM_DOWN - NVRAM mirror is down
- NOTKOVER_OPERATOR_DISABLE_NVRAM - Operator disabled
- NOTKOVER_PARTNER_RESET - A link reset is in progress
- NOTKOVER_FM_TAKEOVER - The failover monitor has declared takeover

- NOTKOVER_NVRAM_MISMATCH - NVRAM sizes mismatch
- NOTKOVER_OPERATOR_DENY - Operator denies takeover
- NOTKOVER_CLUSTER_DISABLE - Cluster is disabled
- NOTKOVER_VERSION - Version mismatch
- NOTKOVER_SHELF_HOT - Disk shelf is too hot
- NOTKOVER_REVERT_IN_PROGRESS - Revert is in progress
- NOTKOVER_HALT_NOTKOVER - Node halted in no-takeover mode
- TKOVER_ON_REBOOT - Takeover on reboot is enabled
- TKOVER_ON_PANIC - Takeover on panic is enabled
- TKOVER_ON_STUTTER_DISABLED - Disable takeover on short uptime
- NFO_DISK_SHELF_ENABLED - Negotiated failover for disk shelf module is enabled
- NFO_ISCSI_ENABLED - Negotiated failover for network interfaces module is enabled
- NFO_FCP_TARGET_ENABLED - Negotiated failover for fcp target module is enabled

[-local-headswap-state <Headswap State>] - Local Head Swap State (privilege: advanced)

Selects the nodes that have the specified local node headswap state. Possible values are:

- HEADSWAP_NONE - head swap not in progress
- HEADSWAP_START - head swap started
- HEADSWAP_CFO_START - CFO phase of head swap started
- HEADSWAP_CFO_END - CFO phase of head swap completed
- HEADSWAP_SFO_START - SFO phase of head swap started

[-partner-headswap-state <Headswap State>] - Partner Head Swap State (privilege: advanced)

Selects the nodes that have the specified partner node headswap state. Possible values are:

- HEADSWAP_NONE - head swap not in progress
- HEADSWAP_START - head swap started
- HEADSWAP_CFO_START - CFO phase of head swap started
- HEADSWAP_CFO_END - CFO phase of head swap completed
- HEADSWAP_SFO_START - SFO phase of head swap started

[-fast-timeout <integer>] - Fast Timeout (privilege: advanced)

Selects the nodes that have the specified fast-timeout configuration setting.

[-slow-timeout <integer>] - Slow Timeout (privilege: advanced)

Selects the nodes that have the specified slow-timeout setting.

[-mailbox-timeout <integer>] - Mailbox Timeout (privilege: advanced)

Selects the nodes that have the specified mailbox-timeout setting.

[-connect-timeout <integer>] - Connect Timeout (privilege: advanced)

Selects the nodes that have the specified connect-timeout setting.

[-operator-timeout <integer>] - Operator Timeout (privilege: advanced)
Selects the nodes that have the specified operator-timeout setting.

[-firmware-timeout <integer>] - Firmware Timeout (privilege: advanced)
Selects the nodes that have the specified firmware-timeout setting.

[-dumpcore-timeout <integer>] - Dumpcore Timeout (privilege: advanced)
Selects the nodes that have the specified dump-core timeout setting.

[-booting-timeout <integer>] - Booting Timeout (privilege: advanced)
Selects the nodes that have the specified booting-timeout setting.

[-transit-timer {true|false}] - Transit Timer Enabled (privilege: advanced)
Selects the nodes that have the specified transit-timer setting.

[-transit-timeout <integer>] - Transit Timeout (privilege: advanced)
Selects the nodes that have the specified transit timeout.

[-firmware-received <integer>] - Firmware Received (privilege: advanced)
Selects the nodes that have the specified firmware-reception time.

[-firmware-received-cycles <integer>] - Firmware Received in CPU Cycles (privilege: advanced)
Selects the nodes that have the specified firmware-reception time in CPU Cycles.

[-booting-received <integer>] - Booting Received (privilege: advanced)
Selects the nodes that have the specified booting-reception time.

[-transit-time <integer>] - Transit Event Time (privilege: advanced)
Selects the nodes whose last failover event occurred at the specified time.

[-pnormal <integer>] - Primary Normal IO Time (privilege: advanced)
Selects the nodes that have the specified normal primary-mailbox I/O time.

[-ptransition <integer>] - Primary Transition IO Time (privilege: advanced)
Selects the nodes that have the specified transitional primary-mailbox I/O time.

[-bnormal <integer>] - Backup Normal IO Time (privilege: advanced)
Selects the nodes that have the specified normal backup-mailbox I/O time.

[-btransition <integer>] - Backup Transition IO Time (privilege: advanced)
Selects the nodes that have the specified transitional backup-mailbox I/O time.

[-logs-unsynched <integer>] - Logs Unsynched Count (privilege: advanced)
Selects the nodes that have the specified count of unsynchronized logs.

[-auto-giveback {true|false}] - Auto Giveback Enabled
Selects the nodes that have the specified auto-giveback setting.

[-detection-time <integer>] - Takeover Detection Time (secs)
Selects the nodes that have the specified detection-time setting.

[-onfailure {true|false}] - Takeover on Failure Enabled (privilege: advanced)
Selects the nodes that have the specified takeover-on-failure setting.

[-onpanic {true|false}] - Takeover on Panic Enabled
Selects the nodes that have the specified takeover-on-panic setting.

[-onshort-uptime {true|false}] - Takeover on Short Uptime Enabled (privilege: advanced)
Selects the storage-failover configurations that match this parameter value.

[-short-uptime <integer>] - Short Uptime (secs) (privilege: advanced)
Selects the nodes that have the specified short-uptime value.

[-attempts <integer>] - Number of Giveback Attempts (privilege: advanced)
Selects the nodes that have the specified number of giveback attempts.

[-attempts-time <integer>] - Giveback Attempts Period (minutes) (privilege: advanced)
Selects the nodes that have the specified time setting for giveback attempts.

[-propagate {true|false}] - Propagate Status via Mailbox (privilege: advanced)
Selects the nodes that have the specified propagate-status-via-mailbox setting.

[-read-interval <integer>] - Node Status Read Interval (secs) (privilege: advanced)
Selects the nodes that have the specified read interval.

[-write-interval <integer>] - Node Status Write Interval (secs) (privilege: advanced)
Selects the nodes that have the specified write interval.

[-onreboot {true|false}] - Takeover on Reboot Enabled
Selects the nodes that have the specified takeover-on-reboot setting.

[-delay-seconds <integer>] - Delay Before Auto Giveback (secs)
Selects the nodes that have the specified delay (in seconds) for the auto giveback.

[-hwassist {true|false}] - Hardware Assist Enabled
Selects the nodes that have the specified hwassist setting.

[-hwassist-partner-ip <IP Address>] - Partner's Hwassist IP
Selects the nodes that have the specified hwassist-partner-ip setting.

[-hwassist-partner-port <integer>] - Partner's Hwassist Port
Selects the nodes that have the specified hwassist-partner-port setting.

[-hwassist-health-check-interval <integer>] - Hwassist Health Check Interval (secs)
Selects the nodes that have the specified hwassist health check interval, in seconds.

`[-hwassist-retry-count <integer>]` - Hwassist Retry Count

Selects the nodes that have the specified hwassist retry count, in seconds.

`[-hwassist-status <text>]` - Hwassist Status

Selects the nodes that have the specified hwassist-status setting.

`[-time-until-autogiveback <integer>]` - Time Until Auto Giveback (secs)

Selects the nodes that have the specified time(in seconds) until auto giveback.

`[-local-mailbox-disks <text>]` - Local Mailbox Disks

Selects the nodes that have the specified mailbox disks on the local node.

`[-partner-mailbox-disks <text>]` - Partner Mailbox Disks

Selects the nodes that have the specified mailbox disks on the partner node.

`[-local-firmware-state <text>]` - Local Firmware State (privilege: advanced)

Selects the nodes that have the specified firmware state on the local node.

`[-local-firmware-progress <integer>]` - Local Firmware Progress Counter (privilege: advanced)

Selects the nodes that have the specified firmware progress counter for the local node.

`[-partner-firmware-state <text>]` - Partner Firmware State (privilege: advanced)

Selects the nodes that have the specified firmware state of the partner node.

`[-partner-firmware-progress <integer>]` - Partner Firmware Progress Counter (privilege: advanced)

Selects the nodes that have the specified firmware progress counter for the partner node.

`[-local-missing-disks <text>]` - Missing Disks on Local Node

Selects the nodes that have the specified missing disks on the local node.

`[-partner-missing-disks <text>]` - Missing Disks on Partner Node

Selects the nodes that have the specified missing disks on the partner node.

`[-reboot-timeout <integer>]` - Reboot Timeout (privilege: advanced)

Selects the nodes that have the specified reboot timeout.

`[-time-since-takeover <text>]` - Time Since Takeover

Selects the nodes that have been in takeover mode for the specified amount of time.

`[-auto-giveback-after-panic {true|false}]` - Auto Giveback After Takeover On Panic

Selects the nodes that have the specified auto-giveback-after-panic setting. If *true* then an automatic giveback operation is invoked when the failover node of an HA pair is repaired and rebooted. The takeover node of the HA pair detects this and initiates a giveback operation automatically.

`[-is-giveback-requested {true|false}]` - Giveback Requested (privilege: advanced)

Selects the nodes that have the specified is-giveback-requested setting. If *true*, a deferred giveback request has been made by the local node.

[-auto-giveback-last-veto-check <integer>] - Auto Giveback Last Veto Check (privilege: advanced)

Selects the nodes that have the specified auto-giveback-last-veto-check time. This setting indicates the time, in milliseconds, when the local node made the most recent giveback veto check.

[-is-auto-giveback-attempts-exceeded {true|false}] - Auto Giveback Attempts Exceeded (privilege: advanced)

Selects the nodes that have the specified is-auto-giveback-attempts-exceeded setting. If *true*, the local node has exceeded the maximum number of allowed auto giveback attempts.

[-was-auto-giveback-done {true|false}] - Was Auto Giveback Done (privilege: advanced)

Selects the nodes that have the specified was-auto-giveback-done setting. If *true*, the last giveback was automatic (as opposed to a manual giveback).

[-is-cifs-auto-giveback-stopping {true|false}] - Is CIFS Auto Giveback Stopping (privilege: advanced)

Selects the nodes that have the specified is-cifs-auto-giveback-stopping setting. If *true*, the local node has initiated CIFS termination as part of an automatic giveback.

[-aggregate-migration-timeout <integer>] - Aggregate Migration Timeout (secs) (privilege: advanced)

Selects the nodes that have the specified aggregate migration timeout.

[-is-mirror-enabled {true|false}] - Is NVRAM Mirroring Enabled (privilege: advanced)

Selects the nodes that have the specified is-mirror-enabled setting. If *true*, then NVRAM mirroring is enabled.

[-is-mirror-consistency-required {true|false}] - Is Mirror Consistency Required (privilege: advanced)

Selects the nodes that have the specified is-mirror-consistency-required setting. If *true*, then NVRAM mirror consistency is required.

[-is-memory-insufficient {true|false}] - Is Memory Insufficient To Takeover (privilege: advanced)

Selects the nodes that have the specified is-memory-insufficient setting. If *true*, the local node does not have enough memory to perform a takeover.

[-memio-state <memio status>] - Current State of Memio Link (privilege: advanced)

Selects the nodes that have the specified memio layer link current state. Possible values are:

- UNINIT - Uninitialized
- CLOSED - Closed
- HB_LISTEN - Listening for connect
- SYN_SENT - Sent generation information
- ESTABLISHED - Connection established

[-is-degraded {true|false}] - Are Partner Mailbox Disks Not Known (privilege: advanced)

Selects the nodes that have the specified is-degraded setting. If *true*, takeovers are deferred because partner mailbox disks are not known.

`[-reserve-policy <reserve policy>]` - FM Reservation Policy (privilege: advanced)

Selects the nodes that have the specified disk reservation policy. Possible values are:

- `RESERVE_NO_DISKS` - no disk reservations made during takeover, nor are disk reservations released during giveback
- `RESERVE_LOCK_DISKS_ONLY` - only mailbox disks are released during takeover and released during giveback
- `RESERVE_ONLY_AT_TAKEOVER` - reservations are issued only at takeover time. All disks are reserved. All reservations are released at giveback
- `RESERVE_ALWAYS_AFTER_TAKEOVER` - reservations are issued at takeover. When disks are subsequently added, they are also reserved. All disks are released at giveback

`[-reset-disks {true|false}]` - Issue Disk Resets during Failover (privilege: advanced)

Selects the nodes that have the specified reset-disks setting. If `true`, disks are reset during takeover/giveback.

`[-total-system-uptime <integer>]` - Total System Uptime (privilege: advanced)

Selects the nodes that have the specified total system uptime, in milliseconds.

`[-current-time <integer>]` - Current System Time (privilege: advanced)

Selects the nodes that have the specified current time on the node.

`[-fm-takeover-state <FM Takeover/Giveback Transition>]` - FM Takeover State (privilege: advanced)

Selects the nodes that have the specified takeover state. Possible values are:

- `FT_NONE` - Not in takeover
- `FT_TAKEOVER_STARTED` - Local node has initiated takeover
- `FT_TAKEOVER_COMMITTED` - Takeover has been committed
- `FT_TAKEOVER_DONE_OK` - Local node successfully completed takeover
- `FT_TAKEOVER_DONE_FAILED` - Takeover failed

`[-fm-giveback-state <FM Takeover/Giveback Transition>]` - FM Giveback State (privilege: advanced)

Selects the nodes that have the specified giveback state. Possible values are:

- `FT_NONE` - Not in giveback
- `FT_GIVEBACK_READY` - Partner node is ready for giveback
- `FT_GIVEBACK_STARTED` - Local node has initiated giveback
- `FT_GIVEBACK_COMMITTED` - Giveback has been committed
- `FT_GIVEBACK_DONE_OK` - Giveback completed successfully

`[-takeover-reason <FM Takeover Reason>]` - Reason why takeover triggered (privilege: advanced)

Selects the nodes that have the specified takeover reason. Possible values are:

- `TAKEOVER_NONE` - Not in takeover

- TAKEOVER_IMMEDIATE - Operator initiated forced takeover
- TAKEOVER_NDU - Takeover initiated as part of NDU
- TAKEOVER_FORCED - Operator initiated forced takeover, possible data loss
- TAKEOVER_EARLY - Takeover occurred during the boot process
- TAKEOVER_OPERATOR_EXP - Takeover occurred after the operator timeout expired
- TAKEOVER_POST_FAILED - Takeover occurred on POST failure
- TAKEOVER_PANIC - Takeover on panic
- TAKEOVER_SHORTUPTIME - Takeover after rapid toggling between up and down states
- TAKEOVER_SPARECORE_EXP - Takeover on panic timeout expiration
- TAKEOVER_REBOOT_EXP - Takeover on reboot timer expiration
- TAKEOVER_BOOTING_EXP - Takeover on booting timer expiration
- TAKEOVER_FIRMWARE_EXP - Takeover on firmware timer expiration
- TAKEOVER_NFO_SHUTDOWN - Takeover on negotiated failover shutdown
- TAKEOVER_NFO_TIMER - Takeover on negotiated failover timer expiration
- TAKEOVER_MDP - Takeover on multi-disk panic
- TAKEOVER_REBOOT - Takeover on reboot
- TAKEOVER_HALT - Takeover on halt
- TAKEOVER_CLAM - CLAM-triggered takeover
- TAKEOVER_HWASSIST - Hardware-assisted takeover
- TAKEOVER_NORMAL - Operator initiated takeover

[*-ha-type* {*none*|*shared_storage*|*non_shared_storage*}] - HA Type

If this parameter is specified, the command selects the nodes that have the specified HA-type setting. If the value is set to *shared_storage*, then the node is in a storage-failover configuration using the shared storage. If it is set to *non_shared_storage*, then the node is in a storage-failover configuration using the unshared storage. If it is set to *none*, then the node is not part of a storage-failover configuration.

Examples

The following example displays information about all storage-failover configurations:

```
cluster1::> storage failover show
```

| Node | Partner | Takeover Possible | State |
|-------|---------|----------------------|--------------------|
| node0 | node1 | true | Connected to node1 |
| node2 | node3 | true | Connected to node3 |
| node1 | node0 | true | Connected to node0 |
| node3 | node2 | true | Connected to node2 |

4 entries were displayed.

storage failover takeover

Take over the storage of a node's partner

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

Description

The `storage failover takeover` command initiates a takeover of the partner node's storage.

Parameters

{ `-ofnode` {<nodename>|local} - Node to Takeover

This specifies the node that is taken over. It is shut down and its partner takes over its storage.

| `-bynode` {<nodename>|local} - Node Initiating Takeover }

This specifies the node that is to take over its partner's storage.

[`-option` <takeover option>] - Takeover Option

This optionally specifies the style of takeover operation. Possible values include the following:

- `normal` - Specifies a normal takeover operation; that is, the partner is given the time to close its storage resources gracefully before the takeover operation proceeds. This is the default value.
- `immediate` - Specifies an immediate takeover. In an immediate takeover, the takeover operation is initiated before the partner is given the time to close its storage resources gracefully. The use of this option results in an immediate takeover which does not do a clean shutdown. In case of NDU this can result in a NDU failure.



If this option is specified, negotiated takeover optimization is bypassed even if the `-bypass -optimization` option is set to false.



If this option is specified, migration of data LIFs from the partner will be delayed even if the `-skip-lif-migration-before-takeover` option is not specified. If possible, migrate the data LIFs to another node prior to specifying this option.

- `allow-version-mismatch` - If this value is specified, the takeover operation is initiated even if the partner is running a version of software that is incompatible with the version running on the node. In this case, the partner is given the time to close its storage resources gracefully before the takeover operation proceeds. However, the takeover operation will not be allowed if the partner has higher WAFL or RAID label versions. Use this value as part of a nondisruptive upgrade or downgrade procedure.
- `force` - If this value is specified, the takeover operation is initiated even if the node detects an error that normally prevents a takeover operation from occurring. This value is available only at the advanced privilege level and higher.



If this option is specified, negotiated takeover optimization is bypassed even if the `-bypass -optimization` option is set to false.



The use of this option can potentially result in data loss. If the HA interconnect is detached or inactive, or the contents of the failover partner's NVRAM cards are unsynchronized, takeover is normally disabled. Using the `-force` option enables a node to take over its partner's storage despite the unsynchronized NVRAM, which can contain client data that can be lost upon storage takeover.

`[-bypass-optimization {true|false}] - Bypass Takeover Optimization (privilege: advanced)`

If this is an operator-initiated planned takeover, this parameter specifies whether the takeover optimization is bypassed. This parameter defaults to `false`.



This parameter is ignored and negotiated takeover optimization automatically bypassed if the `-immediate` option, the `-force` option, or the `-allow-disk-inventory-mismatch` parameter is specified as part of the same `storage failover takeover` command.

`[-allow-disk-inventory-mismatch {true|false}] - Disk inventory`

If this parameter is specified, the takeover operation is initiated even if the local node cannot see the partner's filesystem disks.



If this parameter is specified, negotiated takeover optimization is bypassed even if the `-bypass-optimization` parameter is set to `false`.



The use of this parameter can potentially result in client outage.

`[-skip-lif-migration-before-takeover <true>] - Skip Migrating LIFs Away from Node Prior to Takeover`

This parameter specifies that LIF migration prior to takeover is skipped. However if LIFs on this node are configured for failover, those LIFs may still failover after the takeover has occurred. Without this parameter, the command attempts to synchronously migrate data and cluster management LIFs away from the node prior to its takeover. If the migration fails or times out, the takeover is aborted.

`[-ignore-quorum-warnings <true>] - Skip Quorum Check Before Takeover`

If this parameter is specified, quorum checks will be skipped prior to the takeover. The operation will continue even if there is a possible data outage due to a quorum issue.

`[-override-vetoes <true>] - Override Vetoes`

If this is an operator-initiated planned takeover, this parameter specifies whether the veto should be overridden. If this parameter is not specified, its value is set to `false`.



If this parameter is specified, negotiated takeover will override any vetos to continue with takeover.



The use of this parameter might result in the takeover proceeding even if the node detects issues that can potentially make the takeover dangerous or disruptive.

`[-halt <true>] - Halt the Node That Is Taken Over`

This parameter specifies whether the node being taken over should be halted. If the value is `true`, then the node being taken over is halted. If the value is `false`, then the node being taken over is shutdown and might be rebooted if `AUTOBOOT` is set to `true`. This parameter defaults to `false`.

Examples

The following example causes a node named node0 to initiate a negotiated optimized takeover of its partner's storage:

```
cluster1::> storage failover takeover -bynode node0
```

The following example causes a node named node0 to initiate an immediate takeover of its partner's storage:

```
cluster1::> storage failover takeover -bynode node0 -option immediate
```

storage failover hwassist show

Display hardware-assisted storage failover status

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

Description

The `storage failover hwassist show` command displays information about the hardware-assisted storage failover status on each node. By default, the command displays the following information:

- Node name.
- Partner node name.
- Whether hardware-assisted failover is enabled.
- IP address on which the local node receives hardware-assisted failover alerts.
- Port on which the local node receives hardware-assisted failover alerts.
- Hardware-assisted failover status.
- If the monitor is inactive, the reason it is inactive.
- If the monitor is inactive, the corrective action to make it active.
- Status of keep-alive alerts on the local node.

Hardware-assisted failover establishes a notification channel from each respective node's service processor to the other (HA partner) node. If a node becomes unresponsive, its service processor notifies the HA partner of this condition, accelerating storage failover. By default, hwassist is enabled and configured automatically to use each node's node-mgmt LIF. To modify or show the hardware-assisted storage failover configuration, use the [storage failover modify](#) and [storage failover show](#) commands.

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 you specify this parameter, the command displays information only about hwassist on the specified nodes.

[`-partner-name {<nodename>|local}`] - Name of the Partner Node

If you specify this parameter, the command displays information only about hwassist on nodes with the specified partner node.

[`-enabled {true|false}`] - Local Hardware Assist Enabled

If you specify this parameter, the command displays information only about hwassist on nodes with the specified enabled state.

[`-local-status <text>`] - Local Node's Hwassist Status

If you specify this parameter, the command displays information only about hwassist on nodes with the specified local status (*active* or *inactive*).

[`-local-ip <text>`] - IP Address on Which Local Node is Listening

If you specify this parameter, the command displays information only about hwassist on nodes with the specified local IP address.

[`-local-port <integer>`] - Port on Which Local Node is Listening

If you specify this parameter, the command displays information only about hwassist on nodes with the specified local UDP port.

[`-local-inactive <text>`] - Local Node's Hwassist Inactive Status Reason

If you specify this parameter, the command displays information only about hwassist on nodes with the specified inactive reason.

[`-local-action <text>`] - Corrective Action on Local Node

If you specify this parameter, the command displays information only about hwassist on nodes with the specified corrective action.

Examples

The following example displays the hardware-assisted failover information for node cluster1-01 and its HA partner node cluster1-02:

```

cluster1::> storage failover hwassist show
Node
-----
cluster1-01
                Partner: cluster1-02
        Hwassist Enabled: true
                Hwassist IP: 10.225.248.19
                Hwassist Port: 4444
                Monitor Status: active
                Inactive Reason: -
                Corrective Action: -
        Keep-Alive Status: healthy

cluster1-02
                Partner: cluster1-01
        Hwassist Enabled: true
                Hwassist IP: 10.225.248.21
                Hwassist Port: 4444
                Monitor Status: active
                Inactive Reason: -
                Corrective Action: -
        Keep-Alive Status: healthy

```

Related Links

- [storage failover modify](#)
- [storage failover show](#)

storage failover hwassist test

Test the hwassist functionality

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

Description

The `storage failover hwassist test` command tests the Hardware Assist h/w connectivity between the two nodes in a HA pair. The test result can be one of the following.

- Hardware Assist is not initialized.
- HW assist is not supported.
- Partner is throttling alerts.
- Resource is busy.
- Hardware Assist h/w returned an error.
- No response from partner.Timed out.
- Unexpected abort.

- Partner has taken over.
- Interconnect is down between nodes.
- Partner is not booted up yet.

Parameters

-node {<nodename>|local} - Node

This specifies the node from which a test alert is initiated.

Examples

The following command issues a test alert from the node cluster1-01:

```
cluster1::> storage failover hwassist test -node cluster1-01
Info: Operation successful.
```

storage failover hwassist stats clear

Clear the hwassist statistics

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

Description

The `storage failover hwassist stats clear` command clears the statistics information maintained by Hardware Assist functionality.

Parameters

-node {<nodename>|local} - Node

This specifies the node on which the statistics are to be cleared.

Examples

The following example clears the hwassist statistics on the node cluster1-01:

```
cluster1::> storage failover hwassist stats clear -node cluster1-01
```

storage failover hwassist stats show

Display hwassist statistics

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

Description

The `storage failover hwassist stats show` command displays statistics about the hardware assist alerts processed by a node. The command displays the following information for each alert:

- Locally enabled.
- Partner Inactive Reason.
- Alert type.
- Event that triggered the alert.
- The number of times the alert has been received.
- Whether takeover was possible on receiving the alert.
- The last time at which the alert was received.

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

Selects the hwassist statistics for the specified node.

Examples

The following example displays the hwassist statistics for the node ha1:

```
cluster1:> storage failover hwassist stats show -node ha1
Node: ha1
        Local Enabled: true
        Partner Inactive Reason: -
Alert Type  Alert Event                Count Takeover  Last Received
-----
-----
system_down power_loss                    0 Yes          ---
system_down l2_watchdog_reset          0 Yes          ---
system_down power_off_via_rlm          0 Yes          ---
system_down power_cycle_via_rlm        0 Yes          ---
system_down reset_via_rlm              0 Yes          ---
system_down power_off_via_sp           0 Yes          ---
system_down power_cycle_via_sp         0 Yes          ---
system_down reset_via_sp               0 Yes          ---
system_down post_error                 0 No           ---
system_down abnormal_reboot            0 No           ---
system_down loss_of_heartbeat          0 No           ---
keep_alive periodic_message           10 No          Wed Mar  9 13:41:28
EST 2016
test      test                        0 No           ---
ID_mismatch ---                      0 ---          ---
Key_mismatch ---                    0 ---          ---
Unknown   ---                      0 ---          ---
          alerts_throttled           0 ---          ---
```

The following example displays the hwassist statistics for the node ha1 where hardware assist hardware is not supported.

```
cluster1:> storage failover hwassist stats show -node ha1
Node: ha1
        Local Enabled: false
        Partner Inactive Reason: HW assist is not supported on partner.
Alert Type  Alert Event                Count Takeover  Last Received
-----
-----
-
```

storage failover internal-options show

Display the internal options for storage failover

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

Description

The `storage failover internal-options show` command displays the following information about the storage failover configuration:

- Node name
- Whether automatic giveback is enabled
- Whether partner checking is enabled
- Takeover detection time, in seconds
- Whether takeover on failover is enabled
- Whether takeover on panic is enabled
- Whether takeover on reboot is enabled
- Whether hardware-assisted takeover is enabled
- IP address on which the partner node listens to the hardware-assisted takeover alerts
- Port on which the partner node listens to the hardware-assisted takeover alerts
- Whether takeover on short uptime is enabled (detailed view only)
- Short uptime interval, in seconds (detailed view only)
- Number of giveback attempts (detailed view only)
- Giveback attempt interval, in minutes (detailed view only)
- Whether status is propagated through SFO mailboxes (detailed view only)
- Status read interval, in seconds (detailed view only)
- Status write interval, in seconds (detailed view only)
- Hardware-assisted takeover retry count (detailed view only)
- Hardware-assisted takeover heartbeat period (detailed view only)
- Whether operator-initiated planned takeover is optimized

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.

| [-more] (privilege: advanced)

This parameter displays the following additional information: :

- Node name
- Whether takeover on short uptime is enabled
- Short uptime interval, in seconds
- Number of giveback attempts
- Giveback attempt interval, in minutes
- Whether status is propagated through SFO mailboxes
- Status read interval, in seconds

- Status write interval, in seconds
- Hardware-assisted takeover retry count
- Hardware-assisted takeover heartbeat period

| [-instance] }

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

[-node {<nodename>|local}] - Node (privilege: advanced)

Selects configuration information for the specified node.

[-auto-giveback {true|false}] - Auto Giveback Enabled (privilege: advanced)

Selects configuration information for nodes that have the specified automatic giveback setting.

[-check-partner {true|false}] - Check Partner Enabled (privilege: advanced)

Selects configuration information for nodes that have the specified partner-checking setting.

[-detection-time <integer>] - Takeover Detection Time (secs) (privilege: advanced)

Selects configuration information for nodes that have the specified takeover detection time setting.

[-onfailure {true|false}] - Takeover on Failure Enabled (privilege: advanced)

Selects configuration information for nodes that have the specified takeover-on-failure setting.

[-onpanic {true|false}] - Takeover on Panic Enabled (privilege: advanced)

Selects configuration information for nodes that have the specified takeover-on-panic setting.

[-onshort-uptime {true|false}] - Takeover on Short Uptime Enabled (privilege: advanced)

Selects configuration information for nodes that have the specified takeover-on-short-uptime setting.

[-short-uptime <integer>] - Short Uptime (secs) (privilege: advanced)

Selects configuration information for nodes that have the specified takeover-on-short-uptime time setting.

[-attempts <integer>] - Number of Giveback Attempts (privilege: advanced)

Selects configuration information for nodes that have the specified number of giveback attempts setting.

[-attempts-time <integer>] - Giveback Attempts Minutes (privilege: advanced)

Selects configuration information for nodes that have the specified giveback attempt time setting.

[-propagate {true|false}] - Propagate Status via Mailbox (privilege: advanced)

Selects configuration information for nodes that have the specified setting for propagation of status through Storage Failover mailboxes.

[-read-interval <integer>] - Node Status Read Interval (secs) (privilege: advanced)

Selects configuration information for nodes that have the specified status read interval setting.

[-write-interval <integer>] - Node Status Write Interval (secs) (privilege: advanced)

Selects configuration information for nodes that have the specified status write interval setting.

[`-onreboot {true|false}`] - Takeover on Reboot Enabled (privilege: advanced)

Selects configuration information for nodes that have the specified takeover-on-reboot setting.

[`-delay-seconds <integer>`] - Delay Before Auto Giveback (secs) (privilege: advanced)

If this parameter is specified, the command displays information only about the node or nodes that have the specified delay for auto giveback.

[`-hwassist {true|false}`] - Hwassist Enabled (privilege: advanced)

Selects configuration information for nodes that have the specified hardware-assisted takeover setting.

[`-hwassist-partner-ip <text>`] - Partner's Hwassist IP (privilege: advanced)

Selects configuration information for nodes that have the specified partner IP setting for hardware-assisted takeovers.

[`-hwassist-partner-port <integer>`] - Partner's Hwassist Port (privilege: advanced)

Selects configuration information for nodes that have the specified partner port setting for hardware-assisted takeovers.

[`-hwassist-health-check-interval <integer>`] - Hwassist Health Check Interval (secs) (privilege: advanced)

Selects configuration information for nodes that have the specified health check interval setting for hardware-assisted takeovers

[`-hwassist-retry-count <integer>`] - Hwassist Retry Count (privilege: advanced)

Selects configuration information for nodes that have the specified retry count (in seconds) for hardware-assisted takeovers.

[`-mode {ha|non_ha}`] - HA Mode (privilege: advanced)

If this parameter is specified, the command displays information only about the node or nodes that have the specified HA mode.

[`-bypass-takeover-optimization {true|false}`] - Bypass Takeover Optimization Enabled (privilege: advanced)

Selects configuration information for nodes that have the specified setting for bypass takeover optimization (`_true_` means that optimized operator-initiated planned takeover is bypassed, `false` means that it is enabled). Operator-initiated planned takeover is optimized when SFO aggregates are relocated serially to the partner prior to takeover. This reduces client outage.

Examples

The following example displays detailed information about the internal options for storage failover on a node named node2:

```

cluster1::*> storage failover internal-options show -node node2
Node: node2

Auto Giveback Enabled: false
Check Partner Enabled: true
Takeover Detection Time (secs): 15
Takeover On Failure Enabled: true
Takeover On Panic Enabled: false
Takeover On Short Uptime Enabled: true
Short Uptime (secs): -
Number of Giveback Attempts: 3
Giveback Attempts Minutes: 10
Propagate Status Via Mailbox: true
Node Status Read Interval (secs): 5
Node Status Write Interval (secs): 5
Failover the Storage when Cluster Ports Are Down: -
Failover Interval when Cluster Ports Are Down (secs): -
Takeover on Reboot Enabled: true
Delay Before Auto Giveback (secs): 300
Hardware Assist Enabled: true
Partner's Hw-assist IP:
Partner's Hw-assist Port: 4444
Hw-assist Health Check Interval (secs): 180
Hw-assist Retry count: 2
HA mode: ha
Bypass Takeover Optimization Enabled: true

```

storage failover mailbox-disk show

Display information about storage failover mailbox disks

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

Description

The `storage failover mailbox-disk show` command lists the mailbox disks that are used by storage failover. The command displays the following information:

- Node name
- Whether the mailbox disk is owned by the local node or by its partner
- Disk name
- Disk universal unique identifier (UUID)

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

Parameters

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

If -fields <fieldname>,... is used, the command displays only the specified fields.

| [-instance] }

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

[-node {<nodename>|local}] - Node (privilege: advanced)

Selects the mailbox disks that are associated with the specified node.

[-location {local|partner|tertiary}] - Mailbox Owner (privilege: advanced)

Selects the mailbox disks that have the specified relationship to the node.

[-diskindex <integer>] - Mailbox Disk Index (privilege: advanced)

Selects the mailbox disk that has the specified index number.

[-diskname <text>] - Mailbox Disk Name (privilege: advanced)

Selects the mailbox disks that match the specified disk name.

[-diskuuid <text>] - Mailbox Disk UUID (privilege: advanced)

Selects the mailbox disks that match the specified UUID.

[-physical-location {local|partner|mediator}] - Mailbox Disk Physical Location (privilege: advanced)

Selects the mailbox disks that match the specified physical location.

[-location-id <nvr amid>] - System ID of the Node where the Disk is Attached (privilege: advanced)

Selects the mailbox disks that match the specified location-id.

[-location-name <text>] - Mailbox Disk Location (privilege: advanced)

Selects the mailbox disks that match the specified location-name.

Examples

The following example displays information about the mailbox disks on a node named node1:

```
cluster1::*> storage failover mailbox-disk show -node node1
Node      Location  Index Disk Name      Physical Location  Disk UUID
-----
node1
      local      0 1.0.4      local      20000000:8777E9D6:[...]
      local      1 1.0.6      partner     20000000:8777E9DE:[...]
      partner     0 1.0.1      local      20000000:877BA634:[...]
partner    1 1.0.2      partner     20000000:8777C1F2:[...]
```


storage failover progress-table show

Display status information about storage failover operations

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

Description

The `storage failover progress-table show` displays status information about storage-failover operations. This information is organized in a resource table. The command displays the following information:

- Node name
- Resource-entry index number
- Resource-entry name
- Resource-entry state
- Resource-entry failure code
- Resource-entry time delta

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

Parameters

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

If `-fields <fieldname>`, ... is used, the command will only displays only the specified fields.

| [-instance] }

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

[-node {<nodename>|local}] - Node (privilege: advanced)

Selects the status information for the specified node.

[-index <integer>] - Resource Table Index (privilege: advanced)

Selects the status information for the specified index number.

[-entryname <text>] - Resource Table Entry Name (privilege: advanced)

Selects the status information for the specified entry name.

[-state <text>] - Resource Table Entry State (privilege: advanced)

Selects the status information for the specified state. Possible values include UP, START_RUNNING, START_DONE, START_FAILED, STOP_RUNNING, STOP_FAILED, TAKEOVER_BARRIER, and ONLY_WHEN_INITD.

[-failurecode <text>] - Entry Failure Code (privilege: advanced)

Selects the status information for the specified failure code. Possible values include OK, FAIL, FAIL_ALWAYS, HANG, PANIC, and VETO.

[-timedelta <integer>] - Entry Time Delta (privilege: advanced)

Selects the status information for the specified time delta.

Examples

The following example displays the entire storage-failover resource table:

```
cluster1::*> storage failover progress-table show
```

| Node | Entry Name | State | Time Delta |
|---|-------------------------------------|------------|------------|
| ----- | | | |
| node0 | | | |
| | Pre-rsrctbl: fmdisk_resumePartnerDi | start_done | 6 |
| | Pre-rsrctbl: coredump_get_busy_spar | start_done | 107 |
| | Pre-rsrctbl: raid_preread_labels_be | start_done | 1 |
| | Pre-rsrctbl: fmdisk_reserve_all | start_done | 84 |
| | rsrctbl: fmrsrc_giveback_done | start_done | 0 |
| | rsrctbl: fmic | start_done | 0 |
| | rsrctbl: fmdisk_reserve | start_done | 171 |
| | rsrctbl: fm_partnerSlowTimeout | start_done | 1 |
| | rsrctbl: fmdisk_inventory | start_done | 0 |
| | rsrctbl: fmfsn_reserve | start_done | 0 |
| Press <space> to page down, <return> for next line, or 'q' to quit... | | | |
| Node | Entry Name | State | Time Delta |
| ----- | | | |
| node0 | | | |
| | rsrctbl: rdb-ha | start_done | 36 |
| | rsrctbl: giveback_cleanup_wait | start_done | 0 |
| | rsrctbl: priority_ha | start_done | 0 |
| | rsrctbl: raid | start_done | 113 |
| | rsrctbl: raid_disaster_early | start_done | 0 |
| | rsrctbl: wafl_nvram_replay | start_done | 0 |
| | rsrctbl: takeover_test_1 | start_done | 0 |

storage firmware commands

storage firmware download

Download disk, ACP processor and shelf firmware

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

Description

The `storage firmware download` command downloads the ACP processor, disk, and shelf module firmware to a specified node. This command can also be used to download the disk qualification package (DQP).

Use the [storage disk firmware update](#) command to install downloaded disk firmware.

Use the [storage shelf firmware update](#) command to install downloaded shelf module firmware.

Use the [system node run`](#) local storage download acp` command to install downloaded ACP processor firmware.

Parameters

-node {<nodename>|local} - Node (privilege: advanced)

This specifies the node to which the firmware is to be downloaded.

-package-url <text> - Package URL (privilege: advanced)

This specifies the path to the firmware package.

The packaged ACP processor, disk, and shelf firmware files need to have ".AFW", ".LOD", and ".SFW" file extensions, respectively.

The following URL protocols are supported: ftp, http, tftp and file. The file URL scheme can be used to specify the location of the package to be fetched from an external device connected to the storage controller. Currently, only USB mass storage devices are supported. The USB device is specified as *file://usb0/<filename>*. The package must be present in the root directory of the USB mass storage device.

Examples

The following example downloads a disk firmware package with the path *ftp://example.com/fw/disk-fw-1.2.LOD.zip* to a node named *node1*:

```
cluster1::> storage firmware download -node node1 -package-url
ftp://example.com/fw/disk-fw-1.2.LOD.zip
```

Related Links

- [storage disk firmware update](#)
- [storage shelf firmware update](#)
- [system node run](#)

storage firmware acp delete

Delete an ACP firmware file

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

Description

The `storage firmware acp delete` command deletes the specified ACP processor firmware file from all nodes that are currently part of the cluster.

Parameters

-filename <text> - Firmware Filename (privilege: advanced)

Specifies the firmware file to delete.

Examples

The following example deletes the ACP processor firmware file with the name ACP-IOM3.0150.AFW.FVF on each node:

```
cluster1::*> storage firmware acp delete -filename ACP-IOM3.0150.AFW.FVF
```

storage firmware acp rename

Rename an ACP firmware file

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

Description

The `storage firmware acp rename` command renames the specified ACP processor firmware file on each node.

Parameters

-oldname <text> - Old Filename (privilege: advanced)

This parameter specifies the firmware file to rename.

-newname <text> - New Filename (privilege: advanced)

This parameter specifies the new name of the firmware file.

Examples

The following example renames the ACP processor firmware file with the name ACP-IOM3.0150.AFW.FVF to ACP-IOM3.AFW.FVF on each node:

```
cluster1::*> storage firmware acp rename -oldname ACP-IOM3.0150.AFW.FVF  
-newname ACP-IOM3.AFW.FVF
```

storage firmware acp show

Display the list of ACP firmware files on the given node

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

Description

The `storage firmware acp show` command displays the ACP processor firmware files present on each node.

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

Selects the files that match the specified node name.

[-filename <text>] - Storage Firmware File

Selects the files that match the specified filename.

Examples

The following example displays the ACP processor firmware files on each node:

```
cluster1::> storage firmware acp show
Node: Node1
ACP Firmware Files
-----
ACP-IOM3.0150.AFW.FVF
ACP-IOM3.AFW
ACP-IOM6.0210.AFW
ACP-IOM6.0210.AFW.FVF
Node: Node2
ACP Firmware Files
-----
ACP-IOM3.0150.AFW.FVF
ACP-IOM3.AFW
ACP-IOM6.0210.AFW
ACP-IOM6.0210.AFW.FVF
8 entries were displayed.
```

storage firmware disk delete

Delete a disk firmware file

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

Description

The `storage firmware disk delete` command deletes the specified disk firmware file on each node.

Parameters

-filename <text> - Storage Firmware Filename (privilege: advanced)

Specifies the firmware file to delete.

Examples

The following example deletes the disk firmware file with the name X262_SMOOST25SSX.NA06.LOD on each node:

```
cluster1::*> storage firmware disk delete -filename
X262_SMOOST25SSX.NA06.LOD
```

storage firmware disk rename

Rename a disk firmware file

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

Description

The `storage firmware disk rename` command renames the specified disk firmware file on each node.

Parameters

-oldname <text> - Old Filename (privilege: advanced)

This parameter specifies the firmware file to rename.

-newname <text> - New Filename (privilege: advanced)

This parameter specifies the new name of the firmware file.

Examples

The following example renames the disk firmware file with the name X262_SMOOST25SSX.NA06.LOD to X262_SMOOST25SSX.LOD on each node:

```
cluster1::*> storage firmware disk rename -oldname
X262_SMOOST25SSX.NA06.LOD -newname X262_SMOOST25SSX.LOD
```

storage firmware disk show

Display the list of disk firmware files on the given node

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

Description

The `storage firmware disk show` command displays the disk firmware files present on each node.

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

Selects the files that match the specified node name.

[-filename <text>] - Storage Firmware File

Selects the files that match the specified filename.

Examples

The following example displays the disk firmware files on each node:

```
cluster1::> storage firmware disk show
Node: Node1
Disk Firmware Files
-----
X262_SMOOST25SSX.NA06.LOD
X262_SMOOST25SSX.NA06.LOD.FVF
X267_SMOOST50SSX.NA06.LOD
X267_SMOOST50SSX.NA06.LOD.FVF
Node: Node2
Disk Firmware Files
-----
X262_SMOOST25SSX.NA06.LOD
X262_SMOOST25SSX.NA06.LOD.FVF
X267_SMOOST50SSX.NA06.LOD
X267_SMOOST50SSX.NA06.LOD.FVF
8 entries were displayed.
```

storage firmware shelf delete

Delete a shelf firmware file

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

Description

The `storage firmware shelf delete` command deletes the specified shelf firmware file from all nodes that are currently part of the cluster.

Parameters

-filename <text> - Storage Firmware Filename (privilege: advanced)

Specifies the firmware file to delete.

Examples

The following example deletes the shelf firmware file with the name IOM12.0210.SFW on each node:

```
cluster1::*> storage firmware shelf delete -filename IOM12.0210.SFW
```

storage firmware shelf rename

Rename a shelf firmware file

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

Description

The `storage firmware shelf rename` command renames the specified shelf firmware file on each node.

Parameters

-oldname <text> - Old Filename (privilege: advanced)

This parameter specifies the firmware file to rename.

-newname <text> - New Filename (privilege: advanced)

This parameter specifies the new name of the firmware file.

Examples

The following example renames the shelf firmware file with the name IOM12.0210.SFW to IOM12.000.SFW on each node:

```
cluster1::*> storage firmware shelf rename -oldname IOM12.0210.SFW  
-newname IOM12.000.SFW
```

storage firmware shelf show

Display the list of shelf firmware files on the given node

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

Description

The `storage firmware shelf show` command displays the shelf firmware files present on each node.

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

Selects the files that match the specified node name.

[-filename <text>] - Storage Firmware File

Selects the files that match the specified filename.

Examples

The following example displays the shelf firmware files on each node:

```
cluster1::> storage firmware shelf show
Node: Node1
Shelf Firmware Files
-----
AT-FCX.3800.SFW
AT-FCX.3800.SFW.FVF
ESH4.1400.SFW
ESH4.1400.SFW.FVF
Node: Node2
Shelf Firmware Files
-----
AT-FCX.3800.SFW
AT-FCX.3800.SFW.FVF
ESH4.1400.SFW
ESH4.1400.SFW.FVF
8 entries were displayed.
```

storage iscsi-initiator commands

storage iscsi-initiator add-target

Add an iSCSI target

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

Description

The `storage iscsi-initiator add-target` command adds an iSCSI target to a node's list of targets.

This command is only supported on high-availability shared-nothing virtualized platforms.

Parameters

-node {<nodename>|local} - Node (privilege: advanced)

Specifies the name of the Data ONTAP node to which the iSCSI target will be added.

-label <text> - User Defined Identifier (privilege: advanced)

Specifies a label for the target to be added.

-target-type {external|mailbox|partner|partner2|dr_auxiliary|dr_partner} - Target Type (privilege: advanced)

Specifies the type of the target. It is used by the node to determine how to use the LUNs. There are five target types:

- partner - The partner target should belong to the node's HA partner. This allows the node to access its partner's disks.
- mailbox - A mailbox target's LUNs are used exclusively as HA mailboxes.
- external - External targets' LUNs can be used by the node but do not play a role in HA.
- dr_auxiliary - The DR auxiliary target for MetroCluster over IP. Not a valid target type for the add-target command.
- dr_partner - The DR partner target for MetroCluster over IP. Not a valid target type for the add-target command.

-target-portal <text> - Target Portal (privilege: advanced)

Specifies the target's IP address and listening TCP port. The port is not required if it is the default iSCSI port (3260). Examples of correct target portals are *10.0.0.2* and *10.0.0.2:860*.

-target-name <text> - iSCSI Name (privilege: advanced)

Specifies the iSCSI target name such as an IQN (iSCSI qualified name).

[-status-admin {down|up}] - Administrative Status (default: up) (privilege: advanced)

Use to specify whether the initial administrative status of the connection is up or down. The default setting is *up*.

Examples

The following example adds and connects to an iSCSI target from the specified node.

```
cluster1::*> storage iscsi-initiator add-target -node node1
               -label target1 -target-type external
               -target-portal 10.0.0.2:860
               -target-name iqn.2012-06.com.bsdctl:target0
```

storage iscsi-initiator connect

Connect to an iSCSI target

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

Description

The `storage iscsi-initiator connect` command connects a node to the specified target. This command is only supported on high-availability shared-nothing virtualized platforms.

Parameters

-node {<nodename>|local} - Node (privilege: advanced)

Specifies the name of the Data ONTAP node to which the iSCSI target will be connected.

[-target-type {external|mailbox|partner|partner2|dr_auxiliary|dr_partner}] - Target Type (privilege: advanced)

Selects targets with the specified target type.

-label <text> - User Defined Identifier (privilege: advanced)

Specifies the label of the target to connect to.

Examples

The following example adds and connects to an iSCSI target from the specified node.

```
cluster1::*> storage iscsi-initiator connect -node node1
                    -label target1
```

storage iscsi-initiator disconnect

Disconnect from an iSCSI target

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

Description

The `storage iscsi-initiator disconnect` command disconnects a node from the specified target. This command is only supported on high-availability shared-nothing virtualized platforms.

Parameters

-node {<nodename>|local} - Node (privilege: advanced)

Specifies the name of the Data ONTAP node from which the iSCSI target will be disconnected.

[-target-type {external|mailbox|partner|partner2|dr_auxiliary|dr_partner}] - Target Type (privilege: advanced)

Selects targets with the specified target type.

-label <text> - User Defined Identifier (privilege: advanced)

Specifies the label of the target to disconnect from.

Examples

The following example adds and connects to an iSCSI target from the specified node.

```
cluster1::*> storage iscsi-initiator disconnect -node node1
               -label target1
```

storage iscsi-initiator remove-target

Remove an iSCSI target

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

Description

The `storage iscsi-initiator remove-target` command removes an iSCSI target from a node's list of targets. This command is only supported on high-availability shared-nothing virtualized platforms.

Parameters

-node {<nodename>|local} - Node (privilege: advanced)

Specifies the name of the Data ONTAP node from which the iSCSI target will be removed.

[-target-type {external|mailbox|partner|partner2|dr_auxiliary|dr_partner}] - Target Type (privilege: advanced)

Selects targets with the specified target type.

-label <text> - User Defined Identifier (privilege: advanced)

Specifies the label of the target to be removed.

Examples

The following example adds and connects to an iSCSI target from the specified *node*.

```
cluster1::*> storage iscsi-initiator remove-target -node node1
               -label target1
```

storage iscsi-initiator show

Display the iSCSI targets

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

Description

The `storage iscsi-initiator show` displays the list of iSCSI targets configured for each Data ONTAP node in the cluster. This command is only supported on high-availability shared-nothing virtualized platforms.

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 (privilege: advanced)

Represents the name of the Data ONTAP node for which information is to be displayed. If this parameter is not specified, the command displays information about all nodes in the cluster.

[-target-type {external|mailbox|partner|partner2|dr_auxiliary|dr_partner}] - Target Type (privilege: advanced)

Selects targets with the specified target type.

[-label <text>] - User Defined Identifier (privilege: advanced)

Selects targets with the specified label.

[-target-portal <text>] - Target Portal (privilege: advanced)

Selects targets with the specified portal.

[-target-name <text>] - iSCSI Name (privilege: advanced)

Selects targets with the specified target name.

[-status-admin {down|up}] - Administrative Status (privilege: advanced)

Selects targets with the specified administrative status.

[-status-oper {down|up}] - Operational Status (privilege: advanced)

Selects targets with the specified operational status.

[-failure-reason <text>] - Failure Reason (privilege: advanced)

Selects targets with the specified failure reason.

Examples

The following example displays the list of iSCSI targets for each node in the cluster.

```
cluster1::*> storage iscsi-initiator show

Status
Node Type Label      Target Portal      Target Name
Admin/Op
-----
-----
node1
    mailbox
        mediator 10.235.14.141      iqn.2012-05.local:mailbox.group.1
up/up
    partner
        partner 10.63.7.205:65200  iqn.2012-06.com.bsdctl:target0
up/up
node2
    mailbox
        mediator 10.235.14.141      iqn.2012-05.local:mailbox.group.1
up/up
    partner
        partner 10.63.7.201:65200  iqn.2012-06.com.bsdctl:target0
up/up
4 entries were displayed.
```

storage path commands

storage path quiesce

Quiesce I/O on a path to array

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

Description

The `storage path quiesce` command quiesces I/O on one path to a LUN. It also quiesces the given entire path immediately or can monitor the given path for error threshold before quiesce. After the I/O has been quiesced, no new I/O is sent on the path unless the [storage path resume](#) command is issued to continue I/O.

Parameters

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

The name of the clustered node for which information is being displayed.

-initiator <initiator name> - Initiator Port

Initiator port that the clustered node uses.

-target-wwpn <wwpn name> - Target Port

Target World Wide Port Name. Port on the storage array that is being used.

{ [-lun-number <integer>] - LUN Number

Logical Unit number. The range is: [0...65535]. If this parameter is not specified, Data ONTAP resumes the entire path to an array.

| [-path-failure-threshold <integer>] - Max Number of Path Failures Acceptable During wait-duration

The path failure count, exceeding this value within wait duration will quiesce the path.

[-wait-duration <integer>] - Wait Duration in minutes }

The time duration(minutes) in which path is monitored for path failures.

Examples

The following example suspends I/O between node vbv3170f1b, port 0a and the array port 50001fe1500a8669, LUN 1.

```
node::> storage path quiesce -node vbv3170f1b -initiator 0a -target-wwpn
50001fe1500a8669 -lun-number 1
```

The following example suspends I/O immediately between node vbv3170f1b, port 0a and the array port 50001fe1500a8669.

```
node::> storage path quiesce -node vbv3170f1b -initiator 0a -target-wwpn
50001fe1500a8669
```

The following example suspends I/O between node vbv3170f1b, port 0a and the array port 50001fe1500a8669 after reaching 10 or more errors in duration of 5 mins.

```
node::> storage path quiesce -node vbv3170f1b -initiator 0a -target-wwpn
50001fe1500a8669 -path-failure-threshold 10 -wait-duration 5
```

Related Links

- [storage path resume](#)

storage path resume

Resume I/O on a path to array

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

Description

The `storage path resume` command continues I/O flow to an array LUN on a path or the entire path that was previously quiesced. It also disables the path failures monitoring feature, if it was enabled using the `storage path quiesce`-path-failure-threshold`count` command.

Parameters

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

The name of the clustered node for which information is being displayed.

-initiator <initiator name> - Initiator Port

Initiator port that the clustered node uses.

-target-wwpn <wwpn name> - Target Port

Target World Wide Port Name. Port on the storage array that is being used.

[-lun-number <integer>] - LUN Number

Logical Unit number. The range is: [0...65535]. If this parameter is not specified, Data ONTAP resumes the entire path to an array.

Examples

The following example resumes I/O between node vbv3170f1b, port 0a and the array port 50001fe1500a8669, LUN 1

```
node::> storage path resume -node vbv3170f1b -initiator 0a -target-wwpn
50001fe1500a8669 -lun-number 1
```

The following example resumes I/O between node vbv3170f1b, port 0a and the array port 50001fe1500a8669

```
node::> storage path resume -node vbv3170f1b -initiator 0a -target-wwpn
50001fe1500a8669
```

Related Links

- [storage path quiesce](#)

storage path show-by-initiator

Display a list of paths to attached arrays from the initiator's perspective

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

Description

The `storage path show-by-initiator` command displays path based statistics. The output is similar to the `storage path show` command but the output is listed by initiator.

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}] - Controller name

The name of the clustered node for which information is being displayed.

[-initiator <text>] - Initiator Port

Initiator port that the clustered node uses.

[-target-wwpn <text>] - Target Port

Target World Wide Port Name. Port on the storage array that is being used.

[-initiator-side-switch-port <text>] - Initiator Side Switch Port

Switch port connected to the clustered node.

[-target-side-switch-port <text>] - Target Side Switch Port

Switch port connected to the array.

[-array-name <array name>] - Array Name

Name of the storage array that is connected to the cluster.

[-tpgn <integer>] - Target Port Group Number

TPGN refers to the target port group to which the target port belongs. A target port group is a set of target ports which share the same LUN access characteristics and failover behaviors.

[-port-speed <text>] - Port Speed

Port Speed of the specified port.

[-path-io-kbps <integer>] - Kbytes of I/O per second on Path (Rolling Average)

Rolling average of I/O per second on the path.

[-path-iops <integer>] - Number of IOPS on Path (Rolling Average)

Rolling average of Kbytes of I/O per second on the path

[-initiator-io-kbps <integer>] - Kbytes of I/O per second on Initiator (Rolling Average)

Rolling average of I/O per second on the initiator port.

[-initiator-iops <integer>] - Number of IOPS on Initiator (Rolling Average)

Rolling average of Kbytes of I/O per second on the initiator port.

[-target-io-kbps <integer>] - Kbytes of I/O per second to Target (Rolling Average)

Rolling average of I/O per second on the target port.

[`-target-iops <integer>`] - Number of IOPS to Target (Rolling Average)

Rolling average of Kbytes of I/O per second on the target port.

Examples

```
vnv3070f20b::> storage path show-by-initiator
Node: vnv3070f20b

Initiator I/O      Initiator Side    Path I/O      Target
Side  Target I/O
Initiator      (KB/s)      Switch Port    (KB/s)      Switch
Port          (KB/s)      Target Port Array Name
-----
0a              3 vnbr3850s4:4              3
vnbr3850s5:15      3 200600a0b819e16f IBM_1722_1
                                0
vnbr3850s5:12      0 50060e80004291c0 HITACHI_DF600F_1
0c              35 vnci9124s54:1-6              35
vnci9124s54:1-24    35 200700a0b819e16f IBM_1722_1
                                0
vnci9124s54:1-22    0 50060e80004291c2 HITACHI_DF600F_1
4 entries were displayed.
```

Related Links

- [storage path show](#)

storage path show

Display a list of paths to attached arrays.

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

Description

The `storage path show` command displays path based statistics. The default command shows:

- Node name
- Initiator port
- Target port
- Target IQN
- TPGN
- Port speeds
- Path I/O in Kbytes/sec
- IOPs

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.

| [-array]

Using this option displays:

- Array name
- Target port
- Target IQN
- Target I/O in Kbytes/sec
- Target side switch port
- Path I/O in Kbytes/sec
- Initiator side switch port
- Initiator I/O in Kbytes/sec
- Initiator port

| [-by-target]

Using this option displays the same information as the array option, but grouped by target port.

| [-detail]

Using this option displays the same information as the array and by-target options, but adds the following:

- Target IOPs
- Target LUNs
- Path IOPs
- Path errors
- Path quality
- Path LUNs
- Initiator IOPs
- Initiator LUNs

| [-switch]

Using this option adds switch port information to the default display.

| [-instance] }

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

[-node {<nodename>|local}] - Controller name

The name of the clustered node for which information is being displayed.

[-array-name <array name>] - Array Name

Name of the storage array that is connected to the cluster.

[-target-wwpn <text>] - Target Port

Target World Wide Port Name. Port on the storage array that is being used.

[-initiator <text>] - Initiator Port

Initiator port that the clustered node uses.

[-initiator-side-switch-port <text>] - Initiator Side Switch Port

Switch port connected to the clustered node.

[-tpgn <integer>] - Target Port Group Number

TPGN refers to the target port group to which the target port belongs. A target port group is a set of target ports which share the same LUN access characteristics and failover behaviors.

[-port-speed <text>] - Port Speed

Port Speed of the specified port.

[-path-io-kbps <integer>] - Kbytes of I/O per second on Path (Rolling Average)

Rolling average of I/O per second on the path.

[-path-iops <integer>] - Number of IOPS on Path (Rolling Average)

Rolling average of Kbytes of I/O per second on the path

[-initiator-io-kbps <integer>] - Kbytes of I/O per second on Initiator (Rolling Average)

Rolling average of I/O per second on the initiator port.

[-initiator-iops <integer>] - Number of IOPS on Initiator (Rolling Average)

Rolling average of Kbytes of I/O per second on the initiator port.

[-target-io-kbps <integer>] - Kbytes of I/O per second to Target (Rolling Average)

Rolling average of I/O per second on the target port.

[-target-iops <integer>] - Number of IOPS to Target (Rolling Average)

Rolling average of Kbytes of I/O per second on the target port.

[-target-side-switch-port <text>] - Target Side Switch Port

Switch port connected to the array.

[-path-link-errors <integer>] - Link Error count on path

Fibre Channel link error count.

[-path-quality <integer>] - Percentage of weighted error threshold

A number representing the threshold of errors that is allowed on the path. Path quality is a weighted error value. When the error weight of a path exceeds the threshold, I/O is routed to a different path.

[-path-lun-in-use-count <integer>] - Number of LUNs in the in-use state on this path

Number of LUNs on this path.

[-initiator-lun-in-use-count <integer>] - Number of LUNs in the in-use state on this initiator

Number of LUNs on this initiator.

[-target-lun-in-use-count <integer>] - Number of LUNs in the in-use state on this target

Number of LUNs on this target.

[-vmdisk-device-id <integer>] - Virtual disk device ID

Common device identifier, shared by a VM and its hypervisor, of a virtual disk. On ESX servers, this is the Disk ID component of a virtual device node, with a value of 0 to 15.

[-path-failure-threshold <integer>] - Max number of path failures acceptable in wait-duration

The path failure count, exceeding this value within wait duration will quiesce the path.

[-wait-duration <integer>] - Wait Duration in minutes

The time duration(minutes) in which path is monitored for path failures.

Examples

The following example shows the default display.

```
vbv3170f2a::> storage path show
```

| Path I/O | | Initiator | Array Target | Port | TPGN | Speed |
|---------------|-------|-----------|------------------|-------|-------|--------|
| Node | IOPS | | | | | |
| (KB/s) | | | | | | |
| ----- | ----- | ----- | ----- | ----- | ----- | ----- |
| vbv3170f2a-01 | | 0b | 50001fe1500a866c | | 2 | 2 Gb/s |
| 6 | 2 | | | | | |
| vbv3170f2a-01 | | 0b | 50001fe1500a866d | | 2 | 2 Gb/s |
| 0 | 0 | | | | | |
| vbv3170f2a-01 | | 0c | 50001fe1500a866e | | 4 | 4 Gb/s |
| 0 | 0 | | | | | |
| vbv3170f2b-03 | | 0a | 50001fe1500a866d | | 1 | 2 Gb/s |
| 3 | 1 | | | | | |
| vbv3170f2b-03 | | 0c | 50001fe1500a866f | | 4 | 4 Gb/s |
| 3 | 1 | | | | | |

5 entries were displayed.

The following example shows how the information is displayed with the array option.

```
vnv3070f20b::> storage path show -array
```

```
Node: vnv3070f20b
```

| Path I/O | Initiator Side | Target I/O | Target Side |
|------------------|------------------|---------------|------------------|
| Array Name | Target Port | Initiator I/O | Initiator |
| (KB/s) | Switch Port | (KB/s) | Switch Port |
| Port | | Port | |
| HITACHI_DF600F_1 | 50060e80004291c0 | 0 | vnbr3850s5:12 |
| 0 | vnbr3850s4:4 | 3 | 0a |
| | 50060e80004291c2 | 0 | vnci9124s54:1-22 |
| 0 | vnci9124s54:1-6 | 26 | 0c |
| IBM_1722_1 | 200600a0b819e16f | 3 | vnbr3850s5:15 |
| 3 | vnbr3850s4:4 | 3 | 0a |
| | 200700a0b819e16f | 26 | vnci9124s54:1-24 |
| 26 | vnci9124s54:1-6 | 26 | 0c |

4 entries were displayed.

The following example shows how the information is displayed when grouped by target.

```

vnv3070f20b::> storage path show -by-target
Node: vnv3070f20b
Array Name: HITACHI_DF600F_1

```

| Initiator Side | Target I/O (KB/s) | Initiator I/O (KB/s) | Target Side Initiator | Path I/O (KB/s) |
|------------------|----------------------|-------------------------|--------------------------|--------------------|
| Target Port | | | Switch Port | |
| Switch Port | (KB/s) | Port | | |
| 50060e80004291c0 | 0 | | vnbr3850s5:12 | 0 |
| vnbr3850s4:4 | 3 | 0a | | |
| 50060e80004291c2 | 0 | | vnci9124s54:1-22 | 0 |
| vnci9124s54:1-6 | 26 | 0c | | |

```

Node: vnv3070f20b
Array Name: IBM_1722_1

```

| Initiator Side | Target I/O (KB/s) | Initiator I/O (KB/s) | Target Side Initiator | Path I/O (KB/s) |
|------------------|----------------------|-------------------------|--------------------------|--------------------|
| Target Port | | | Switch Port | |
| Switch Port | (KB/s) | Port | | |
| 200600a0b819e16f | 3 | | vnbr3850s5:15 | 3 |
| vnbr3850s4:4 | 3 | 0a | | |
| 200700a0b819e16f | 26 | | vnci9124s54:1-24 | 26 |
| vnci9124s54:1-6 | 26 | 0c | | |

4 entries were displayed.

The following example shows how the information is displayed with the switch option.

```
vbv3170f2b::> storage path show -switch
```

| Initiator Side | | | Path I/O | Target Side |
|----------------|-----------|------------------|-------------|----------------|
| Node | Initiator | Array | Target Port | Switch Port |
| Switch Port | TPGN | Speed | (KB/s) | IOPS |
| ----- | ----- | ----- | ----- | ----- |
| ----- | | | | |
| vbv3170f2a-01 | 0b | 50001fe1500a866c | | vbbr300s1:6 |
| vbbr300s1:2 | 2 | 2 Gb/S | 9 | 3 |
| vbv3170f2a-01 | 0b | 50001fe1500a866d | | vbbr300s1:7 |
| vbbr300s1:2 | 2 | 2 Gb/S | 0 | 0 |
| vbv3170f2a-01 | 0c | 50001fe1500a866e | | vbc19124s2:1-7 |
| vbc19124s2:1-3 | 4 | 4 Gb/S | 0 | 0 |
| vbv3170f2b-03 | 0a | 50001fe1500a866d | | vbbr300s1:7 |
| vbbr300s1:3 | 1 | 2 Gb/S | 4 | 1 |
| vbv3170f2b-03 | 0c | 50001fe1500a866f | | vbc19124s2:1-8 |
| vbc19124s2:1-4 | 4 | 4 Gb/S | 4 | 1 |

5 entries were displayed.

storage pool commands

storage pool add

Add disks to a storage pool

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

Description

The `storage pool add` command increases the total capacity of an existing storage pool by adding the specified SSDs to the storage pool. The disks are split into four equal partitions and added to each of the allocation units of the storage pool. If any allocation units from the storage pool have already been allocated to an aggregate, the cache capacity of that aggregate is increased automatically.

For example, if an SSD with a usable size of 745 GB is added to a storage pool that is part of four aggregates, each aggregate will grow its cache capacity by 186.2 GB. If a different allocation is desired, create a new storage pool using the [storage pool create](#) command.



Storage pools are not supported in all-flash aggregates.

Parameters

-storage-pool <storage pool name> - Storage Pool Name

This parameter specifies the storage pool to which disks are to be added.

{ [-disk-count <integer>] - Number of Disks to Add in Storage Pool

This parameter specifies the number of disks that are to be added to the storage pool. The disks to be added come from the pool of spare disks.

[-nodes {<nodename>|local}] - Nodes From Which Spares Should be Selected

This parameter specifies a list of nodes from which SSD disks are selected for addition to the storage pool. If this parameter is not specified, disks to be added to the storage pool can be selected from both the nodes sharing the storage pool. Use this parameter to restrict the selection of spare disks to one particular node.

[-disk-list <disk path name>, ...] - List of Spare Disks }

This parameter specifies a list of disks to be added to the storage pool. In an HA configuration, SSDs being added to a storage pool can be owned by either node in the HA pair.

{ [-quiet <true>] - Confirmations off (privilege: advanced)

When set to *true*, this parameter specifies the operation should be executed without pausing for confirmation.

[-simulate <true>] - Simulate Storage Pool Addition

When set to *true*, this parameter specifies the operation should be performed as a simulation. The command reports which aggregates would grow automatically as a result of adding the disks to the storage pool. The disks are not added to the storage pool.

Examples

In this example, the user requests a report detailing the changes that would occur if a new disk is added to the storage pool *SP1*. In this case, 186.2 GB of cache is added to the Flash Pool aggregates *nodeA_flashpool_1* and *nodeB_flashpool_1*. There are two unprovisioned allocation units in the storage pool and therefore the storage pool available capacity also grows by 372.5 GB.

```
cluster1::> storage pool add -storage-pool SP1 -disk-list 1.0.23 -simulate
```

This operation will result in capacity being allocated in the following way:

| Container Name | Capacity To Be Added | Current Size | New Size |
|----------------------|-------------------------|-----------------|-------------|
| ----- | ----- | ----- | ----- |
| nodeA_flashpool_1 | 186.2GB | 558.7GB | 744.9GB |
| nodeB_flashpool_1 | 186.2GB | 558.7GB | 744.9GB |
| (Available Capacity) | 372.5GB | 1.09TB | 1.45TB |

The following example adds one disk to a storage pool named *SP1*. The spare disks are selected from either local node or its partner or both based on spare availability.

```
cluster-1::> storage pool add -storage-pool SP1 -disk-count 1

Info: The following disks will be added to storage pool "SP1":
Disk                Size  Type  Owner
-----
1.0.12              744.9GB SSD    cluster-1-01
New Allocation Unit Size: 744.8GB
Capacity will be allocated in the following way:
Container           Capacity  Current      New
Name               To Be Added      Size      Size
-----
nodeA_flashpool_1   186.2GB    558.7GB    744.9GB
nodeB_flashpool_1   186.2GB    558.7GB    744.9GB
(Available Capacity) 372.5GB    1.09TB     1.45TB
Are you sure you want to continue with this operation?
{y|n}: y
[Job 48] Job succeeded: storage pool add job for "SP1" completed
successfully
```

Related Links

- [storage pool create](#)

storage pool create

Create a new storage pool

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

Description

The `storage pool create` command creates an SSD storage pool using a given list of spare SSDs.

When a storage pool is created, Data ONTAP splits the capacity provided by the SSDs into four equally-sized allocation units. In an HA configuration, two allocation units (containing 50% of the total capacity) are assigned to each node in the HA pair. This assignment can be modified using the [storage pool reassign](#) command.

After the storage pool is created, its allocation units can be provisioned into Flash Pool aggregates using the [storage aggregate add-disks](#) command and the `-storage-pool` parameter.



Storage pools are not supported in All-Flash aggregates.

Parameters

-storage-pool <storage pool name> - Storage Pool Name

This parameter specifies the name of the storage pool that is to be created. The SSDs are partitioned and placed into the new storage pool.

{ [-nodes {<nodename>|local}] - Nodes Sharing the Storage Pool

This parameter specifies a list of nodes from which SSD disks are selected to create the storage pool. If two nodes are specified then they need to be in HA configuration. Spare disks are selected from either node or its partner or both. If this parameter is not specified, storage pool will be created by selecting disks from either the node or its partner or both from where command is run.

-disk-count <integer> - Number of Disks in Storage Pool

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

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

This parameter specifies the size of the disks on which the storage pool is to be created. Disks with a usable size between 95% and 105% of the specified size are selected.

[-disk-list <disk path name>,... - Disk List for Storage Pool Creation }

This parameter specifies a list of SSDs to be included in the new storage pool. The SSDs must be spare disks and can be owned by either node in an HA pair.

[-simulate <true>] - Simulate Storage Pool Creation

This option simulates the storage pool creation and prints the allocation unit size that would be used for the storage pool.

Examples

The following example creates a storage pool named SP1. The storage pool contains 3 SSD disks, the spare disks selected are from either local node, or its partner or both based on spare availability.

```
cluster1::> storage pool create -storage-pool SP1 -disk-count 3
```

The following example creates a storage pool named SP2. The storage pool contains 3 SSD disks, the spare disks selected are from either node0, or its partner node1 or both based on spare availability.

```
cluster1::> storage pool create -storage-pool SP2 -disk-count 3 -nodes  
node0,node1
```

The following example creates a storage pool named SP3 from four SSDs using disk list.

```
cluster1::> storage pool create -storage-pool SP3 -disk-list 1.0.13,  
1.0.15, 1.0.17, 1.0.19
```

Related Links

- [storage pool reassign](#)
- [storage aggregate add-disks](#)

storage pool delete

Delete an existing storage pool

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

Description

The `storage pool delete` command deletes an existing SSD storage pool. At the end of the operation, the SSDs are converted back to spare disks.

Parameters

-storage-pool <storage pool name> - Storage Pool Name

This parameter specifies the storage pool that you want to delete. You can delete the storage pool only if all of the allocation units in the storage pool are available.

Examples

Verify that storage pool *SP3* is ready for deletion by confirming it has four available allocation units and then delete it.

```
cluster1::> storage pool show-available-capacity -storage-pool SP3
```

| Node | Storage Pool | Storage Type | SyncMirror Pool | Allocation Unit size | Unit Count | Total Usable Size |
|--------|--------------|--------------|-----------------|----------------------|------------|-------------------|
| node-a | SP3 | SSD | Pool0 | 372.5GB | 2 | 744.9GB |
| node-b | SP3 | SSD | Pool0 | 372.5GB | 2 | 744.9GB |

```
2 entries were displayed.
```

```
cluster1::> storage pool delete -storage-pool SP3
```

```
Warning: Are you sure you want to delete storage pool "SP3"? {y|n}: y
[Job 313] Job succeeded: storage pool delete job for "SP3" completed
successfully
```

storage pool reassign

Reassign capacity from one node to another node in storage pool

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

Description

The `storage pool reassign` command changes the ownership of unprovisioned (available) storage pool allocation units from one HA partner to the other for an existing storage pool.

Parameters

-storage-pool <storage pool name> - Storage Pool Name

This parameter specifies the storage pool within which available capacity is reassigned from one node to another.

-from-node {<nodename>|local} - Reassign Available Capacity from This Node

This parameter specifies the name of the node that currently owns the allocation units.

-to-node {<nodename>|local} - Reassign Available Capacity to This Node

This parameter specifies the name of the node that will now own the allocation units.

-allocation-units <integer> - Allocation Units

This parameter specifies the number of allocation units to be reassigned.

Examples

Move an available allocation unit from node-b to node-a in preparation for provisioning the allocation units on node-a.

```
cluster1::*> storage pool show-available-capacity -storage-pool SP2
```

| Node | Storage Pool | Storage Type | SyncMirror Pool | Allocation Unit size | Unit Count | Total Usable Size |
|--------|--------------|--------------|-----------------|----------------------|------------|-------------------|
| node-a | SP2 | SSD | Pool0 | 744.9GB | 2 | 1.45TB |
| node-b | SP2 | SSD | Pool0 | 744.9GB | 1 | 744.9GB |

2 entries were displayed.

```
cluster1::*> storage pool reassign -storage-pool SP2 -from-node node-b -to
-node node-a -allocation-units 1
[Job 310] Job succeeded: storage pool reassign job for "SP2" completed
successfully
```

```
cluster1::*> storage pool show-available-capacity -storage-pool SP2
```

| Node | Storage Pool | Storage Type | SyncMirror Pool | Allocation Unit size | Unit Count | Total Usable Size |
|--------|--------------|--------------|-----------------|----------------------|------------|-------------------|
| node-a | SP2 | SSD | Pool0 | 744.9GB | 3 | 2.18TB |
| node-b | SP2 | SSD | Pool0 | 744.9GB | 0 | 0B |

2 entries were displayed.

storage pool rename

Rename storage pool

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

Description

The `storage pool rename` command changes the name of the storage pool.

Parameters

-storage-pool <storage pool name> - Storage Pool Name

This parameter specifies the storage pool name.

-new-name <storage pool name> - New Name of the Storage Pool

This parameter specifies the new name of the storage pool.

Examples

Renaming storage pool "sp1" to "sp2"

```
cluster-1::> storage pool show
Storage
Storage Pool      Type      #Disks  Nodes              Total Size
-----
sp1                SSD           4 node-a,
                   node-b              10.44GB

cluster-1::> storage pool rename -storage-pool sp1 -new-name sp2

cluster-1::> storage pool show
Storage
Storage Pool      Type      #Disks  Nodes              Total Size
-----
sp2                SSD           4 node-a,
                   node-b              10.44GB
```

storage pool show-aggregate

Display aggregates provisioned from storage pools

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

Description

The `storage pool show-aggregate` command displays allocation information for SSD storage pools in the cluster. The command output depends upon the parameter or parameters specified with the command. If no parameters are specified, the command displays information about allocations of all storage pools 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.

[-storage-pool <storage pool name>] - Name of Storage Pool

Selects the storage pools that match this parameter value.

[-aggregate <aggregate name>] - Aggregate

Selects the storage pools that match this parameter value.

[-capacity {<integer>[KB|MB|GB|TB|PB]}] - Capacity

Selects the storage pools that match this parameter value.

Capacity includes space provided by data and parity portions of each allocation unit. Only the data portions of each allocation unit contribute to the cache or usable capacity of Flash Pool.

[-allocated-unit-count <integer>] - Number of AU's Assigned to This Aggregate

Selects the storage pools that match this parameter value.

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

Selects the storage pools that match this parameter value.

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

Selects the storage pools that match this parameter value.

Examples

Display information about the aggregate or aggregates using a storage pool called *SP2* :

```
cluster1::> storage pool show-aggregate -storage-pool SP2 -instance
Name of Storage Pool: SP2
                        Aggregate: node2_flashpool_1
                        Capacity: 744.9GB
Number of AU's Assigned to This Aggregate: 1
                        Original Owner Name: node2
                        Node: node2
```

storage pool show-available-capacity

Display available capacity of storage pools

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

Description

The `storage pool show-available-capacity` command displays information about available capacity in SSD storage pools on each node in the cluster. The command output depends upon the parameter or parameters specified with the command. If no parameters are specified, the command displays information about available capacities in all shared pools in the cluster.

Storage pool available capacity is data storage space that has not yet been provisioned into Flash Pool. Allocation units might be provisioned into aggregates using the [storage aggregate add-disks](#) command and the `-storage-pool` parameter.



All storage pool available capacity can be provisioned into aggregates. Available capacity within a storage pool is not used to protect against a disk failure. In the case of an SSD failure or predicted failure, Data ONTAP moves a suitable whole SSD spare disk from outside the storage pool into the storage pool and begins the recovery process (using either reconstruction or Rapid RAID Recovery, whichever is appropriate).

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.

[-storage-pool <storage pool name>] - Name of Storage Pool

Selects the available capacities that match this parameter value.

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

Selects the available capacities that match this parameter value.

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

Selects the available capacities that match this parameter value.

Allocation units are the units of storage capacity that are available to be provisioned into aggregates.

[-storage-type <SSD>] - Type of Storage Pool

Selects the available capacities that match this parameter value. Only the `SSD` type is supported for this version of Data ONTAP.

[-allocation-unit-count <integer>] - Number of Allocation Units Available

Selects the available capacities that match this parameter value.

Allocation units are the units of storage capacity that are available to be provisioned into aggregates. Each allocation unit is one minimum unit of allocation (MUA) and its capacity is given as `allocation-unit-size`.

[-syncmirror-pool <text>] - SyncMirror Pool

Selects the available capacities that match this parameter value.

The SyncMirror pool of an allocation unit must match the SyncMirror pool of the disks of the aggregate when adding allocation units into an aggregate.

Mirroring of aggregates that are provisioned from SSD storage pools is not supported.

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

Selects the available capacities that match this parameter value.

The `available-size` is the sum of the capacities of the allocation units that are assigned but not yet provisioned. The amount of `available-size` that is contributed to the cache or usable capacity of an aggregate depends upon the RAID type used when provisioning the allocation units.

Examples

In this example, two nodes of an HA pair share available capacity from two storage pools, *SP1* and *SP2* . There are a total of 5 allocation units that have not yet been provisioned.

```
cluster1::> storage pool show-available-capacity
```

| | | Storage | SyncMirror | Allocation | Unit | Total |
|--------|--------------|---------|------------|------------|-------|-------------|
| Node | Storage Pool | Type | Pool | Unit size | Count | Usable Size |
| node-a | SP1 | SSD | Pool0 | 558.7GB | 1 | 558.7GB |
| node-b | SP1 | SSD | Pool0 | 558.7GB | 1 | 558.7GB |
| node-a | SP2 | SSD | Pool0 | 744.9GB | 2 | 1.45TB |
| node-b | SP2 | SSD | Pool0 | 744.9GB | 1 | 744.9GB |

Related Links

- [storage aggregate add-disks](#)

storage pool show-disks

Display disks in storage pools

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

Description

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

[-storage-pool <storage pool name>] - Name of Storage Pool

Selects the storage pools that match this parameter value.

[-disk <disk path name>] - Name of the disk

Selects the storage pools with the disks that match this parameter value.

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

Selects the storage pools with the disks that match this parameter value. Only the `SSD` type is supported for this version of Data ONTAP.

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

Selects the storage pools with the disks that match this parameter value.

In this command, `usable-size` refers to the sum of the capacities of all of the partitions on the disk.

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

Selects the storage pools with the disks that match this parameter value.

[-node-list <nodename>,...] - List of Nodes

Selects the storage pools with the disks that are visible to all of the specified nodes.

Examples

Show information about SSDs in a storage pool called `SP2`.

```
cluster1::> storage pool show-disks -storage-pool SP2
```

Storage Pool Name: SP2

Storage

| Disk | Type | Usable Size | Total Size |
|--------|-------|-------------|------------|
| ----- | ----- | ----- | ----- |
| 1.0.16 | SSD | 745.0GB | 745.2GB |
| 1.0.18 | SSD | 745.0GB | 745.2GB |
| 1.0.20 | SSD | 745.0GB | 745.2GB |
| 1.0.22 | SSD | 745.0GB | 745.2GB |

storage pool show

Display details of storage pools

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

Description

The `storage pool show` command displays information about SSD storage pools in the cluster. By default, the command displays information about all storage pools in the cluster. You can specify parameters to limit the output to a specific set of storage pools.

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.

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

Selects the storage pools that match this parameter value.

[-storage-pool-uuid <UUID>] - UUID of Storage Pool

Selects the storage pools that match this parameter value.

[-nodes {<nodename>|local}] - Nodes Sharing the Storage Pool

Selects the storage pools that match this parameter value.

In an HA pair, either node name may be specified.

[-disk-count <integer>] - Number of Disks in Storage Pool

Selects the storage pools that match this parameter value.

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

Selects the storage pools that match this parameter value.

Allocation units represent the unit of storage allocated to aggregates from this storage pool.

[-allocation-unit-data-size-raid4 {<integer>[KB|MB|GB|TB|PB]}] - Allocation Unit Data Size for RAID4

This parameter shows the amount of additional data capacity provided if an allocation unit from this storage pool was added to an aggregate with `-raidtype` as `raid4`.

[-allocation-unit-data-size-raid-dp {<integer>[KB|MB|GB|TB|PB]}] - Allocation Unit Data Size for RAID-DP

This parameter shows the amount of additional data capacity provided if an allocation unit from this storage pool was added to an aggregate with `-raidtype` as `raid_dp`.

[-allocation-unit-data-size-raid-tec {<integer>[KB|MB|GB|TB|PB]}] - Allocation Unit Data Size for RAID-TEC

This parameter shows the amount of additional data capacity provided if an allocation unit from this storage pool was added to an aggregate with `-raidtype` as `raid_tec`.

[`-storage-type <SSD>`] - Storage Type

Selects the storage pools that match this parameter value.

Only the `SSD` type is supported for this version of Data ONTAP.

[`-pool-usable-size {<integer>[KB|MB|GB|TB|PB]}`] - Storage Pool Usable Size

Selects the storage pools that match this parameter value.

The `pool-usable-size` is the sum of the capacities of the allocation units that are assigned to nodes but not yet provisioned. The amount of `pool-usable-size` that is contributed to the cache or usable capacity of an aggregate depends upon the RAID type used when provisioning the allocation units.

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

Selects the storage pools that match this parameter value.

The `pool-total-size` is the sum of the capacities of allocation units belonging to this storage pool.

[`-is-healthy {true|false}`] - Is Pool Healthy?

Selects the storage pools that match this parameter value.

For storage pools with `is-healthy`false` , the ``unhealthy-reason` parameter provides more information.

`is-healthy` must be `true` to provision allocation units from a storage pool into an aggregate.

[`-pool-state <State of the Storage Pool>`] - State of the Storage Pool

Selects the storage pools that match this parameter value. Possible states are:

- `normal` - the storage pool is operating normally.
- `degraded` - the storage pool has one or more failed disks.
- `creating` - the storage pool is being created.
- `deleting` - the storage pool is being deleted.
- `reassigning` - allocation units are being reassigned from one node to another.
- `growing` - allocation units in the storage pool are expanding due to the addition of new capacity into the storage pool.

[`-unhealthy-reason <text>`] - Reason for Storage Pool Being Unhealthy

Selects the storage pools that match this parameter value.

The message provided gives additional details about why the storage pool is unhealthy.

[`-current-operation-job-id <integer>`] - Job ID of the Currently Running Operation

Selects the storage pools that match this parameter value.

Long-running operations associated with storage pools will be managed via jobs. For example, if you provision allocation units from a storage pool into an aggregate and the disks associated with the storage pool need to be zeroed, the operation will be completed via a job.

Examples

Display the storage pools in the cluster.

```
cluster1::> storage pool show
Storage Pool      Type  #Disks Nodes              Total Size
-----
LargeSP           SSD    10 noda-a,node-b      7.27TB
SmallSP           SSD     2 noda-a,node-b      1.45TB
2 entries were displayed.
```

The following example displays the details of a storage pool named SmallSP. Only one of its four allocation unit has been provisioned, so 75% of its size is available (usable).

```
cluster1::> storage pool show -storage-pool SmallSP
Storage Pool Name: SmallSP
                UUID of Storage Pool: 60f2f1b9-e60f-11e3-a5e7-
00a0981899a2
        Nodes Sharing the Storage Pool: node-a, node-b
    Number of Disks in Storage Pool: 2
        Allocation Unit Size: 372.5GB
                Storage Type: SSD
        Storage Pool Usable Size: 1.09TB
        Storage Pool Total Size: 1.45TB
                Is Pool Healthy?: true
        State of the Storage Pool: normal
    Reason for storage pool being unhealthy: -
    Job ID of the Currently Running Operation: -
```

storage port commands

storage port disable

Disable a storage port

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

Description

The `storage port disable` command disables a specified storage port.

Parameters

-node {<nodename>|local} - Node

Use this parameter to specify the node on which the port resides.

-port <text> - Port

Use this parameter to specify the port that needs to be disabled.

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

Use this optional parameter to force the disabling of the storage port. The parameter can be used to disable the specified port even if some devices can only be accessed using this port. Note that doing so might cause multiple device failures.

Examples

The following example disables port 0a on node node1:

```
cluster1::> storage port disable -node node1 -port 0a
```

storage port enable

Enable a storage port

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

Description

The `storage port enable` command enables a specified storage port.

Parameters

-node {<nodename>|local} - Node

Use this parameter to specify the node on which the port resides.

-port <text> - Port

Use this parameter to specify the port that needs to be enabled.

Examples

The following example enables port 0a on node node1:

```
cluster1::> storage port enable -node node1 -port 0a
```

storage port modify

Modify a storage port

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

Description

The `storage port modify` command modifies the attributes of a storage port.

Parameters

-node {<nodename>|local} - Node

This parameter specifies the name of the node.

-port <text> - Port

This parameter specifies the name of the port.

[-mode {storage|network}] - Mode

This parameter specifies the usage mode for an Ethernet port - storage or network use. The mode parameter only applies to Ethernet ports that are not dedicated to storage connectivity.

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

This parameter allows the operation to override warnings.

Examples

The following example places a nondedicated Ethernet port in storage mode:

```

cluster1::> storage port show
Speed
Node          Port Type  Mode      VLAN      State   Status      ID
-----
cluster1-01
              e0b  ENET  network    - -        -        -
              e0c  ENET  -          100 enabled online      30
              e0e  ENET  storage    0 enabled offline    30
              e0f  ENET  -          100 enabled online    30
cluster1-02
              e0b  ENET  storage    0 enabled offline    30
              e0c  ENET  -          100 enabled online    30
              e0e  ENET  storage    0 enabled offline    30
              e0f  ENET  -          100 enabled online    30
8 entries were displayed.

cluster1::> storage port modify -node cluster1-01 -port e0b -mode storage

cluster1::> storage port show
Speed
Node          Port Type  Mode      VLAN      State   Status      ID
-----
cluster1-01
              e0b  ENET  storage    0 enabled offline    30
              e0c  ENET  -          100 enabled online    30
              e0e  ENET  storage    0 enabled offline    30
              e0f  ENET  -          100 enabled online    30
cluster1-02
              e0b  ENET  storage    0 enabled offline    30
              e0c  ENET  -          100 enabled online    30
              e0e  ENET  storage    0 enabled offline    30
              e0f  ENET  -          100 enabled online    30
8 entries were displayed.

```

storage port rescan

Rescan a storage port

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

Description

The `storage port rescan` command rescans a specified storage port. This command is not supported on Ethernet storage ports (type = ENET).

Parameters

-node {<nodename>|local} - Node (privilege: advanced)

Use this parameter to specify the node on which the port resides.

-port <text> - Port (privilege: advanced)

Use this parameter to specify the port that needs to be rescanned.

Examples

The following example rescans port 0a on node node1:

```
cluster1::> storage port rescan -node node1 -port 0a
```

storage port reset-device

Reset a device behind a storage port

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

Description

The ``storage port reset-device`` command resets a device behind a port. If the device is behind a SAS port, you need to specify the shelf name and bay ID where the device resides. If the device is behind a FC port, you need to specify the loop ID of the device. This command is not supported on Ethernet storage ports (type = ENET).

Parameters

-node {<nodename>|local} - Node (privilege: advanced)

Use this parameter to specify the node on which the port resides.

-port <text> - Port (privilege: advanced)

Use this parameter to specify the port used to reset the device.

{ -shelf-name <text> - Shelf Name (privilege: advanced)

Use this parameter to specify the shelf where the device resides.

-bay-id <integer> - Bay ID (privilege: advanced)

Use this parameter to specify the bay where the device resides.

| -loop-id <integer> - Loop ID (privilege: advanced) }

Use this parameter to specify the loop ID of the device.

Examples

The following example resets a device behind SAS port 0a on node node1:

```
cluster1::> storage port reset-device -node node1 -port 0a -shelf-name 1.0  
-bay-id 10
```

The following example resets a device behind FC port 1b on node node1:

```
cluster1::> storage port reset-device -node node1 -port 1b -loop-id 20
```

storage port reset

Reset a storage port

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

Description

The `storage port reset` command resets a specified storage port. This command is not supported on Ethernet storage ports (type = ENET).

Parameters

-node {<nodename>|local} - Node (privilege: advanced)

Use this parameter to specify the node on which the port resides.

-port <text> - Port (privilege: advanced)

Use this parameter to specify the port that needs to be reset.

Examples

The following example resets port 0a on node node1:

```
cluster1::> storage port reset -node node1 -port 0a
```

storage port show

Show storage port information

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

Description

The `storage port show` command displays information about the storage ports in the cluster. If no parameters are specified, the default command displays the following information about the storage ports:

- Node
- Port
- Type
- Speed
- State
- Status

To display detailed profile information about a single storage port, use the `-node` and `-port` parameters.

Parameters

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

Displays the specified fields for all the storage ports, in column style output.

| [-errors]

Displays the following error status information about the storage ports which have errors:

- Error type
- Error severity
- Error description

| [-instance] }

Displays expanded information about all the storage ports in the system. If a storage port is specified, then this parameter displays detailed information for that port only.

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

Displays detailed information about the storage ports on the specified node.

[-port <text>] - Port

Selects the ports with the specified port name.

[-port-type {Unknown|SAS|FC|ENET}] - Port Type

Selects the ports of the specified type.

[-port-speed {0|1|1.5|2|2.5|3|4|5|6|8|10|12|14|16|25|32|40|100}] - Port Speed

Selects the ports with the specified speed.

[-state {enabled|disabled|enable-pending|disable-pending}] - Port State

Selects the ports with the specified state.

[-status {unknown|online|online-degraded|offline|link-down}] - Port Status

Selects the ports with the specified operational status.

[-description <text>] - Description

Selects the ports with the specified description.

[-firmware-rev <text>] - Firmware Revision

Selects the ports with the specified firmware revision.

[-serial-number <text>] - Serial Number

Selects the ports with the specified serial number.

[-is-dedicated {true|false}] - Is Dedicated Storage Port?

Selects the ports that match the specified value for storage-only ports. This value is always true for FC and SAS ports, as well as for ENET ports that are dedicated to storage.

[-connection-mode {Unknown|Loop|Point-to-point}] - Connection Mode

Selects the ports with the specified connection mode.

[-wwnn <FC WWN>] - World Wide Node Name

Selects the ports with the specified World Wide Node Name.

[-wwpn <FC WWN>] - World Wide Port Name

Selects the ports with the specified World Wide Port Name.

[-board-name <text>] - Board Name

Selects the ports with the specified board name.

[-connector-capabilities <integer>,...] - Connector Capabilities

Selects the ports with the specified list of connector capabilities.

[-wwn <FC WWN>] - Base World Wide Name

Selects the ports with the specified World Wide Name.

[-mfg-part-number <text>] - MFG Part Number

Selects the ports with the specified manufacturer part number.

[-nvdata-rev <text>] - NVDATA Revision

Selects the ports with the specified NVDATA revision.

[-part-number <text>] - Part Number

Selects the ports with the specified part number.

[-date-code <text>] - Date Code

Selects the ports with the specified date code.

[-connector-technology {active-copper|passive-copper|optical}] - Connector Technology

Selects the ports with the specified connector technology.

[-phy-id <integer>,...] - Phy ID

Selects the ports that have phys with the specified phy ID.

[-phy-state {enabled|disabled}] - Phy State

Selects the ports that have phys with the specified state.

[`-phy-status` {`unknown`|`online`|`offline`|`speed-negotiation-failed`|`sata-oob-failed`}] -

Phy Status

Selects the ports that have phys with the specified status.

[`-phy-speed` {`0`|`1`|`1.5`|`2`|`2.5`|`3`|`4`|`5`|`6`|`8`|`10`|`12`|`14`|`16`|`25`|`32`|`40`|`100`}] - Phy Speed

Selects the ports that have phys with the specified speed.

[`-mac-address` <text>] - MAC Address

Selects ports that match the specified MAC address.

[`-vlan-id` <integer>] - VLAN ID

Selects the ports with the specified VLAN ID.

[`-vendor-id` <text>] - Vendor ID

Selects the ports with the specified vendor ID.

[`-vendor-part-id` <text>] - Vendor part ID

Selects the ports with the specified vendor part ID.

[`-device-type` <text>] - Device type

Selects ports that match the specified device type.

[`-error-type` {`unknown`|`online`|`online-degraded`|`offline`|`link-down`}] - Error Type

Selects the ports with the specified error type.

[`-error-severity` {`unknown`|`notice`|`warning`|`error`|`critical`}] - Error Severity

Selects the ports with the specified error severity.

[`-error-text` <text>] - Error Text

Selects the ports with the specified error text.

[`-corrective-action` <text>] - Corrective Action

Selects the ports with the specified corrective action.

[`-cable-length` <text>] - Cable Length

Selects the ports with the specified cable length.

[`-cable-identifier` <text>] - Cable Identifier

Selects the ports with the specified cable identifier.

[`-cable-end-id` {`end_0`|`end_1`}] - Cable End Identifier

Selects the ports with the specified cable end identifier.

[`-connector-type` {`QSFP`|`QSFP+`|`QSFP28`|`Mini-SAS` `HD`|`SFP`}] - Connector Type

Selects the ports with the specified connector type.

[`-connector-vendor` <text>] - Connector Vendor

Selects the ports with the specified connector vendor.

[-connector-part-number <text>] - Connector Part Number

Selects the ports with the specified connector part number.

[-connector-serial-number <text>] - Connector Serial Number

Selects the ports with the specified connector serial number.

[-mode {storage|network}] - Mode

Selects the ports with the specified mode.

Examples

The following example displays information about all storage ports in the cluster:

```
cluster1::> storage port show
Speed
Node          Port Type  Mode   (Gb/s) State   Status   ID
-----
cluster1-01
    4a   SAS   -       0 enabled offline  -
    4b   SAS   -       0 enabled offline  -
    4c   SAS   -       0 enabled offline  -
    4d   SAS   -       0 enabled offline  -
    e3a  ENET  network - -      -        -
    e3b  ENET  storage 100 enabled online   30
    e5a  ENET  storage 100 enabled online   30
    e5b  ENET  network - -      -        -
cluster1-02
    4a   SAS   -       0 enabled offline  -
    4b   SAS   -       0 enabled offline  -
    4c   SAS   -       0 enabled offline  -
    4d   SAS   -       0 enabled offline  -
    e3a  ENET  network - -      -        -
    e3b  ENET  storage 100 enabled online   30
    e5a  ENET  storage 100 enabled online   30
    e5b  ENET  network - -      -        -
16 entries were displayed.
```

The following example displays detailed information about port e5a on node node1:

```
cluster1::> storage port show -node cluster1-01 -port e5a
Node: cluster1-01
      Port: e5a
      Port Type: ENET
      Mode: storage
      Description: 40G/100G Ethernet Controller CX5
      Firmware Revision: 16.23.1020
      MAC Address: ec:0d:9a:65:e4:44
      Is Dedicated: false
      Serial Number: MT1730X00227
      Connector Vendor: Molex Inc.
      Connector Part Number: 112-00322
      Connector Serial Number: 532120266
      Port Speed: 100 Gb/s
      Port State: enabled
      Port Status: online
```

storage raid-options commands

storage raid-options modify

Modify a RAID option

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

Description

The `storage raid-options modify` command is used to modify the available RAID options for each node in a cluster. The options are described in the `storage raid-options` manual page.

Parameters

-node {<nodename>|local} - Node

This parameter specifies the node on which the RAID option is to be modified.

-name <text> - Option Name

This parameter specifies the RAID option to be modified. To see the list of RAID options that can be modified, use the [storage raid-options show](#) command.

[-value <text>] - Option Value

This parameter specifies the value of the selected RAID option.

Related Links

- [storage raid-options show](#)

storage raid-options show

Display RAID options

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

Description

The `storage raid-options show` command displays information about all the RAID options in a cluster. The options are described in the `storage raid-options` manual page.

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

Selects information about all the RAID options on the specified node.

[-name <text>] - Option Name

Selects information about the RAID options that have the specified name.

[-value <text>] - Option Value

Selects information about all the RAID options that have the specified value.

[-recommended-value <text>] - Recommended Value

Selects information about all the RAID options that have the specified recommended value.

[-constraint <text>] - Option Constraint

Selects information about all the RAID options that have the specified constraint. The 'constraint' field indicates the expected setting for a RAID option across both nodes of an HA pair. The possible values are:

- *none* - no constraint on the value of this RAID option; nodes can have different values
- *same_preferred* - the same value should be used on both nodes of an HA pair, otherwise the next takeover may not function correctly
- *same_required* - the same value must be used on both nodes of an HA pair, otherwise the next takeover will not function correctly
- *only_one* - the same value should be used on both nodes of an HA pair. If the values are different and a takeover is in progress, the value of the RAID option on the node that is taking over will be used
- *unknown* - no information about constraints for this RAID option

Examples

The following example shows the RAID scrub settings for a node named node1:


```
cluster1::> storage raid-options show -node node1 -name raid.scrub*

Node      Option                               Current      Recommended
Constraint
-----
node1     raid.scrub.perf_impact                low          low
only_one
node1     raid.scrub.enable                     on           off
none
2 entries were displayed.
```

storage shelf commands

storage shelf show

Display a list of storage shelves

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

Description

The `storage shelf show` command displays information about all the storage shelves in the storage system. If no parameters are specified, the default command displays the following information about the storage shelves:

- Shelf Name
- Shelf ID
- Serial Number
- Model
- Module Type
- Status

To display detailed profile information about a single storage shelf, use the `-shelf` parameter.

Parameters

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

Displays the specified fields for all the storage shelves, in column style output.

| [-bay]

Displays the following details about the disk bays in the storage shelf:

- The unique positional identifier of the disk bay
- Whether a disk drive is installed in the bay

- Bay type
- Operational status of the disk bay

| [-connectivity]

Displays the following details about the connectivity from the node to the storage shelf:

- Node name
- Initiator side switch port
- Target side switch port
- World-wide port name
- Target Port Group Number (TPGN)

| [-cooling]

Displays the following details about the cooling elements and temperature sensors of the storage shelf:

- Element ID of the cooling fan
- The current speed of the cooling fan in revolutions per minute (rpm)
- Operational status of the cooling fan
- Sensor ID of the temperature sensor element
- Temperature at the sensor in degrees Celsius
- Whether the current temperature at the sensor is the ambient temperature
- Low critical threshold value for the temperature sensor
- Low warning threshold value for the temperature sensor
- High critical threshold value for the temperature sensor
- High warning threshold value for the temperature sensor
- Operational status for the temperature sensor

| [-errors]

Displays the following error status information about the storage shelves that have errors:

- Error type
- Error description

| [-module]

Displays the following details about the I/O modules attached to the storage shelf:

- Module ID
- Module part number
- Serial number of the Enclosure Services Controller Electronics element
- Whether monitoring is enabled on this module
- Whether this module is the SAS expander master module
- Whether this module is the element reporting
- Version of the firmware installed on the module

- Latest firmware revision
- Number of times, since the last boot, that this module has been swapped
- Operational status of the module

| [-port]

Displays the following details about the storage shelf ports:

- Expander phy element identifier
- SAS shelf port type
- World-wide Port Name of the SAS port
- Operational physical link rate of the SAS port in Gb/s
- Negotiated physical link rate of the SAS port in Gb/s
- Power status of the SAS port
- Status of the SAS port
- Fibre Channel shelf port ID
- Fibre Channel shelf port type
- Fibre Channel shelf port status

| [-power]

Displays the following details about the power supplies, voltage sensors, and current sensors of the storage shelf:

- Power Supply Unit (PSU) number
- PSU type
- PSU part number
- PSU serial number
- PSU power rating in watts
- PSU crest factor
- Power drawn from the PSU in watts
- Whether the PSU can be reset via software control
- Whether the auto power reset of the PSU is enabled
- PSU firmware revision
- Operational status of the PSU
- Voltage sensor number
- Voltage detected by the voltage sensor, in volts (V)
- Operational status of the voltage sensor
- Current sensor number
- Current detected by the current sensor, in milliamps (mA)
- Operational status of the current sensor

| [-instance]

Displays expanded information about all the storage shelves in the system.

[-shelf <text>] - Shelf Name

Displays information only about the storage shelves that match the names you specify.

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

Displays information only about the storage shelves that are attached to the node you specify.

[-shelf-uid <text>] - Shelf UID

Displays information only about the storage shelf that matches the shelf UID you specify. Example:

`50:05:0c:c0:02:10:64:26`

[-stack-id {<integer>|-}] - Stack ID

Displays information only about the storage shelves that are attached to the stack that matches the stack ID you specify

[-shelf-id <text>] - Shelf ID

Displays information only about the storage shelves that match the shelf ID you specify.

[-module-type

**{unknown|atfcx|esh4|iom3|iom6|iom6e|iom12|iom12e|iom12f|nsm100|nsm8e|psm3e|iom12b
}] - Shelf Module Type**

Displays information only about the storage shelves that match the module-type you specify.

[-connection-type {unknown|fc|sas|nvme}] - Shelf Connection Type

Displays information only about the storage shelves that match the connection type you specify. Example:
FC or SAS.

[-is-local-attach {true|false}] - Is the Shelf Local to This Cluster?

Displays information only about the storage shelves that are local (TRUE) or remote (FALSE) to this cluster.

[-vendor <text>] - Shelf Vendor

Displays information only about the storage shelves that match the vendor you specify.

[-product-id <text>] - Shelf Product Identification

Displays information only about the storage shelves that match the product ID you specify.

[-serial-number <text>] - Shelf Serial Number

Displays information only about the storage shelf that matches the serial number you specify.

[-disk-count {<integer>|-}] - Disk Count

Displays information only about the storage shelves that have the disk count you specify.

[-state {unknown|no-status|init-required|online|offline|missing}] - Shelf State

Displays information only about the storage shelves that are in the state you specify.

[-op-status {unknown|normal|warning|error|critical|standby-power}] - Shelf Operational Status

Displays information only about the storage shelves that are currently operating under the status condition

you specify.

[-bay-id {<integer>|-}] - Bay ID

Displays information only about the storage shelves that have bays that match the bay ID you specify.

[-bay-type {unknown|single-disk|multi-lun}] - Bay Type

Displays information only about the storage shelves that have bays that match the type of bay you specify.

[-bay-has-disk {true|false}] - Bay Has Disk

Displays information only about the storage shelves that have bays with disk drives inserted in them (true) or empty bays (false).

[-bay-op-status {unknown|normal|error}] - Bay Operational Status

Displays information only about the storage shelves that have bays that match the operational state you specify.

[-controller {<nodename>|local}] - Controller Name

Displays information only about the storage shelves that are connected to the node you specify.

[-controller-uuid <text>,...] - Controller UUID

Displays information only about the storage shelves that are connected to the node UUID you specify.

[-initiator <text>,...] - Initiator

Displays information only about the storage shelves that are visible to the initiator you specify.

[-initiator-wwpn <text>,...] - Initiator WWPN

Displays information only about the storage shelves that are visible to the initiator WWPN you specify.

[-initiator-side-switch-port <text>,...] - Initiator Side Switch Port

Displays information only about the storage shelves that are visible to an initiator connected to the switch port you specify.

[-target-side-switch-port <text>,...] - Target Side Switch Port

Displays information only about the storage shelves visible on target ports identified by the switch port to which they are connected.

[-target-port <text>,...] - Target Port

Displays information only about the storage shelves visible on the specified target ports identified by their World-Wide Port Name (WWPN).

[-tpgn {<integer>|-}] - Target Port Group Number

Displays information only about the storage shelves that belong to the Target Port Group Name (TPGN) you specify.

[-port-speed {<integer>|-}] - Port Speed

Displays information only about the storage shelves with ports that match the port speed you specify.

[-io-kbps {<integer>|-}] - Kbytes/sec on Storage Shelf

Displays information only about the storage shelves visible to an initiator that has executed I/O at the throughput you specify.

[-iops {<integer>|-}] - Number IOPS per Second on Storage Shelf

Displays information only about the storage shelves visible to an initiator that has executed the number of IOPs you specify.

[-current-sensor-id {<integer>|-}] - Current Sensor ID

Displays information only about the storage shelves with current sensor that matches the current sensor ID you specify.

[-current-sensor-location <text>,...] - Current Sensor Location

Displays information only about the storage shelves with current sensors installed at the location you specify.

[-current-sensor-reading {<integer>|-}] - Current Sensor Reading

Displays information only about the storage shelves with current sensors that match the current reading you specify.

[-current-op-status {unknown|normal|over-current-critical|under-current-critical|not-supported|not-installed}] - Operational Status

Displays information only about the storage shelves with current sensors that match the operational status you specify.

[-fan-id {<integer>|-}] - Fan ID

Displays information only about the storage shelves with cooling fans that match the fan IDs you specify.

[-fan-location <text>,...] - Fan Location

Displays information only about the storage shelves with cooling fans installed.

[-fan-rpm {<integer>|-}] - Fan Rotation Per Minute

Displays information only about the storage shelves with cooling fans that match the rpm rate you specify.

[-fan-op-status {unknown|normal|off|error|not-supported|not-installed}] - Fan Operational Status

Displays information only about the storage shelves with cooling fans that match the operational status you specify.

[-module-id <text>,...] - Module ID

Displays information only about the storage shelves with an I/O module that matches the module ID you specify.

[-module-location <text>,...] - Module Location

Displays information only about the storage shelves with I/O modules in the specified shelf module slots.

[-module-part-number <text>,...] - Module Part Number

Displays information only about the storage shelves with I/O modules that match the module part numbers you specify.

[-is-sas-master-module {true|false}] - Is SAS Expander Master Module?

Displays information only about the storage shelves with a SAS master I/O module (true) or an I/O module that is not a SAS master (false). This parameter applies only to SAS shelves.

[-is-monitor-active {true|false}] - Is Monitor Active?

Displays information only about the storage shelves whose monitoring is enabled (true) or disabled (false).

[-enclosure-type <text>,...] - Module Enclosure Type

Displays information only about the storage shelves that match the enclosure types you specify.

[-es-serial-number <text>,...] - ES Electronics Element Serial Number

Displays information only about the storage shelves with I/O modules that match the electronics serial numbers you specify.

[-module-fru-id <text>,...] - Field Replaceable Unit ID

Displays information only about the storage shelves with I/O modules that match the field replaceable unit (FRU) IDs you specify.

[-module-is-reporting-element {true|false}] - Is Reporting Element?

Displays information only about the storage shelves with element reporting I/O modules (true) or not (false).

[-module-fw-revision <text>,...] - Firmware Revision

Displays information only about the storage shelves with I/O modules that match the firmware revision you specify.

[-module-latest-fw-revision <text>,...] - Latest Firmware Revision

Displays information only about the storage shelves with I/O modules that match the latest firmware revision you specify.

[-module-fw-progress {not-available|ready|in-progress|failed}] - Module Firmware Progress

Displays information only about the storage shelves with I/O modules that match the specified firmware update progress.

[-module-swap-count {<integer>|-}] - Module Swap Count

Displays information only about the storage shelves whose I/O modules have been swapped the specified number of times.

[-module-op-status {unknown|normal|warning|error}] - Module Operational Status

Displays information only about the storage shelves with I/O modules that match the operational status you specify.

[-sas-port-id <text>,...] - Port ID

Displays information only about the storage shelves with SAS Ports that match the port IDs you specify.

[-sas-port-type {unknown|circle|square|sil|disk|in|out|unused|host|dcm|aux1|aux2|hi_ho|a_to_b|b_to_a}] - Port Type

Displays information only about the storage shelves with SAS Ports that match the SAS port type you specify.

[-sas-port-wwpn <text>,...] - Port World Wide Port Name

Displays information only about the storage shelves with SAS Ports that match the World-Wide Port Names you specify.

[-sas-port-speed <text>,...] - Port Speed

Displays information only about the storage shelves with SAS Ports that match the port speed you specify.

[-sas-negotiated-port-speed <text>,...] - Negotiated Port Speed

Displays information only about the storage shelves with SAS Ports that match the negotiated port speed you specify.

[-sas-port-power-status <text>,...] - Port Power Status

Displays information only about the storage shelves with SAS Ports that match the power status you specify.

[-sas-port-op-status {error|normal|off|unknown|byp-bad-term|bad-zone-recovery|byp_clk_thr|byp_comma_los|byp_crc_brst_thr|byp_data_timeout|byp_drv_fault|byp_drv_pcycle|byp_drv_pwr|byp_drv_self|byp_gen|byp_init|byp_lip_brst_thr|byp_lip_f8|byp_lip_rate_thr|byp_lipf7|byp_ltbi|byp_man|byp_no_drive|byp_osc|byp_other_thr|byp_rec_los|byp_rport|byp_stall_thr|byp_wrd_brst_thr|byp_wrd_rate_thr|byp_xmit_fault|diag_transmit|inserted|loopback|status_unknown|warn_high_clk_delta|warn_high_crc_rate|warn_high_lip|warn_high_wrd_rate|term|phy_dis_clk_fault|phy_dis_crc_err|phy_dis_crc_err_burst|phy_dis_disparity|phy_dis_disparity_burst|phy_dis_emulate_reserve|phy_dis_inval_dword|phy_dis_inval_dword_burst|phy_dis_loss_dword|phy_dis_loss_dword_burst|phy_dis_man_smp|phy_dis_manual|phy_dis_mirrored|empty|phy_dis_phy_change|phy_dis_phy_change_burst|phy_dis_phy_reset|phy_dis_phy_reset_burst|phy_dis_phy_unused|phy_ena|phy_ena_not_attach|phy_ena_unknown|phy_unknown|phy_dis_illegal}}] - Port Operational Status

Displays information only about the storage shelves with SAS Ports that match the operational status you specify.

[-sas-port-module-id {A|B}] - Port Module ID

Displays information only about the storage shelves with SAS Ports that match the module ID you specify.

[-fc-port-id <text>,...] - Fibre Channel Port ID

Displays information only about the storage shelves with FC Ports that match the port IDs you specify.

[-fc-port-mode

{unknown|circle|square|sil|disk|in|out|unused|host|dcm|aux1|aux2|hi_ho|a_to_b|b_to_a}}] - Fibre Channel Port Mode

Displays information only about the storage shelves with FC Ports that match the port modes you specify.

[-fc-port-op-status {error|normal|off|unknown|byp-bad-term|bad-zone-recovery|byp_clk_thr|byp_comma_los|byp_crc_brst_thr|byp_data_timeout|byp_drv_fault|byp_drv_pcycle|byp_drv_pwr|byp_drv_self|byp_gen|byp_init|byp_lip_brst_thr|byp_lip_f8|byp_lip_rate_thr|byp_lipf7|byp_ltbi|byp_man|byp_no_drive|byp_osc|byp_other_thr|byp_rec_los|byp_rport|byp_stall_thr|byp_wrd_brst_thr|byp_wrd_rate_thr|byp_xmit_fault|diag_transmit|inserted|loopback|status_unknown|warn_high_clk_delta|warn_high_crc_rate|warn_high_lip|warn_high_wrd_rate|term|phy_dis_clk_fault|phy_dis_crc_err|phy_dis_crc_err_burst|phy_dis_disparity|phy_dis_disparity_burst|phy_dis_emulate_reserve|phy_dis_inval_dword|phy_dis_inval_dword_burst|phy_dis_loss_dword|phy_dis_loss_dword_burst|phy_dis_man_smp|phy_dis_manual|phy_dis_mirrored|empty|phy_dis_phy_change|phy_dis_phy_change_burst|phy_dis_phy_reset|phy_dis_phy_reset_burst|phy_dis_phy_unused|phy_ena|phy_ena_not_attach|phy_ena_unknown|phy_unknown|phy_dis_illegal}}] - Fibre Channel Port Operational Status

Displays information only about the storage shelves with FC Ports that match the operational status you specify.

[-psu-id {<integer>|-}] - Power Supply Unit ID

Displays information only about the storage shelves with power supply units (PSU) that match the unit IDs you specify.

[-psu-location <text>,...] - Power Supply Unit Location

Displays information only about the storage shelves with PSUs that are located at the specified location inside the shelf.

[-psu-type <text>,...] - Power Supply Unit Type

Displays information only about the storage shelves with PSUs that match the PSU types you specify.

[-psu-part-number <text>,...] - Power Supply Unit Part Number

Displays information only about the storage shelves with PSUs that match the PSU part number you specify.

[-psu-serial-number <text>,...] - Power Supply Unit Serial Number

Displays information only about the storage shelves with PSUs that match the PSU serial numbers you specify.

[-psu-reset-capable {true|false}] - Power Supply Unit Reset Capability

Displays information only about the storage shelves with reset capable PSUs (true) or reset incapable PSUs (false).

[-psu-is-enabled {true|false}] - Power Supply Unit Enable/Disable Status

Displays information only about the storage shelves with PSUs that are enabled (true) or disabled (false).

[-psu-fw-version <text>,...] - Power Supply Unit Firmware Version

Displays information only about the storage shelves with PSUs that have the firmware version you specify.

[-psu-op-status {unknown|normal|error|dc-over-voltage|dc-under-voltage|dc-over-current|over-temperature-error|failed|off|not-supported|not-installed}] - Operational Status

Displays information only about the storage shelves with PSUs that match the operational status you specify.

[-psu-power-rating {<integer>|-}] - Power Supply Power Ratings In Watts

Displays information only about the storage shelves with PSUs that match the power rating you specify.

[-psu-crest-factor {<integer>|-}] - Power Supply Crest Factor

Displays information only about the storage shelves with PSUs that match the crest factor value you specify.

[-psu-power-drawn {<integer>|-}] - Power Drawn From PSU In Watts

Displays information only about the storage shelves with PSUs that match the drawn power you specify.

[-temp-sensor-id {<integer>|-}] - Sensor Name

Displays information only about the storage shelves with temperature sensors that match the sensor IDs

you specify.

[-temp-sensor-location <text>,...] - Sensor Location

Displays information only about the storage shelves with temperature sensors that match the specified sensor locations inside the shelf.

[-temp-sensor-reading {<integer>|-}] - Temperature Reading

Displays information only about the storage shelves with temperature sensors that match the temperature reading you specify.

[-temp-is-ambient {true|false}] - Temperature Reading at Ambient Value

Displays information only about the storage shelves with temperature sensors whose current temperature reading is ambient (true) or not (false).

[-temp-high-critical-threshold {<integer>|-}] - High Critical Threshold

Displays information only about the storage shelves with temperature sensors that match the high critical threshold you specify.

[-temp-high-warning-threshold {<integer>|-}] - High Warning Threshold

Displays information only about the storage shelves with temperature sensors that match the high warning threshold you specify.

[-temp-low-warning-threshold {<integer>|-}] - Low Warning Threshold

Displays information only about the storage shelves with temperature sensors that match the low warning threshold you specify.

[-temp-low-critical-threshold {<integer>|-}] - Low Critical Threshold

Displays information only about the storage shelves with temperature sensors that match the low critical threshold you specify.

[-temp-op-status {unknown|normal|under-temperature|over-temperature|error|not-supported|not-installed}] - Operational Status

Displays information only about the storage shelves with temperature sensors that match the operational status you specify.

[-voltage-sensor-id {<integer>|-}] - Voltage Sensor ID

Displays information only about the storage shelves with voltage sensors that match the sensor IDs you specify.

[-voltage-sensor-location <text>,...] - Voltage Sensor Location

Displays information only about the storage shelves with voltage sensors that match the specified sensor locations inside the shelf.

[-voltage-sensor-reading <text>,...] - Voltage Current Reading

Displays information only about the storage shelves with voltage sensors that match the voltage reading you specify.

[-voltage-op-status {unknown|normal|over-voltage-critical|under-voltage-critical|not-supported|not-installed|not-recoverable}] - Operational Status

Displays information only about the storage shelves with voltage sensors that match the operational status you specify.

[-nsm-port-module-id {A|B}] - Port Module ID

Displays information only about the storage shelves with PCIe Ports from the specified module.

[-nsm-port-id <integer>,...] - Port ID

Displays information only about the storage shelves with PCIe Ports that match the specified ID.

[-nsm-port-type {cpu|disk|cx5|ethernet}] - Port Type

Displays information only about the storage shelves with PCIe Ports that match the specified type.

[-nsm-port-state {ok|off-link-disabled|off-dll-link|link-down|no-drive}] - Port State

Displays information only about the storage shelves with PCIe Ports that match the specified state.

[-nsm-port-bay <integer>,...] - Port Bay

Displays information only about the storage shelves with PCIe Ports that match the specified bay.

[-nsm-port-disk-id <integer>,...] - Port Disk ID

Displays information only about the storage shelves with PCIe Ports that match the specified disk ID.

[-nsm-port-is-installed {true|false}] - Port Is Disk Installed

Displays information only about the storage shelves with PCIe Ports that have a disk installed.

[-nsm-port-is-error {true|false}] - Port Has Error

Displays information only about the storage shelves with PCIe Ports that have errors.

[-nsm-port-speed {2.5|5.0|8.0}] - Port Speed

Displays information only about the storage shelves with PCIe Ports that match the specified speed.

[-nsm-port-speed-max {2.5|5.0|8.0}] - Max Port Speed

Displays information only about the storage shelves with PCIe Ports that match the specified maximum speed.

[-nsm-port-lane-width <integer>,...] - Port Lane Width

Displays information only about the storage shelves with PCIe Ports that match the specified lane width.

[-nsm-port-lane-width-max <integer>,...] - Max Port Lane Width

Displays information only about the storage shelves with PCIe Ports that match the specified maximum lane width.

[-dimm-module-id {A|B}] - DIMM Module ID

Displays information only about the storage shelves with DIMMs from the specified module.

[-dimm-id <integer>,...] - DIMM ID

Displays information only about the storage shelves with DIMMs that match the specified ID.

[-dimm-serial-number <text>,...] - DIMM Serial Number

Displays information only about the storage shelves with DIMMs that match the specified serial number.

[-dimm-part-number <text>,...] - DIMM Part Number

Displays information only about the storage shelves with DIMMs that match the specified part number.

[-dimm-vendor <text>,...] - DIMM Vendor

Displays information only about the storage shelves with DIMMs that match the specified vendor.

[-dimm-type <text>,...] - DIMM Type

Displays information only about the storage shelves with DIMMs that match the specified type.

[-dimm-size <text>,...] - DIMM Size

Displays information only about the storage shelves with DIMMs that match the specified size.

[-dimm-speed <text>,...] - DIMM Speed

Displays information only about the storage shelves with DIMMs that match the specified speed.

[-dimm-location <text>,...] - DIMM Location

Displays information only about the storage shelves with DIMMs that match the specified location.

[-dimm-op-status {unknown|normal|error|not-supported|not-installed}] - DIMM Operational Status

Displays information only about the storage shelves with DIMMs that match the specified operational status.

[-boot-device-module-id {A|B}] - Boot Device Module ID

Displays information only about the storage shelves with boot devices from the specified module.

[-boot-device-id <integer>,...] - Boot Device ID

Displays information only about the storage shelves with boot devices that match the specified ID.

[-boot-device-serial-number <text>,...] - Boot Device Serial Number

Displays information only about the storage shelves with boot devices that match the specified serial number.

[-boot-device-part-number <text>,...] - Boot Device Part Number

Displays information only about the storage shelves with boot devices that match the specified part number.

[-boot-device-vendor <text>,...] - Boot Device Vendor

Displays information only about the storage shelves with boot devices that match the specified vendor.

[-boot-device-type <text>,...] - Boot Device Type

Displays information only about the storage shelves with boot devices that match the specified type.

[-boot-device-size <text>,...] - Boot Device Size

Displays information only about the storage shelves with boot devices that match the specified size.

[-boot-device-op-status {unknown|normal|error|not-supported|not-installed}] - Boot Device Operational Status

Displays information only about the storage shelves with boot devices that match the specified operational status.

[-coin-battery-module-id {A|B}] - Coin Battery Module ID

Displays information only about the storage shelves with coin batteries from the specified module.

[-coin-battery-id <integer>,...] - Coin Battery ID

Displays information only about the storage shelves with coin batteries that match the specified ID.

[-coin-battery-voltage <integer>,...] - Coin Battery Voltage (mV)

Displays information only about the storage shelves with coin batteries that match the specified voltage.

[-coin-battery-op-status {unknown|normal|error|low|high|not-supported|not-installed}] - Coin Battery Operational Status

Displays information only about the storage shelves with coin batteries that match the specified operational status.

[-error-type {Unknown|ACPP|Bootdevice|Coinbattery|Configuration|Current|Dimm|Expander|Fan|Module|PCM|Power|Temperature|Voltage}] - Error Type

Displays information only about the storage shelves with errors that match the error type you specify.

[-error-severity {unknown|notice|warning|error|critical}] - Error Severity

Displays information only about the storage shelves with errors that match the error severity you specify.

Examples

The following example displays information about all storage shelves:

```
cluster1::> storage shelf show
Module Operational
      Shelf Name  Shelf ID  Serial Number    Model    Type
Status
-----
Critical          1.1         1  6000832415     DS2246   IOM6
Normal            1.2         2  6000647652     DS2246   IOM6
Normal            1.3         3  6000003844     DS2246   IOM6
Normal            1.4         4  SHJ000000013A9E DS4246   IOM6
Normal            1.5         5  SHJ000000013A84 DS4246   IOM6
Normal            1.6         6  6000005555     DS2246   IOM6
Normal
      6 entries were displayed.
cluster1::>
```

The following example displays expanded information about a storage shelf named 1.2:

```
cluster1::> storage shelf show -shelf 1.2 -instance
```

Shelf Name: 1.2

Stack ID: 1

Shelf ID: 2

Shelf UID: 50:0a:09:80:01:b9:75:41

Serial Number: 6000647652

Module Type: IOM6

Model: DS2246

Shelf Vendor: NETAPP

Disk Count: 12

Connection Type: SAS

Shelf State: Online

Status: Normal

Modules:

Module is

Monitor Is Reporting FW Update

Latest Swap Operational Module

ID Part No. ES Serial No. is Active Master Element Progress

FW Rev. FW Rev. Count Status Location

```
-----
a 111-00190+A0 8006437891 true false false not-
available 0191 - 0 normal rear of the shelf at the top
left
b 111-00190+A0 8006435180 true true true not-
available 0191 - 0 normal rear of the shelf at the top
right
```

Paths:

Speed

| Controller | Initiator | Initiator | Side | Switch | Port | Target | Side |
|-------------|-------------|-----------|------|--------|------|--------|------|
| Switch Port | Target Port | TPGN | Gb/s | I/O | KB/s | IOPS | |

```
-----
stsw-8020-01 0a - - - -
- - - - -
stsw-8020-01 2b - - - -
- - - - -
stsw-8020-02 0a - - - -
- - - - -
stsw-8020-02 2b - - - -
- - - - -
```

Power Supply Units:

Crest Power

Reset PSU Operational

ID Type Part# Serial# Power Rating Factor Drawn

Capable Enabled Firmware Status PSU Location

```

-----
-----
1 9C 114-00065+A1 XXT131052637 - - -
false true 020F normal rear of the shelf at the bottom
left
2 9C 114-00065+A1 XXT131052551 - - -
false true 020F normal rear of the shelf at the bottom
right
Voltage Sensors:
Voltage Operational
ID (V) Status Sensor Location
-----
1 5.70 normal rear of the shelf on the lower left
power supply
2 12.300 normal rear of the shelf on the lower left
power supply
3 5.70 normal rear of the shelf on the lower
right power supply
4 12.180 normal rear of the shelf on the lower
right power supply
Current Sensors:
Current Operational
ID (mA) Status Sensor Location
-----
1 0 normal rear of the shelf on the lower left
power supply
2 0 normal rear of the shelf on the lower left
power supply
3 0 normal rear of the shelf on the lower
right power supply
4 0 normal rear of the shelf on the lower
right power supply
Fans:
Speed Operational
ID (RPM) Status Fan Location
-----
1 3000 normal rear of the shelf on the lower left
power supply
2 2970 normal rear of the shelf on the lower left
power supply
3 3000 normal rear of the shelf on the lower
right power supply
4 2970 normal rear of the shelf on the lower
right power supply
Temperature:
-- Thresholds °C --

```

| Location | ID | Temp °C | Is Ambient | Low Crit | Low Warn | High Crit | High Warn | Operational Status | Sensor |
|---|----|---------|------------|----------|----------|-----------|-----------|--------------------|---------------|
| shelf on the left, on the OPS panel | 1 | 23 | true | 0 | 5 | 42 | 40 | normal | front of the |
| shelf on the midplane | 2 | 26 | false | 5 | 10 | 55 | 50 | normal | inside of the |
| shelf on the lower left power supply | 3 | 24 | false | 5 | 10 | 55 | 50 | normal | rear of the |
| shelf on the lower left power supply | 4 | 39 | false | 5 | 10 | 70 | 65 | normal | rear of the |
| shelf on the lower right power supply | 5 | 25 | false | 5 | 10 | 55 | 50 | normal | rear of the |
| shelf on the lower right power supply | 6 | 36 | false | 5 | 10 | 70 | 65 | normal | rear of the |
| shelf at the top left, on shelf module A | 7 | 25 | false | 5 | 10 | 60 | 55 | normal | rear of the |
| shelf at the top right, on shelf module B | 8 | 26 | false | 5 | 10 | 60 | 55 | normal | rear of the |

SAS Ports:

| Port | Phy # | IOM Port | Type | WWPN | Operational | Negotiated | Status |
|---------|-------|----------|--------|------------------|-------------|------------|--------|
| Enabled | 0 | A | Square | 500a098004b063b0 | 6.0 | - | - |
| Enabled | 1 | A | Square | 500a098004b063b0 | 6.0 | - | - |
| Enabled | 2 | A | Square | 500a098004b063b0 | 6.0 | - | - |
| Enabled | 3 | A | Square | 500a098004b063b0 | 6.0 | - | - |
| Enabled | 4 | A | Circle | 500a09800569f03f | 6.0 | - | - |
| Enabled | 5 | A | Circle | 500a09800569f03f | 6.0 | - | - |
| Enabled | 6 | A | Circle | 500a09800569f03f | 6.0 | - | - |
| Enabled | 7 | A | Circle | 500a09800569f03f | 6.0 | - | - |
| Enabled | 8 | A | Disk | 500605ba00c1cb8d | 6.0 | 6.0 | on |

| | | | | | | | |
|---------|----|---|------|------------------|-----|-----|----|
| Enabled | 9 | A | Disk | 500605ba00c1ea8d | 6.0 | 6.0 | on |
| Enabled | 10 | A | Disk | 500605ba00c1d111 | 6.0 | 6.0 | on |
| Enabled | 11 | A | Disk | 500605ba00c1bc49 | 6.0 | 6.0 | on |
| Enabled | 12 | A | Disk | 500605ba00c1cdfd | 6.0 | 6.0 | on |
| Enabled | 13 | A | Disk | 500605ba00c1c531 | 6.0 | 6.0 | on |
| Enabled | 14 | A | Disk | 500605ba00c1eb05 | 6.0 | 6.0 | on |
| Enabled | 15 | A | Disk | 500605ba00c1ec29 | 6.0 | 6.0 | on |
| Enabled | 16 | A | Disk | 500605ba00c1bc29 | 6.0 | 6.0 | on |
| Enabled | 17 | A | Disk | 500605ba00c1c471 | 6.0 | 6.0 | on |
| Enabled | 18 | A | Disk | 500605ba00c039a9 | 6.0 | 6.0 | on |
| Enabled | 19 | A | Disk | 500605ba00c1c4dd | 6.0 | 6.0 | on |
| Empty | 20 | A | Disk | - | - | - | - |
| Empty | 21 | A | Disk | - | - | - | - |
| Empty | 22 | A | Disk | - | - | - | - |
| Empty | 23 | A | Disk | - | - | - | - |
| Empty | 24 | A | Disk | - | - | - | - |
| Empty | 25 | A | Disk | - | - | - | - |
| Empty | 26 | A | Disk | - | - | - | - |
| Empty | 27 | A | Disk | - | - | - | - |
| Empty | 28 | A | Disk | - | - | - | - |
| Empty | 29 | A | Disk | - | - | - | - |
| Empty | 30 | A | Disk | - | - | - | - |
| Empty | 31 | A | Disk | - | - | - | - |

| | | | | | |
|----------|---|--------|------------------|-----|--------|
| 32 | A | SIL | - | - | - - |
| Disabled | | | | | |
| 33 | A | SIL | - | - | - - |
| Disabled | | | | | |
| 34 | A | SIL | - | - | - - |
| Disabled | | | | | |
| 35 | A | SIL | - | - | - - |
| Disabled | | | | | |
| 0 | B | Square | 500a098004af9e30 | 6.0 | - - |
| Enabled | | | | | |
| 1 | B | Square | 500a098004af9e30 | 6.0 | - - |
| Enabled | | | | | |
| 2 | B | Square | 500a098004af9e30 | 6.0 | - - |
| Enabled | | | | | |
| 3 | B | Square | 500a098004af9e30 | 6.0 | - - |
| Enabled | | | | | |
| 4 | B | Circle | 500a098005688dbf | 6.0 | - - |
| Enabled | | | | | |
| 5 | B | Circle | 500a098005688dbf | 6.0 | - - |
| Enabled | | | | | |
| 6 | B | Circle | 500a098005688dbf | 6.0 | - - |
| Enabled | | | | | |
| 7 | B | Circle | 500a098005688dbf | 6.0 | - - |
| Enabled | | | | | |
| 8 | B | Disk | 500605ba00c1cb8e | 6.0 | 6.0 on |
| Enabled | | | | | |
| 9 | B | Disk | 500605ba00c1ea8e | 6.0 | 6.0 on |
| Enabled | | | | | |
| 10 | B | Disk | 500605ba00c1d112 | 6.0 | 6.0 on |
| Enabled | | | | | |
| 11 | B | Disk | 500605ba00c1bc4a | 6.0 | 6.0 on |
| Enabled | | | | | |
| 12 | B | Disk | 500605ba00c1cdfc | 6.0 | 6.0 on |
| Enabled | | | | | |
| 13 | B | Disk | 500605ba00c1c532 | 6.0 | 6.0 on |
| Enabled | | | | | |
| 14 | B | Disk | 500605ba00c1eb06 | 6.0 | 6.0 on |
| Enabled | | | | | |
| 15 | B | Disk | 500605ba00c1ec2a | 6.0 | 6.0 on |
| Enabled | | | | | |
| 16 | B | Disk | 500605ba00c1bc2a | 6.0 | 6.0 on |
| Enabled | | | | | |
| 17 | B | Disk | 500605ba00c1c472 | 6.0 | 6.0 on |
| Enabled | | | | | |
| 18 | B | Disk | 500605ba00c039aa | 6.0 | 6.0 on |
| Enabled | | | | | |

| | | | | | | | |
|----------|----|---|------|------------------|-----|-----|----|
| Enabled | 19 | B | Disk | 500605ba00c1c4de | 6.0 | 6.0 | on |
| Empty | 20 | B | Disk | - | - | - | - |
| Empty | 21 | B | Disk | - | - | - | - |
| Empty | 22 | B | Disk | - | - | - | - |
| Empty | 23 | B | Disk | - | - | - | - |
| Empty | 24 | B | Disk | - | - | - | - |
| Empty | 25 | B | Disk | - | - | - | - |
| Empty | 26 | B | Disk | - | - | - | - |
| Empty | 27 | B | Disk | - | - | - | - |
| Empty | 28 | B | Disk | - | - | - | - |
| Empty | 29 | B | Disk | - | - | - | - |
| Empty | 30 | B | Disk | - | - | - | - |
| Empty | 31 | B | Disk | - | - | - | - |
| Disabled | 32 | B | SIL | - | - | - | - |
| Disabled | 33 | B | SIL | - | - | - | - |
| Disabled | 34 | B | SIL | - | - | - | - |
| Disabled | 35 | B | SIL | - | - | - | - |

FC Ports:

| ID | Port | Type | Status |
|-----|-------|-------|--------|
| --- | ----- | ----- | ----- |
| - | - | - | - |

Bays:

| Has | ID | Disk | Bay | Type | Operational Status |
|-----|-----|------|-----|-------------|--------------------|
| --- | --- | --- | --- | --- | ----- |
| | 0 | true | | single-disk | normal |
| | 1 | true | | single-disk | normal |
| | 2 | true | | single-disk | normal |

```

3 true  single-disk normal
4 true  single-disk normal
5 true  single-disk normal
6 true  single-disk normal
7 true  single-disk normal
8 true  single-disk normal
9 true  single-disk normal
10 true single-disk normal
11 true single-disk normal
12 false single-disk normal
13 false single-disk normal
14 false single-disk normal
15 false single-disk normal
16 false single-disk normal
17 false single-disk normal
18 false single-disk normal
19 false single-disk normal
20 false single-disk normal
21 false single-disk normal
22 false single-disk normal
23 false single-disk normal
cluster1::>

```

The following example displays information about the power supplies, voltage sensors and current sensors of the storage shelf 1.1:

```

cluster1::> storage shelf show -shelf 1.1 -power
Shelf Name: 1.1
      Stack ID: 1
      Shelf ID: 1
      Shelf UID: 50:0a:09:80:01:cb:d6:84
      Serial Number: 6000832415
      Module Type: IOM6
      Model: DS2246
      Shelf Vendor: NETAPP
      Disk Count: 12
      Connection Type: SAS
      Shelf State: Online
      Status: Normal
Power Supply Units:

```

| Reset | PSU | Operational | Crest | Power |
|---------|---------|-------------|---------|--------------|
| ID | Type | Part# | Serial# | Power Rating |
| Capable | Enabled | Firmware | Status | Factor |
| Drawn | | | | |
| --- | ---- | ----- | ----- | ----- |

```

-----
      1 9C 114-00065+A1 XXT132835072 - - -
false true 020F normal
      2 9C 114-00065+A1 XXT132835073 - - -
false true 020F normal
Voltage Sensors:
      Voltage Operational
      ID      (V) Status
-----
      1      5.70 normal
      2     12.180 normal
      3      5.70 normal
      4     12.300 normal
Current Sensors:
      Current Operational
      ID      (mA) Status
-----
      1         0 normal
      2         0 normal
      3      3900 normal
      4         0 normal
Errors:
-----
      Critical condition is detected in storage shelf power supply unit "1".
The unit might fail.
      Critical over temperature failure for temperature sensor "1". Current
temperature: "75" C ("167" F).
cluster1::>

```

The following example displays information about the cooling elements and temperature sensors inside the storage shelf 1.2:

```

cluster1::> storage shelf show -shelf 1.2 -cooling
Shelf Name: 1.2
      Stack ID: 1
      Shelf ID: 2
      Shelf UID: 50:0a:09:80:01:b9:75:41
      Serial Number: 6000647652
      Module Type: IOM6
      Model: DS2246
      Shelf Vendor: NETAPP
      Disk Count: 12
      Connection Type: SAS
      Shelf State: Online
      Status: Normal

Fans:
      Speed Operational
      ID (RPM) Status
      -- -----
      1  3000 normal
      2  3000 normal
      3  3000 normal
      4  2970 normal

Temperature:
      -- Thresholds °C --
      Temp Is      Low  Low High High Operational
      ID   °C Ambient Crit Warn Crit Warn Status
      --- ---
      1   23 true    0   5  42  40 normal
      2   26 false   5  10  55  50 normal
      3   24 false   5  10  55  50 normal
      4   39 false   5  10  70  65 normal
      5   25 false   5  10  55  50 normal
      6   36 false   5  10  70  65 normal
      7   25 false   5  10  60  55 normal
      8   27 false   5  10  60  55 normal

Errors:
      -----
      -
cluster1::>

```

The following example displays information about the connectivity from the node to the storage shelf 1.2:


```
cluster1::> storage shelf show -shelf 1.2 -bay
Shelf Name: 1.2
          Stack ID: 1
          Shelf ID: 2
          Shelf UID: 50:0a:09:80:01:b9:75:41
    Serial Number: 6000647652
          Module Type: IOM6
          Model: DS2246
          Shelf Vendor: NETAPP
          Disk Count: 12
    Connection Type: SAS
          Shelf State: Online
          Status: Normal
```

Bays:

| Has | ID | Disk | Bay | Type | Status |
|-----|----|-------|-------------|--------|--------|
| | 0 | true | single-disk | normal | |
| | 1 | true | single-disk | normal | |
| | 2 | true | single-disk | normal | |
| | 3 | true | single-disk | normal | |
| | 4 | true | single-disk | normal | |
| | 5 | true | single-disk | normal | |
| | 6 | true | single-disk | normal | |
| | 7 | true | single-disk | normal | |
| | 8 | true | single-disk | normal | |
| | 9 | true | single-disk | normal | |
| | 10 | true | single-disk | normal | |
| | 11 | true | single-disk | normal | |
| | 12 | false | single-disk | normal | |
| | 13 | false | single-disk | normal | |
| | 14 | false | single-disk | normal | |
| | 15 | false | single-disk | normal | |
| | 16 | false | single-disk | normal | |
| | 17 | false | single-disk | normal | |
| | 18 | false | single-disk | normal | |
| | 19 | false | single-disk | normal | |
| | 20 | false | single-disk | normal | |
| | 21 | false | single-disk | normal | |
| | 22 | false | single-disk | normal | |
| | 23 | false | single-disk | normal | |

Errors:

```
-----
-
```

```
cluster1::>
```


The following example displays information about the ports of the storage shelf 1.2:

```
cluster1::> storage shelf show -shelf 1.2 -port
Shelf Name: 1.2
      Stack ID: 1
      Shelf ID: 2
      Shelf UID: 50:0a:09:80:01:b9:75:41
      Serial Number: 6000647652
      Module Type: IOM6
      Model: DS2246
      Shelf Vendor: NETAPP
      Disk Count: 12
      Connection Type: SAS
      Shelf State: Online
      Status: Normal

SAS Ports:
-- Port Speeds Gb/s -- Power
Port
Phy # IOM Port Type WWPN      Operational Negotiated Status
Status
-----
-----
      0  A  Square  500a098004b063b0      6.0      - -
Enabled
      1  A  Square  500a098004b063b0      6.0      - -
Enabled
      2  A  Square  500a098004b063b0      6.0      - -
Enabled
      3  A  Square  500a098004b063b0      6.0      - -
Enabled
      4  A  Circle  500a09800569f03f      6.0      - -
Enabled
      5  A  Circle  500a09800569f03f      6.0      - -
Enabled
      6  A  Circle  500a09800569f03f      6.0      - -
Enabled
      7  A  Circle  500a09800569f03f      6.0      - -
Enabled
      8  A  Disk    500605ba00c1cb8d      6.0      6.0 on
Enabled
      9  A  Disk    500605ba00c1ea8d      6.0      6.0 on
Enabled
     10  A  Disk    500605ba00c1d111      6.0      6.0 on
Enabled
     11  A  Disk    500605ba00c1bc49      6.0      6.0 on
Enabled
```

| | | | | | | | |
|----------|----|---|------|------------------|-----|-----|----|
| Enabled | 12 | A | Disk | 500605ba00c1cdfd | 6.0 | 6.0 | on |
| Enabled | 13 | A | Disk | 500605ba00c1c531 | 6.0 | 6.0 | on |
| Enabled | 14 | A | Disk | 500605ba00c1eb05 | 6.0 | 6.0 | on |
| Enabled | 15 | A | Disk | 500605ba00c1ec29 | 6.0 | 6.0 | on |
| Enabled | 16 | A | Disk | 500605ba00c1bc29 | 6.0 | 6.0 | on |
| Enabled | 17 | A | Disk | 500605ba00c1c471 | 6.0 | 6.0 | on |
| Enabled | 18 | A | Disk | 500605ba00c039a9 | 6.0 | 6.0 | on |
| Enabled | 19 | A | Disk | 500605ba00c1c4dd | 6.0 | 6.0 | on |
| Empty | 20 | A | Disk | - | - | - | - |
| Empty | 21 | A | Disk | - | - | - | - |
| Empty | 22 | A | Disk | - | - | - | - |
| Empty | 23 | A | Disk | - | - | - | - |
| Empty | 24 | A | Disk | - | - | - | - |
| Empty | 25 | A | Disk | - | - | - | - |
| Empty | 26 | A | Disk | - | - | - | - |
| Empty | 27 | A | Disk | - | - | - | - |
| Empty | 28 | A | Disk | - | - | - | - |
| Empty | 29 | A | Disk | - | - | - | - |
| Empty | 30 | A | Disk | - | - | - | - |
| Empty | 31 | A | Disk | - | - | - | - |
| Disabled | 32 | A | SIL | - | - | - | - |
| Disabled | 33 | A | SIL | - | - | - | - |
| Disabled | 34 | A | SIL | - | - | - | - |

| | | | | | |
|----------|---|--------|------------------|-----|--------|
| 35 | A | SIL | - | - | - - |
| Disabled | | | | | |
| 0 | B | Square | 500a098004af9e30 | 6.0 | - - |
| Enabled | | | | | |
| 1 | B | Square | 500a098004af9e30 | 6.0 | - - |
| Enabled | | | | | |
| 2 | B | Square | 500a098004af9e30 | 6.0 | - - |
| Enabled | | | | | |
| 3 | B | Square | 500a098004af9e30 | 6.0 | - - |
| Enabled | | | | | |
| 4 | B | Circle | 500a098005688dbf | 6.0 | - - |
| Enabled | | | | | |
| 5 | B | Circle | 500a098005688dbf | 6.0 | - - |
| Enabled | | | | | |
| 6 | B | Circle | 500a098005688dbf | 6.0 | - - |
| Enabled | | | | | |
| 7 | B | Circle | 500a098005688dbf | 6.0 | - - |
| Enabled | | | | | |
| 8 | B | Disk | 500605ba00c1cb8e | 6.0 | 6.0 on |
| Enabled | | | | | |
| 9 | B | Disk | 500605ba00c1ea8e | 6.0 | 6.0 on |
| Enabled | | | | | |
| 10 | B | Disk | 500605ba00c1d112 | 6.0 | 6.0 on |
| Enabled | | | | | |
| 11 | B | Disk | 500605ba00c1bc4a | 6.0 | 6.0 on |
| Enabled | | | | | |
| 12 | B | Disk | 500605ba00c1cdfc | 6.0 | 6.0 on |
| Enabled | | | | | |
| 13 | B | Disk | 500605ba00c1c532 | 6.0 | 6.0 on |
| Enabled | | | | | |
| 14 | B | Disk | 500605ba00c1eb06 | 6.0 | 6.0 on |
| Enabled | | | | | |
| 15 | B | Disk | 500605ba00c1ec2a | 6.0 | 6.0 on |
| Enabled | | | | | |
| 16 | B | Disk | 500605ba00c1bc2a | 6.0 | 6.0 on |
| Enabled | | | | | |
| 17 | B | Disk | 500605ba00c1c472 | 6.0 | 6.0 on |
| Enabled | | | | | |
| 18 | B | Disk | 500605ba00c039aa | 6.0 | 6.0 on |
| Enabled | | | | | |
| 19 | B | Disk | 500605ba00c1c4de | 6.0 | 6.0 on |
| Enabled | | | | | |
| 20 | B | Disk | - | - | - - |
| Empty | | | | | |
| 21 | B | Disk | - | - | - - |
| Empty | | | | | |

```

22  B  Disk  -  -  -
Empty
23  B  Disk  -  -  -
Empty
24  B  Disk  -  -  -
Empty
25  B  Disk  -  -  -
Empty
26  B  Disk  -  -  -
Empty
27  B  Disk  -  -  -
Empty
28  B  Disk  -  -  -
Empty
29  B  Disk  -  -  -
Empty
30  B  Disk  -  -  -
Empty
31  B  Disk  -  -  -
Empty
32  B  SIL   -  -  -
Disabled
33  B  SIL   -  -  -
Disabled
34  B  SIL   -  -  -
Disabled
35  B  SIL   -  -  -
Disabled
FC Ports:
      Port
      ID Port Type Status
      ---
      - - -
Errors:
-----
-
cluster1::>

```

The following example displays error information about the storage shelves that have errors:

```
cluster1::> storage shelf show -errors
Shelf Name: 1.1
    Shelf UID: 50:0a:09:80:01:cb:d6:84
    Serial Number: 6000832415
Error Type      Description
-----
Power           Critical condition is detected in storage shelf
power supply unit "1". The unit might fail.
Temperature     Critical over temperature failure for temperature
sensor "1". Current temperature: "75" C ("167" F).
```

storage shelf acp configure

Configure alternate control path (ACP)

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

Description

Configure the ACP connectivity on the cluster.

Parameters

-is-enabled {true|false} - Is Enabled?

Configures the connectivity to the specified state.

[-subnet <IP Address>] - Subnet

Configures the connectivity to the specified subnet.

[-netmask <IP Address>] - Netmask

Configures the connectivity to the specified netmask.

[-channel {out-of-band|in-band}] - Channel

Configures the connectivity to the specified channel.

Examples

The following example configures out-of-band ACP connectivity on each node:

```
cluster1::> storage shelf acp configure -is-enabled true -channel out-of-
band -subnet 192.168.0.1 -netmask 255.255.255.0
```

The following example configures in-band ACP connectivity on each node:

```
cluster1::> storage shelf acp configure -is-enabled true -channel in-band
```

The following example disables ACP connectivity on each node:

```
cluster1::> storage shelf acp configure -is-enabled false
```

storage shelf acp show

Show connectivity information

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

Description

Displays information about the ACP connectivity on each node

Parameters

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

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

| [-errors]

If you specify the -errors parameter, the command displays detailed information about all modules with errors.

| [-instance] }

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

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

Selects the nodes that match this parameter value.

[-is-enabled {true|false}] - Is Enabled?

Selects the nodes that are enabled or disabled.

[-port <text>] - Port

Selects the nodes that match the specified port on which ACP is configured.

[-address <IP Address>] - IP Address

Selects the nodes with the specified IP address.

[-subnet <IP Address>] - Subnet

Selects the nodes with the specified subnet.

[-netmask <IP Address>] - Netmask

Selects the nodes with the specified netmask.

[-connection-status {no-connectivity|partial-connectivity|full-connectivity|additional-connectivity|unknown-connectivity|not-available|connection-disabled}] - Connection Status

Selects the nodes with the specified connection status.

[-error-id <integer>] - Error ID

Selects the node with the specified error ID.

[-error-type {No-Error|Connection-Issue|Connection-Activity|Module-Error|Shelf-Error}] - Error Type

The error type, in case of a connection error.

[-error-severity {unknown|notice|warning|error|critical}] - Error Severity

The error severity, in case of a connection error.

[-error-text <text>] - Error Text

Selects the node with the specified error text.

[-corrective-action <text>] - Corrective Action

Selects the node with the specified corrective action.

[-channel {unknown|out-of-band|in-band}] - Channel

Selects the nodes that has channel configured out-of-band or in-band.

Examples

The following example displays ACP connectivity on each node (in-band):

```
fas2750-2n-rtp-1::> storage shelf acp show
Node                               Channel                               Connectivity
-----
fas2750-rtp-1a                    in-band                             active
fas2750-rtp-1b                    in-band                             active
2 entries were displayed.
```

The following example displays ACP connectivity on each node (out of band):

```
fas2750-2n-rtp-1::> storage shelf acp show
Node                               Channel                               Connectivity
-----
fas2750-rtp-1a                    out-of-band                         full-connectivity
fas2750-rtp-1b                    out-of-band                         full-connectivity
2 entries were displayed.
```

The following example displays the -instance output of the storage acp show (in-band) command. Use this command to display details on connectivity and configuration.

```
fas2750-2n-rtp-1::> storage shelf acp show -instance
Node: fas2750-rtp-1a
Channel: in-band
Enable Status: true
Connection Status: active
Node: fas2750-rtp-1b
Channel: in-band
Enable Status: true
Connection Status: active
2 entries were displayed.
```

The following example displays the `-instance` output of the `storage acp show (out-of-band)` command. Use this command to display details on connectivity and configuration.

```
fas2750-2n-rtp-1::> storage shelf acp show -instance
Node: fas2750-rtp-1a
Channel: out-of-band
Enable Status: true
Port: e0P
IP Address: 192.168.1.74
Subnet: 192.168.0.1
Netmask: 255.255.252.0
Connection Status: full-connectivity
Node: fas2750-rtp-1b
Channel: out-of-band
Enable Status: true
Port: e0P
IP Address: 192.168.1.75
Subnet: 192.168.0.1
Netmask: 255.255.252.0
Connection Status: full-connectivity
2 entries were displayed.
```

storage shelf acp module show

Show modules connected to the cluster

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

Description

Displays information about the modules connected to each node

Parameters

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

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

| [-errors]

If you specify the -errors parameter, the command displays detailed information about all modules with errors.

| [-instance] }

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

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

Selects the modules that match this parameter value.

[-mac-address <text>] - MAC Address

Selects the module that match the specified MAC address.

[-module-name <text>] - Module name

Selects the module that match the specified module name.

[-module-address <IP Address>] - IP Address

Selects the module that match the specified IP address.

[-protocol-version <text>] - Protocol Version

Selects the modules that match the specified protocol version.

[-firmware-version <text>] - Firmware Version

Selects the modules that match the specified firmware version.

[-acpa-id <integer>] - ACPA assigner ID

Selects the modules that match the specified ACPA ID.

[-shelf-serial-number <text>] - Shelf Serial Number

Selects the modules that match the specified shelf serial number.

[-iom-type {Unknown|iom3|iom6|iom6e|iom12|iom12e|iom12f|iom12b}] - IOM Type

Selects the modules that match the specified IOM type (IOM3/IOM6/IOM6E).

[-last-contact <integer>] - Last Contact (secs)

Selects the modules that match the specified last contact.

[-state {unknown|initializing|discovery-complete|awaiting-inband|no-inband|active|awaiting-bootp|updating-firmware|connection-error|firmware-update-required|rebooting|fail|unsupported|degraded|shelf-off}] - Local Node State

Selects the modules that match the specified state.

[-stack-id {<integer>|-}] - Stack ID

Selects the modules that match the specified stack ID.

[-shelf-id <text>] - Shelf ID

Selects the modules that match the specified shelf ID.

[-adapter-name <text>] - Adapter Name

Selects the modules that match the specified adapter name.

[-error-id <integer>,...] - Error ID

Selects the modules that match the specified error ID.

[-error-text <text>,...] - Error Text

The error text, in case of a module error.

[-corrective-action <text>,...] - Corrective Action

The corrective action, in case of a module error.

[-error-type {No-Error|Connection-Issue|Connection-Activity|Module-Error|Shelf-Error}] - Error Type

Selects the modules that match the specified error type.

[-error-severity {unknown|notice|warning|error|critical}] - Error Severity

Selects the modules that match the specified error severity.

[-power-cycle-count <integer>] - Power Cycle count

Number of times a shelf power cycle has been performed on a shelf

[-power-off-count <integer>] - Power Off count

Number of times a shelf power off has been performed on a shelf

[-power-on-count <integer>] - Power On count

Number of times a shelf power on has been performed on a shelf

[-expander-reset-count <integer>] - Expander reset count

Number of times an expander reset has been performed on a module

[-expander-power-cycle-count <integer>] - Expander power cycle count

Number of times an expander power cycle has been performed on a module

Examples

The following example displays the ACP modules connected to each node:

```
cluster1::> storage shelf acp module show
```

| Node | Module Name | State |
|------------------|-------------|--------|
| stor-v4-1a-1b-01 | 1.10.A | Active |
| | 1.10.B | Active |
| | 1.254.B | Active |
| | 1.254.A | Active |
| stor-v4-1a-1b-02 | 1.10.A | Active |
| | 1.10.B | Active |
| | 1.254.B | Active |
| | 1.254.A | Active |

8 entries were displayed.

The following example displays the -instance output of the storage shelf acp module show. More details on each module can be seen [here](#).

```
cluster1::> storage shelf acp module show -instance
```

```
Node: stor-v4-1a-1b-01
    Module Name: 1.10.A
    Mac Address: 00:a0:98:19:53:ee
    IOM Type: IOM6E
Shelf Serial Number: SHJMS000000001A
    IP Address: 192.168.3.239
    Protocol Version: 2.1.1.21
    Assigner ID: 2.1.1.21
    State: Active
    Last Contact: 203
    Power Cycle Count: 0
    Power Off Count: 0
    Power On Count: 0
    Expander Reset Count: 0
Expander Power Cycle Count: 0
Node: stor-v4-1a-1b-01
    Module Name: 1.10.B
    Mac Address: 00:a0:98:19:55:16
    IOM Type: IOM6E
Shelf Serial Number: SHJMS000000001A
    IP Address: 192.168.1.23
    Protocol Version: 2.1.1.21
    Assigner ID: 2.1.1.21
    State: Active
    Last Contact: 206
    Power Cycle Count: 0
    Power Off Count: 0
```

```
        Power On Count: 0
    Expander Reset Count: 0
Expander Power Cycle Count: 0
Node: stor-v4-1a-1b-01
    Module Name: 1.254.B
    Mac Address: 00:a0:98:32:d6:ac
        IOM Type: IOM6
    Shelf Serial Number: 6000368103
        IP Address: 192.168.2.173
    Protocol Version: 1.2.2. 8
        Assigner ID: 1.2.2. 8
            State: Active
        Last Contact: 215
    Power Cycle Count: 0
    Power Off Count: 0
    Power On Count: 0
    Expander Reset Count: 0
Expander Power Cycle Count: 0
Node: stor-v4-1a-1b-01
    Module Name: 1.254.A
    Mac Address: 00:a0:98:32:d6:dc
        IOM Type: IOM6
    Shelf Serial Number: 6000368103
        IP Address: 192.168.2.221
    Protocol Version: 1.2.2. 8
        Assigner ID: 1.2.2. 8
            State: Active
        Last Contact: 218
    Power Cycle Count: 0
    Power Off Count: 0
    Power On Count: 0
    Expander Reset Count: 0
Expander Power Cycle Count: 0
Node: stor-v4-1a-1b-02
    Module Name: 1.106.A
    Mac Address: 00:a0:98:19:53:ee
        IOM Type: IOM6E
    Shelf Serial Number: SHJMS000000001A
        IP Address: 192.168.3.239
    Protocol Version: 2.1.1.21
        Assigner ID: 2.1.1.21
            State: Initializing
        Last Contact: 206
    Power Cycle Count: 0
    Power Off Count: 0
    Power On Count: 0
```

```
    Expander Reset Count: 0
Expander Power Cycle Count: 0
Node: stor-v4-1a-1b-02
    Module Name: 1.106.B
    Mac Address: 00:a0:98:19:55:16
    IOM Type: IOM6E
Shelf Serial Number: SHJMS000000001A
    IP Address: 192.168.1.23
    Protocol Version: 2.1.1.21
    Assigner ID: 2.1.1.21
    State: Initializing
    Last Contact: 209
    Power Cycle Count: 0
    Power Off Count: 0
    Power On Count: 0
Expander Reset Count: 0
Expander Power Cycle Count: 0
Node: stor-v4-1a-1b-02
    Module Name: 1.10.B
    Mac Address: 00:a0:98:32:d6:ac
    IOM Type: IOM6
Shelf Serial Number: 6000368103
    IP Address: 192.168.2.173
    Protocol Version: 1.2.2.8
    Assigner ID: 1.2.2.8
    State: Initializing
    Last Contact: 217
    Power Cycle Count: 0
    Power Off Count: 0
    Power On Count: 0
Expander Reset Count: 0
Expander Power Cycle Count: 0
Node: stor-v4-1a-1b-02
    Module Name: 1.10.A
    Mac Address: 00:a0:98:32:d6:dc
    IOM Type: IOM6
Shelf Serial Number: 6000368103
    IP Address: 192.168.2.221
    Protocol Version: 1.2.2.8
    Assigner ID: 1.2.2.8
    State: Initializing
    Last Contact: 220
    Power Cycle Count: 0
    Power Off Count: 0
    Power On Count: 0
Expander Reset Count: 0
```

```
Expander Power Cycle Count: 0
```

```
8 entries were displayed.
```

storage shelf drawer show-phy

Display a list of PHYs per drawer

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

Description

The `storage shelf drawer show-phy` command displays information for drawer PHYs in the storage system. If no parameters are specified, the default command displays the following information about PHYs:

- Shelf Name
- Drawer Number
- PHY Number
- Type
- SAS Address
- State

To display detailed information about a single PHY, use the `-shelf`, `-drawer`, and `-phy` parameters.

Parameters

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

Displays the specified fields for all drawer PHYs, in column style output.

| [-instance]

Displays expanded information for all drawer PHYs in the system. If a shelf, drawer, and PHY are specified, then this parameter displays the same detailed information for the PHY you specify as does the `-shelf`, `-drawer`, and `-phy` parameters.

[-shelf <text>] - Shelf Name

Displays the PHYs in the storage shelf that matches the specified shelf name.

[-drawer <integer>] - Drawer Number

Displays the PHYs in the drawers that match the specified drawer number.

[-phy <integer>] - PHY Number

Displays the PHYs that match the specified PHY number.

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

Displays the PHYs that are present for the specified node.

[-type {unknown|disk|virtual|input}] - Type

Displays the PHYs with the specified type.

[-physical-id <integer>] - Physical ID

Displays the PHYs that match the specified physical-id.

[-sas-address <text>] - Attached SAS Address

Displays the PHYs with the specified attached sas address.

[-state-a {unknown|enabled|disabled}] - State Module A

Displays the PHYs for which module A has the specified state.

[-state-b {unknown|enabled|disabled}] - State Module B

Displays the PHYs for which module B has the specified state.

[-status-a <Drawer PHY Status>] - Status Module A

Displays the PHYs with module A currently operating under the specified status.

[-status-b <Drawer PHY Status>] - Status Module B

Displays the PHYs with module B currently operating under the specified status.

Examples

The following example displays information about all drawer PHYs:

```
cluster1::> storage shelf drawer show-phy
```

| Shelf | Drawer | PHY # | Type | SAS Address | PHY State A/B |
|-------|--------|-------|---------|------------------|-----------------|
| 2.5 | 1 | 0 | disk | 00c5005079183f85 | enabled/enabled |
| | | 1 | disk | - | enabled/enabled |
| | | 2 | disk | - | enabled/enabled |
| | | 3 | disk | 00c50050e1183f85 | enabled/enabled |
| | | 4 | disk | - | enabled/enabled |
| | | 5 | disk | - | enabled/enabled |
| | | 6 | disk | 00c50050dd183f85 | enabled/enabled |
| | | 7 | disk | - | enabled/enabled |
| | | 8 | disk | - | enabled/enabled |
| | | 9 | disk | 00c500502d163f85 | enabled/enabled |
| | | 10 | disk | - | enabled/enabled |
| | | 11 | disk | - | enabled/enabled |
| | | 12 | input | 80090a5045e46f06 | enabled/enabled |
| | | 13 | input | 80090a5045e46f06 | enabled/enabled |
| | | 14 | input | 80090a5045e46f06 | enabled/enabled |
| | | 15 | input | 80090a5045e46f06 | enabled/enabled |
| | | 16 | virtual | 8a090a503dd01b17 | enabled/enabled |

2

| | | | |
|----|---------|------------------|-----------------|
| 0 | disk | 00c500503d0e3d85 | enabled/enabled |
| 1 | disk | - | enabled/enabled |
| 2 | disk | - | enabled/enabled |
| 3 | disk | 00c50050e9173f85 | enabled/enabled |
| 4 | disk | - | enabled/enabled |
| 5 | disk | - | enabled/enabled |
| 6 | disk | 00c50050a9163f85 | enabled/enabled |
| 7 | disk | - | enabled/enabled |
| 8 | disk | - | enabled/enabled |
| 9 | disk | 00c5005021173f85 | enabled/enabled |
| 10 | disk | - | enabled/enabled |
| 11 | disk | - | enabled/enabled |
| 12 | input | 80090a5045e46f06 | enabled/enabled |
| 13 | input | 80090a5045e46f06 | enabled/enabled |
| 14 | input | 80090a5045e46f06 | enabled/enabled |
| 15 | input | 80090a5045e46f06 | enabled/enabled |
| 16 | virtual | 8a090a503d90fd16 | enabled/enabled |

3

| | | | |
|----|---------|------------------|-----------------|
| 0 | disk | 00c500503d163f85 | enabled/enabled |
| 1 | disk | - | enabled/enabled |
| 2 | disk | - | enabled/enabled |
| 3 | disk | 00c50050bd163f85 | enabled/enabled |
| 4 | disk | - | enabled/enabled |
| 5 | disk | - | enabled/enabled |
| 6 | disk | 00c50050c1d44085 | enabled/enabled |
| 7 | disk | - | enabled/enabled |
| 8 | disk | - | enabled/enabled |
| 9 | disk | 00c50050f1d54085 | enabled/enabled |
| 10 | disk | - | enabled/enabled |
| 11 | disk | - | enabled/enabled |
| 12 | input | 80090a5045e46f06 | enabled/enabled |
| 13 | input | 80090a5045e46f06 | enabled/enabled |
| 14 | input | 80090a5045e46f06 | enabled/enabled |
| 15 | input | 80090a5045e46f06 | enabled/enabled |
| 16 | virtual | 8a090a503d202a17 | enabled/enabled |

4

| | | | |
|---|------|------------------|-----------------|
| 0 | disk | 00c50050fdd54085 | enabled/enabled |
| 1 | disk | - | enabled/enabled |
| 2 | disk | - | enabled/enabled |
| 3 | disk | 00c50050d9d44085 | enabled/enabled |
| 4 | disk | a0cc0050e5973712 | enabled/enabled |
| 5 | disk | - | enabled/enabled |
| 6 | disk | 00c500506dd34085 | enabled/enabled |
| 7 | disk | - | enabled/enabled |
| 8 | disk | - | enabled/enabled |

| | | | | |
|---|----|---------|------------------|-----------------|
| | 9 | disk | 00c5005045d64085 | enabled/enabled |
| | 10 | disk | - | enabled/enabled |
| | 11 | disk | - | enabled/enabled |
| | 12 | input | 80090a5045e46f06 | enabled/enabled |
| | 13 | input | 80090a5045e46f06 | enabled/enabled |
| | 14 | input | 80090a5045e46f06 | enabled/enabled |
| | 15 | input | 80090a5045e46f06 | enabled/enabled |
| | 16 | virtual | 8a090a503d100b17 | enabled/enabled |
| 5 | | | | |
| | 0 | disk | 00c50050c9d54085 | enabled/enabled |
| | 1 | disk | - | enabled/enabled |
| | 2 | disk | - | enabled/enabled |
| | 3 | disk | 00c50050f9d44085 | enabled/enabled |
| | 4 | disk | - | enabled/enabled |
| | 5 | disk | - | enabled/enabled |
| | 6 | disk | 00c5005081d34085 | enabled/enabled |
| | 7 | disk | - | enabled/enabled |
| | 8 | disk | - | enabled/enabled |
| | 9 | disk | 00c500505dd64085 | enabled/enabled |
| | 10 | disk | - | enabled/enabled |
| | 11 | disk | - | enabled/enabled |
| | 12 | input | 80090a5045e46f06 | enabled/enabled |
| | 13 | input | 80090a5045e46f06 | enabled/enabled |
| | 14 | input | 80090a5045e46f06 | enabled/enabled |
| | 15 | input | 80090a5045e46f06 | enabled/enabled |
| | 16 | virtual | 8a090a503df00a17 | enabled/enabled |

85 entries were displayed.
cluster1::>

The following example displays expanded information for PHY 0 of drawer 1 in shelf 2.5:

```
cluster1::> storage shelf drawer show-phy -shelf 2.5 -drawer 1 -phy 0
Shelf: 2.5
  Drawer ID: 1
  PHY Number: 0
    Type: disk
  Physical ID: 1
  SAS Address: 00c5005079183f85
    State A: enabled
    State B: enabled
  Status A: enabled-12gbs
  Status B: enabled-12gbs
cluster1::>
```

storage shelf drawer show-slot

Display a map between bay number and drawer/slot number

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

Description

The `storage shelf drawer show-slot` command maps each drawer and slot number to the corresponding bay number.

Parameters

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

Displays the specified fields in column style output.

| [-instance]

Displays all slot information.

[-shelf <text>] - Shelf Name

Displays the slots in the shelf that matches the specified shelf name.

[-bay <integer>] - Bay Number

Displays the slots that have the specified bay number.

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

Displays the slots that are present for the specified node.

[-drawer <integer>] - Drawer Number

Displays the slots in the drawers that match the specified drawer number.

[-slot <integer>] - Slot Number

Displays the slots that match the specified slot number.

[-is-installed {yes|no}] - Is Disk Installed

Displays the slots that have a disk installed.

Examples

The following example displays the mapping from drawer and slot number to bay number:

```
cluster1:> storage shelf drawer show-slot
Shelf  Drawer  Slot  Bay  Installed?
-----
2.5
      1
        0    0  yes
        1    1  no
        2    2  no
```

| | | | |
|---|----|----|-----|
| | 3 | 3 | yes |
| | 4 | 4 | no |
| | 5 | 5 | no |
| | 6 | 6 | yes |
| | 7 | 7 | no |
| | 8 | 8 | no |
| | 9 | 9 | yes |
| | 10 | 10 | no |
| | 11 | 11 | no |
| 2 | | | |
| | 0 | 12 | yes |
| | 1 | 13 | no |
| | 2 | 14 | no |
| | 3 | 15 | yes |
| | 4 | 16 | no |
| | 5 | 17 | no |
| | 6 | 18 | yes |
| | 7 | 19 | no |
| | 8 | 20 | no |
| | 9 | 21 | yes |
| | 10 | 22 | no |
| | 11 | 23 | no |
| 3 | | | |
| | 0 | 24 | yes |
| | 1 | 25 | no |
| | 2 | 26 | no |
| | 3 | 27 | yes |
| | 4 | 28 | no |
| | 5 | 29 | no |
| | 6 | 30 | yes |
| | 7 | 31 | no |
| | 8 | 32 | no |
| | 9 | 33 | yes |
| | 10 | 34 | no |
| | 11 | 35 | no |
| 4 | | | |
| | 0 | 36 | yes |
| | 1 | 37 | no |
| | 2 | 38 | no |
| | 3 | 39 | yes |
| | 4 | 40 | yes |
| | 5 | 41 | no |
| | 6 | 42 | yes |
| | 7 | 43 | no |
| | 8 | 44 | no |
| | 9 | 45 | yes |

```

          10    46  no
          11    47  no
5
          0    48  yes
          1    49  no
          2    50  no
          3    51  yes
          4    52  no
          5    53  no
          6    54  yes
          7    55  no
          8    56  no
          9    57  yes
         10    58  no
         11    59  no
60 entries were displayed.
cluster1::>

```

storage shelf drawer show

Display a list of drawers

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

Description

The `storage shelf drawer show` command displays information for storage shelf drawers in the storage system. If no parameters are specified, the default command displays the following information for the drawers:

- Shelf Name
- Drawer Number
- Status
- Closed/Open
- Disk Count
- Firmware

To display detailed information for a single drawer, use the `-shelf` and `-drawer` parameters.

Parameters

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

Displays the specified fields for all drawers, in column style output.

| [-errors]

Displays the following error status information about the drawers that have errors:

- Status

- Error Description

[`-instance`] -

Displays expanded information for all drawers in the system. If a shelf and drawer are specified, then this parameter displays the same detailed information for the specified drawer as does the `-shelf` and `-drawer` parameters.

[`-shelf <text>`] - Shelf Name

Displays the drawers in the storage shelf that matches the specified shelf name.

[`-drawer <integer>`] - Drawer Number

Displays the drawers that match the specified drawer number.

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

Displays the drawers that are present for the specified node.

[`-disk-count <integer>`] - Drawer Disk Count

Displays the drawers that have the specified disk count.

[`-part-number <text>`] - Part Number

Displays the drawers that have the specified part number.

[`-serial-number <text>`] - Serial Number

Displays the drawer that matches the specified serial number.

[`-is-closed {open|closed}`] - Drawer is Closed?

Displays the drawers that are closed or open.

[`-firmware-a <text>`] - Firmware A

Displays the drawers for which module A has the specified firmware version.

[`-firmware-b <text>`] - Firmware B

Displays the drawers for which module B has the specified firmware version.

[`-path-a {unknown|ok|degraded|none}`] - Path A

Displays the drawers for which module A has the specified path status.

[`-path-b {unknown|ok|degraded|none}`] - Path B

Displays about drawers for which module B has the specified path status.

[`-is-supported {yes|no}`] - Drawer is Supported?

Displays the drawers that are supported (TRUE) or not supported (FALSE).

[`-vendor <text>`] - Vendor Name

Displays the drawers that match the specified vendor.

[`-mfg-date <text>`] - Mfg. Date

Displays the drawers that match the specified manufactured date.

[-fru-type <text>] - FRU Type

Displays the drawers that match the specified FRU type.

[-status-a {unknown|normal|warning|error|critical}] - Status A

Displays the drawers with module A currently operating under the specified status.

[-status-b {unknown|normal|warning|error|critical}] - Status B

Displays the drawers with module B currently operating under the specified status.

[-error <text>] - Error

Displays the drawers that match the specified error description.

Examples

The following example displays information about all drawers:

```
cluster1::> storage shelf drawer show
Drawer  Disk
Shelf Drawer    Status A/B    Closed? Count    Firmware A/B
-----
2.5
      1  normal/normal  closed      4  00000634/00000634
      2  normal/normal  closed      4  00000634/00000634
      3  normal/normal  closed      4  00000634/00000634
      4  normal/normal  closed      5  00000634/00000634
      5  normal/normal  closed      4  00000634/00000634
5 entries were displayed.
cluster1::>
```

The following example displays expanded information about drawer 1 in shelf 2.5:

```

cluster1::> storage shelf drawer show -shelf 2.5 -drawer 1
Shelf: 2.5
        Drawer ID: 1
        Part Numer: 111-03071
        Serial Number: 021604008153
Drawer is Closed?: closed
        Disk Count: 4
        Firmware A: 00000634
        Firmware B: 00000634
        Path A: ok
        Path B: ok
        Status A: normal
        Status B: normal
Drawer is Supported?: yes
        Vendor Name: NETAPP
        Mfg. Date: 02/2016
        FRU Type: SASDRWR
Error Description: -
cluster1::>

```

The following example displays error information about the drawers that have errors:

```

cluster1::> storage shelf drawer show -errors
Shelf Drawer      Status A/B      Error Description
-----
2.5
        2  warning/warning  Drawer open.
cluster1::>

```

storage shelf firmware show-update-status

Display the Shelf Firmware Update (SFU) Status.

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

Description

The `storage shelf firmware show-update-status` command displays the state of the Shelf Firmware Update process.

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 (privilege: advanced)

Selects the node that matches this parameter value.

[-update-status {running|idle}] - Disk Shelf Firmware Update Status (privilege: advanced)

Selects the nodes whose SFU process status matches this parameter value. Possible values are:

- running - Disk shelf firmware update is in progress.
- idle - Disk shelf firmware update is not in progress.

[-in-progress-count <integer>] - Number of Shelves with Earlier Revisions Being Updated (privilege: advanced)

Selects the nodes that matches the number of shelves the SFU process is updating to this parameter value. This specifies the number of shelves with earlier revisions that are being updated.

Examples

```
cluster1::*> storage shelf firmware show-update-status
```

| Node | Update Status | In-Progress Count |
|------------|---------------|-------------------|
| cluster-n1 | running | 10 |
| cluster-n2 | idle | - |
| cluster-n3 | running | 7 |

storage shelf firmware update

Update Shelf Firmware

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

Description

The `storage shelf firmware update` command updates the firmware on one or more shelves. You can download the latest firmware by using the [storage firmware download](#) command. You can specify a shelf whose firmware is to be updated by using the `-shelf` parameter. You can update the firmware on all the shelves by not providing the `-shelf` parameter. All the shelves of a specific module type can be updated by providing a value to the `-module-type` parameter.

Parameters

{ [-shelf <text>] - Shelf Name (privilege: advanced)

This specifies the name of the shelf whose firmware is to be updated.

[`-module-type {atfcx|esh4|iom3|iom6|iom6e|iom12|iom12e|iom12b|nsm100|nsm8e}`] - Shelf Module Type (privilege: advanced) }

Update the firmware on the shelves that match the module-type you specify.

[`-refresh <true>`] - Refresh (privilege: advanced)

Forces an update on the shelf with the highest revision of the applicable firmware, resulting in a refresh of the firmware image already present on the shelf.

Examples

The following example updates the firmware on all the shelves in the cluster:

```
cluster1::*> storage shelf firmware update
```

The following example updates the firmware on all shelves with the IOM6 module type:

```
cluster1::*> storage shelf firmware update -module-type IOM6
```

The following example updates the firmware on shelf 1.2:

```
cluster1::*> storage shelf firmware update -shelf 1.2
```

The following example refreshes the firmware on all shelves with the IOM6 module type:

```
cluster1::*> storage shelf firmware update -refresh -module-type IOM6
```

The following example refreshes the firmware on shelf 1.2:

```
cluster1::*> storage shelf firmware update -refresh -shelf 1.2
```

Related Links

- [storage firmware download](#)

storage shelf location-led modify

Modify the state of the shelf Location LED

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

Description

The `storage shelf location-led modify` command modifies the on/off state of the shelf location LED.

Parameters

-shelf-name <text> - Shelf Name

This parameter specifies the shelf whose LED is to be turned on or turned off.

[-led-status {on|off}] - Location LED

This parameter specifies whether the shelf location LED needs to be turned on or turned off.

Examples

The following example turns on the shelf location LED of the specified shelf.

```
cluster1::> storage shelf location-led modify -node node1 -shelf-name 1.0  
-led-status on
```

```
Info: Shelf locate request successful for shelf "1.0".
```

The following example turns off the shelf location LED of the specified shelf.

```
cluster1::> storage shelf location-led modify -node node1 -shelf-name 1.0  
-led-status off
```

```
Info: Shelf locate request successful for shelf "1.0".
```

storage shelf location-led show

Display the Location LED status

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

Description

The `storage shelf location-led show` command displays the state of shelf location LED.

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.

[-shelf-name <text>] - Shelf Name

Selects the shelves whose shelf-name matches this parameter value.

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

Selects the nodes that match this parameter value.

[`-stack-id <integer>`]- Stack ID

Selects the shelves whose stack-id matches this parameter value.

[`-shelf-id <integer>`]- Shelf ID

Selects the shelves whose shelf-id matches this parameter value.

[`-led-status {on|off}`}] - Location LED

Shows the state of the shelf location LED.

Examples

The following example shows the state of the shelf location LED for each shelf.

```
cluster1::> storage shelf location-led show
Shelf Name Stack ID Shelf ID LED Status
-----
      8.2         8         2 off
      8.3         8         3 off
      6.0         6         0 unsupported
      8.1         8         1 off
4 entries were displayed.
```

storage shelf port show

Display storage shelf ports

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

Description

The `storage shelf port show` command displays information for storage shelf ports in the storage system. If no parameters are specified, the default command displays the following information for the ports:

- Shelf Name
- ID
- Module
- State
- Internal?

To display detailed information for a single port, use the `-shelf` and `-id` parameters.

Parameters

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

Displays output in column style about the specified fields for all shelf ports.

| [-cables]

Displays information about all cables connected to the shelf ports.

| [-instance]

Displays expanded information for all shelf ports in the system. If a shelf and ID are specified, then this parameter displays the same detailed information for the specified port as does the -shelf and -id parameters.

[-shelf <text>] - Shelf Name

Displays the ports in the storage shelf that matches the specified shelf name.

[-id <integer>] - Port ID

Displays the ports that match the specified ID.

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

Displays the ports that are present for the specified node.

[-module-id {A|B}] - Module ID

Displays the ports from the specified shelf module ID.

[-is-internal {true|false}] - Is Port Internal?

Displays the ports that are internal.

[-location <text>] - Location

Displays the ports with the specified location.

[-is-cable-connected {true|false}] - Is Cable Connected?

Displays the ports that have cables connected to them.

[-is-error {true|false}] - Any Errors?

Displays the ports for which errors have been logged.

[-connector-state {connected|disconnected|error}] - Connector State

Displays the ports with the specified connector state.

[-connector-serial-number <text>] - Connector Serial Number

Displays the ports with the specified connector serial number.

[-connector-type {QSFP|QSFP+|QSFP28|Mini-SAS-HD}] - Connector Type

Displays the ports with the specified connector type.

[-cable-vendor <text>] - Cable Vendor

Displays the ports that are connected to a cable from the specified vendor.

[-cable-part-number <text>] - Cable Part Number

Displays the ports that are connected to a cable with the specified part number.

[-cable-technology {active-copper|passive-copper|optical}] - Cable Technology

Displays the ports that are connected to a cable with the specified technology.

[-cable-length <text>] - Cable Length

Displays the ports that are connected to a cable with the specified length.

[-cable-id <text>] - Cable ID

Displays the ports that are connected to a cable with the specified ID.

[-cable-end {end_0|end_1}] - Cable End

Displays the ports that are connected to a cable with the specified cable end.

[-designator <text>] - Designator

Displays the ports with the specified designator.

[-wwn <text>] - Local Device WWN

Displays the ports with the specified WWN.

[-remote-wwn <text>] - Remote Device WWN

Displays the ports connected to the specified remote WWN.

[-remote-phy <text>] - Remote Phy

Displays the ports connected to the specified remote PHY.

[-swap-count <integer>] - Swap Count

Displays the ports with the specified swap count.

[-mac <MAC Address>] - Local MAC Address

Displays the ports with the specified MAC address.

[-remote-mac <MAC Address>] - Remote MAC Address

Displays the ports connected to the specified MAC address.

[-remote-port <text>] - Remote Port

Displays the ports connected to the specified port.

[-remote-chassis <text>] - Remote Chassis

Displays the ports connected to the specified chassis.

[-remote-device <text>] - Remote Device

Displays the ports connected to the specified device.

[-vlan-id <integer>] - VLAN ID

Displays the ports with the specified VLAN ID.

Examples

The following example displays information about all shelf ports:

```
cluster1::> storage shelf port show
```

| Shelf | ID | Module | State | Internal? |
|-------|----|--------|-----------|-----------|
| 1.4 | | | | |
| | 0 | A | connected | false |
| | 1 | A | connected | false |
| | 2 | B | connected | false |
| | 3 | B | connected | false |

4 entries were displayed.

The following example displays expanded information about port 0 in shelf 1.4:

```
cluster1::> storage shelf port show -shelf 1.4 -id 0
Shelf Name: 1.4
      Port ID: 0
      Module ID: A
      Is Port Internal?: false
      Location: rear of the shelf at the top left, on shelf
module A
      Is Cable Connected?: true
      Any Errors?: false
      Connector State: connected
Connector Serial Number: 616930439
      Connector Type: qsfp+
      Cable Vendor: Molex Inc.
Cable Part Number: 112-00431+A0
      Cable Technology: passive-copper
      Cable Length: 5m
      Cable ID: 500a0980000b6c3f-50000d1703544b80
      Cable End: end_1
      Designator: sqr
      Local Device WWN: 500A0980000B6C3F
      Remote Device WWN: 50000D1703544B80
      Remote Phy: 12
      Swap Count: 0
```

The following example displays information about the cables:

```
cluster1::> storage shelf port show -cables
```

```
Shelf: 1.4
```

| ID Number | Vendor | Part Number | Technology | Length | Type | Serial |
|-----------|------------|--------------|----------------|--------|-------|-----------|
| 0 | Molex Inc. | 112-00431+A0 | passive-copper | 5m | qsfp+ | 616930439 |
| 1 | Molex Inc. | 112-00431+A0 | passive-copper | 5m | qsfp+ | 616930364 |
| 2 | Molex Inc. | 112-00431+A0 | passive-copper | 5m | qsfp+ | 616930452 |
| 3 | Molex Inc. | 112-00431+A0 | passive-copper | 5m | qsfp+ | 616930474 |

4 entries were displayed.

storage tape commands

storage tape offline

Take a tape drive offline

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

Description

This command takes the specified tape drive offline.

Parameters

-node {<nodename>|local} - Node

Use this parameter to specify the node to which the tape drive is attached.

{ -name <text> - Tape Drive Device Name

Use this parameter to specify the device name of the tape drive that needs to be taken offline. The format of the device `-name` name includes a prefix to specify how the tape cartridge is handled and a suffix to describe the density of the tape. The prefix suggests 'r', 'nr' or 'ur' for rewind, no rewind, or unload/reload and a suffix shows density of 'l', 'm', 'h' or 'a'. For example, a tape device name for this operation might have the form "nrst8m" where 'nr' is the 'no rewind' prefix, 'st8' is the alias-name and 'm' is the tape density. You can use the 'storage tape show -device-names' command to find more information about device names of tape drives attached to a node.

| -device-id <text> - Tape Drive Device ID }

Use this parameter to specify the device ID of the tape drive that needs to be taken offline.

Examples

The following example takes the tape drive with device name 'nrst8m' offline. This tape drive is attached to cluster1-01.

```
cluster1::> storage tape offline -node cluster1-01 -name nrst8m
```

storage tape online

Bring a tape drive online

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

Description

This command brings a specified tape drive online.

Parameters

-node {<nodename>|local} - Node

Use this parameter to specify the node to which the tape drive is attached.

{ -device-id <text> - Tape Drive Device ID

Use this parameter to specify the device ID of the tape drive that needs to be brought online.

| -name <text> - Tape Drive Device Name }

Use this parameter to specify the device name of the tape drive that needs to be brought online. The format of the device `-name` name includes a prefix to specify how the tape cartridge is handled and a suffix to describe the density of the tape. The prefix suggests 'r', 'nr' or 'ur' for rewind, no rewind, or unload/reload and a suffix shows density of 'l', 'm', 'h' or 'a'. For example, a tape device name for this operation might have the form "nrst8m" where 'nr' is the 'no rewind' prefix, 'st8' is the alias-name and 'm' is the tape density. You can use the 'storage tape show -device-names' command to find more information about device names of tape drives attached to a node.

Examples

The following example brings the tape drive with device id sw4:2.126L4 attached to the node, cluster1-01, online.

```
cluster1::> storage tape online -node cluster1-01 -device-id sw4:2.126L4
```

storage tape position

Modify a tape drive cartridge position

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

Description

This command changes the tape drive cartridge position.

Parameters

-node {<nodename>|local} - Node

Use this parameter to specify the node to which the tape drive is attached.

-name <text> - Tape Drive Device Name

Use this parameter to specify the device name of the tape drive whose cartridge position needs to be changed. The format of the device `-name` includes a prefix to specify how the tape cartridge is handled and a suffix to describe the density of the tape. The prefix suggests 'r', 'nr' or 'ur' for rewind, no rewind, or unload/reload and a suffix shows density of 'l', 'm', 'h' or 'a'. For example, a tape device name for this operation might have the form "nrst8m" where 'nr' is the 'no rewind' prefix, 'st8' is the alias-name and 'm' is the tape density. You can use the 'storage tape show -device-names' command to find more information about device names of tape drives attached to a node.

-operation {weof|fsf|bsf|fsr|bsr|rewind|erase|eom} - Tape Position Operation

Use this parameter to specify the tape positioning operation. The possible values for `-operation` are:

- weof - Write end-of-file marks
- fsf - Forward space end-of-file marks
- bsf - Backward space end-of-file marks
- fsr - Forward space records
- bsr - Backward space records
- rewind - Rewind the tape
- erase - Erase then entire tape media from current position
- eom - Position the tape at end of data (end of media if full)

[-count <integer>] - Count for Positioning

Use this parameter to specify the count for a tape positioning operation. You can specify this parameter only with the following operations: weof, fsf, bsf, fsr, and bsr. The default value of this parameter is one.

Examples

The following example specifies a rewind operation on a tape device. Note the `-count` parameter does not need to be specified for this type of operation.

```
cluster1::> storage tape position -node cluster1-01 -name nrst8m
-operation rewind
```

The following example specifies an fsf (forward space filemark) operation on a tape device. Note the `-count` parameter specifies 5 forward space filemarks for this operation.

```
cluster1::> storage tape position -node cluster1-01 -name nrst1a
-operation fsf -count 5
```

The following example specifies an eom (end-of-media) operation on a tape device. The 'eom' positions a tape at end of data (end of media if full). Note the `-count` parameter does not need to be specified for this type of

operation.

```
cluster1::> storage tape position -node cluster1-01 -name rst0h -operation eom
```

storage tape reset

Reset a tape drive

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

Description

This command resets a specified tape drive.

Parameters

-node {<nodename>|local} - Node (privilege: advanced)

Use this parameter to specify the node to which the tape drive is attached.

-device-id <text> - Tape Drive Device ID (privilege: advanced)

Use this parameter to specify the device ID of the tape drive is to be reset.

Examples

The following example resets the tape drive with device ID sw4:2.126L3 attached to the node, cluster1-01.

```
cluster1::> storage tape reset -node cluster1-01 -device-id sw4:2.126L3
```

storage tape show-errors

Display tape drive errors

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

Description

The `storage tape show-errors` command displays error information about tape drives attached to the cluster. By default, this command displays the following information about all tape drives:

- Node to which the tape drive is attached
- Device ID of the tape drive
- Type of device(tape drive)
- Description of the tape drive
- Alias name of the tape drive
- Tape drive errors

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

Displays detailed information about tape drives on the specified node.

[-device-id <text>] - Device ID

Selects the tape drive with the specified device ID.

[-device-type <text>] - Device Type

Selects the devices with the specified type of tape drive.

[-description <text>] - Description

Selects the tape drives with the specified description.

[-alias-name <text>] - Alias Name

Selects the tape drive with the specified alias name.

[-wwnn <text>] - World Wide Node Name

Selects the tape drives with the specified World Wide Node Name.

[-wwpn <text>] - World Wide Port Name

Selects the tape drive with the specified World Wide Port Name.

[-serial-number <text>] - Serial Number

Selects the tape drive with the specified serial number.

[-error <text>] - Tape Drive Error Description

Selects the tape drives with the specified error string.

[-initiator <text>] - Initiator Port

Selects the tape drives with the specified initiator port.

Examples

The following example shows error information for all tape drives attached to cluster1.

```
cluster1::> storage tape show-errors
      Node: node1
      Device ID: 0d.125
      Device Type: tape drive
      Description: Hewlett-Packard LTO-5
      Alias: st0
Errors: hardware error; repair or replace tape drive
Node: node1
      Device ID: 2d.0
      Device Type: tape drive
      Description: IBM LTO-6 ULT3580
      Alias: st2
Errors: -
```

The following example shows error information for tape drive sw4:2.126L1 attached to the node, node1.

```
cluster1::> storage tape show-errors -device-id sw4:2.126L1 -node node1
Node: node1
      Device ID: sw4:2.126L1
      Device Type: tape drive
      Description: Hewlett-Packard LTO-3
      Alias: st3
Errors: -
```

storage tape show-media-changer

Display information about media changers

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

Description

This `storage tape show-media-changer` command displays information about media changers attached to the cluster. By default, this command displays the following information about all media changers:

- Device ID of media changer
- Description of media changer
- World Wide Node Name of media changer
- World Wide Port Name of media changer
- Serial number of media changer
- Media changer errors
- Node to which the media changer is attached
- Initiator port which hosts the media changer

- Alias name of media changer
- Operational state of media changer
- Functional status of media changer

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.

[-device-id <text>] - Device ID

Selects the media changer with the specified device ID.

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

Displays detailed information about media changers on the specified node.

[-description <text>] - Description

Selects the media changers with the specified description.

[-alias-name <text>] - Alias Name

Selects the media changer with the specified alias name.

[-wwnn <text>] - World Wide Node Name

Selects the media changers with the specified World Wide Node Name.

[-wwpn <text>] - World Wide Port Name

>Selects the media changer with the specified World Wide Port Name.

[-serial-number <text>] - Serial Number

Selects the media changer with the specified serial number.

[-device-if-type {unknown|fibre-channel|SAS|pSCSI}] - Device If Type

Selects the media changers with the specified interface type.

[-device-state {unknown|available|ready-write-enabled|ready-write-protected|offline|in-use|error|reserved-by-another-host|normal}] - Operational State of Device

Selects the media changers with the specified operational state.

[-error <text>] - Media Changer Error Description

Selects the media changers with the specified error string.

[-initiator <text>] - Initiator Port

Selects the media changers with the specified initiator port.

Examples

The following example displays information about all media changers attached to the cluster:

```
cluster1::> storage tape show-media-changer
```

```
Media Changer: sw4:10.11L1
```

```
Description: PX70-TL
```

```
WWNN: 2:00a:000e11:10b919
```

```
WWPN: 2:00b:000e11:10b919
```

```
Serial Number: 00FRU7800000_LL1
```

```
Errors: -
```

```
Paths:
```

| Node | Initiator | Alias | Device State |
|-------------|-----------|-------|--------------|
| cluster1-01 | 2b | mc0 | in-use |

```
Media Changer: sw4:12.4L1
```

```
Description: NEO-TL
```

```
WWNN: 2:001:000e11:10b919
```

```
WWPN: 2:002:000e11:10b919
```

```
Serial Number: 00FRU7800000_LL0
```

```
Errors: -
```

```
Paths:
```

| Node | Initiator | Alias | Device State |
|-------------|-----------|-------|--------------|
| cluster1-01 | 5a | mc1 | available |

storage tape show-supported-status

Displays the qualification and supported status of tape drives

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

Description

This command displays the supported and qualification status of all tape drives recognized by Data ONTAP attached to a node in the cluster. This includes nonqualified tape drives. Such tape drives do not have a Tape Configuration File (TCF) on the storage system. A nonqualified tape drive can be used if the tape drive emulates a qualified tape drive or if the appropriate TCF for the nonqualified tape drive is downloaded from the

NetApp Support Site to the storage 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.

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

Selects the tape drives that match this parameter value.

[-tape-drive <text>] - Tape Drive Name

Selects the tape drives that match this parameter value.

[-is-supported {true|false}] - Tape Drive Supported

Selects the tape drives that match this parameter value.

[-status <text>] - Supported Status

Selects the tape drives that match this parameter value.

Examples

The following example displays support and qualification status of tape drives recognized by Data ONTAP. The command also identifies tape drives attached to the node that are nonqualified (not supported).

```
cluster1::> storage tape show-supported-status
```

```
Node: Node1
```

| Tape Drive | Is Supported | Support Status |
|---------------------------|--------------|-------------------------|
| sw4:2.126L6 | false | Nonqualified tape drive |
| Hewlett-Packard C1533A | true | Qualified |
| Hewlett-Packard C1553A | true | Qualified |
| Hewlett-Packard Ultrium 1 | true | Qualified |
| Sony SDX-300C | true | Qualified |
| Sony SDX-500C | true | Qualified |
| StorageTek T9840C | true | Dynamically Qualified |
| StorageTek T9840D | true | Dynamically Qualified |
| Tandberg LTO-2 HH | true | Dynamically Qualified |

The following example displays support and qualification status of tape drives selected by `□tape-drive`. The command identifies the supported status of the selected tape drive.

```
cluster1::> storage tape show-supported-status -tape-drive "Sony SDX-300C"
```

```
Node: Node1
```

| Tape Drives | Is Supported | Support Status |
|---------------|--------------|----------------|
| Sony SDX-300C | true | Qualified |

storage tape show-tape-drive

Display information about tape drives

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

Description

This `storage tape show-tape-drive` command displays information about tape drives attached to the cluster. By default, this command displays the following information about all tape drives:

- Device ID of tape drive
- Description of tape drive
- World Wide Node Name of tape drive
- World Wide Port Name of tape drive
- Serial Number of tape drive
- Tape drive errors
- Node to which the tape drive is attached

- Initiator port which hosts the tape drive
- Alias name of tape drive
- Operational state of tape drive
- Functional status of tape drive

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.

[-device-id <text>] - Device ID

Selects the tape drive with the specified device ID.

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

Displays detailed information about tape drives on the specified node.

[-description <text>] - Description

Selects the tape drives with the specified description.

[-alias-name <text>] - Alias Name

Selects the tape drive with the specified alias name.

[-wwnn <text>] - World Wide Node Name

Selects the tape drives with the specified World Wide Node Name.

[-wwpn <text>] - World Wide Port Name

Selects the tape drive with the specified World Wide Port Name.

[-serial-number <text>] - Serial Number

Selects the tape drive with the specified serial number.

[-device-if-type {unknown|fibre-channel|SAS|pSCSI}] - Device If Type

Selects the tape drives with the specified interface type.

[-device-state {unknown|available|ready-write-enabled|ready-write-protected|offline|in-use|error|reserved-by-another-host|normal}] - Operational State of Device

Selects the tape drives with the specified operational state.

[-error <text>] - Tape Drive Error Description

Selects the tape drives with the specified error string.

[-initiator <text>] - Initiator Port

Selects the tape drives with the specified initiator port.

[*-resv-type* {*off*|*persistent*|*scsi*}] - Reservation type for device

Selects the tape drives with the specified type.

Examples

The following example displays information about all tape drives attached to the cluster:

```
cluster1::> storage tape show-tape-drive
Tape Drive: sw4:11.126
  Description: StorageTek T10000C
    WWNN: 5:001:04f000:b39ec8
    WWPN: 5:001:04f000:b39ec9
  Serial Number: 576004000041
  Errors: -

  Paths:
  Node           Initiator  Alias      Device State
  Status
  -----
  -----
cluster1-01      2a          st0        ready-write-enabled
normal
Tape Drive: sw4:12.4
  Description: HP LTO-3
    WWNN: 2:001:000e11:10b919
    WWPN: 2:002:000e11:10b919
  Serial Number: 1068000371
  Errors: -

  Paths:
  Node           Initiator  Alias      Device State
  Status
  -----
  -----
cluster1-01      0b          st1        ready-write-enabled
normal
```

storage tape show

Display information about tape drives and media changers

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

Description

The `storage tape show` command displays information about tape drives and media changers attached to the cluster. Where it appears in the remainder of this document "device" may refer to either a tape drive or a

media changer. By default, this command displays the following information about all tape drives and media changers:

- Node to which the tape drive/media changer is attached
- Device ID of the tape drive/media changer
- Description of the tape drive/media changer
- Type of device: tape drive or media changer
- Functional status of the tape drive/media changer

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.

| [-alias]

Displays the tape drive/media changer alias with the following details:

- Node to which tape drive/media changer is attached
- Device ID of the tape drive/media changer
- Alias name of the tape drive/media changer
- Alias mapping for tape drive/media changer

| [-connectivity]

Displays the connectivity from the node to the tape drive/media changer with the following details:

- Node to which tape drive/media changer is attached
- Device ID of the tape drive/media changer
- Tape drive/media changer description
- Type of device: tape drive or media changer
- Interface type for the tape drive/media changer
- World Wide Node Name of tape drive/media changer
- World Wide Port Name of tape drive/media changer
- Serial Number of tape drive/media changer
- Tape drive/media changer errors
- Initiator port which hosts the tape drive/media changer
- Alias name of the tape drive/media changer
- Operational state of tape drive/media changer
- Functional status of tape drive/media changer

| [-device-names]

Displays the tape drive names for used tape positioning using the following details: rewind, no rewind, unload/reload and density

- Node to which tape drive/media changer is attached
- Device ID of the tape drive/media changer
- Tape drive/media changer description
- Device Names that include Rewind, no Rewind, Unload/Reload

| [-status]

Displays the status of tape drive/media changer with the following details:

- Device ID of the tape drive/media changer
- Tape drive/media changer description
- World Wide Node Name of tape drive/media changer
- World Wide Port Name of tape drive/media changer
- Serial Number of tape drive/media changer
- Alias name of the tape drive/media changer
- Format used for tape cartridge mounted by tape drive
- Tape drive/media changer errors
- Node to which tape drive/media changer is attached
- Operational state of tape drive/media changer
- File number following last tape drive I/O operation
- Block number following last tape drive I/O operation
- Residual count following last tape drive I/O operation

| [-instance] }

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

[-device-id <text>] - Device ID

Selects the tape drive/media changer with the specified device ID.

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

Displays detailed information about tape drives or media changers on the specified node.

[-device-type <text>] - Device Type

Selects the devices with the specified type of tape drive or media changer.

[-description <text>] - Description

Selects the tape drives/media changers with the specified description.

[-alias-name <text>] - Alias Name

Selects the tape drive/media changer with the specified alias name.

[-alias-mapping <text>] - Alias Mapping

Selects the tape drive/media changer with the specified alias mapping.

[-wwnn <text>] - World Wide Node Name

Selects the tape drives/media changers with the specified World Wide Node Name.

[-wwpn <text>] - World Wide Port Name

Selects the tape drive/media changer with the specified World Wide Port Name.

[-serial-number <text>] - Serial Number

Selects the tape drive/media changer with the specified serial number.

[-functional-status {unknown|normal|error}] - Functional Status

Selects the tape drives/media changers with the specified functional status of the device.

[-device-if-type {unknown|fibre-channel|SAS|pSCSI}] - Device Interface Type

Selects the tape drives/media changers with the specified interface type.

[-device-state {unknown|available|ready-write-enabled|ready-write-protected|offline|in-use|error|reserved-by-another-host|normal}] - Operational State of Device

Selects the tape drives/media changers with the specified operational state.

[-format <text>,...] - Tape Cartridge Format

Selects the tape drives with the specified tape format.

[-error <text>] - Tape Error

Selects the tape drives/media changers with the specified error string.

[-initiator <text>] - Initiator Port

Selects the tape drives/media changers with the specified initiator port.

[-file-number <integer>] - File Number

Selects the tape drives/media changers with the specified file number. The file number is the number of file marks between the beginning of media and current logical position. File number gets modified on write file mark, and forward or backward space file operations. A value of -1 indicates unknown position on the tape media or tape not loaded in the tape drive.

[-block-number <integer>] - Block Number

Selects the tape drives/media changers with the specified block number. The block number is the number of logical blocks between the beginning of tape media or the prior file mark and the current logical position on the tape media. Block number gets modified on writes, reads, and forward or backward space over records (blocks). The block number also gets reset to zero when a file mark is crossed or another file mark is written that designates a new file. If the tape is back spaced to a prior file mark, the block number might be zeroed. A value of -1 indicates unknown position on the tape media or that a tape not loaded in the tape drive.

[-residual-count <integer>] - Residual Count of Last I/O Operation

Selects the tape drives with the specified residual count.

[-device-name-r <text>,...] - Device Name for Rewind

Selects the tape drives with the specified device name for rewind.

[-device-name-nr <text>,...] - Device Name for No Rewind

Selects the tape drives with the specified device name for no rewind.

[-device-name-ur <text>,...] - Device Name for Unload Reload

Selects the tape drives with the specified device name for unload/reload.

[-resv-type {off|persistent|scsi}] - Reservation Type for device

Selects the tape drives with the specified type.

Examples

The following example displays information about all tape drives and media changers attached to the cluster:

```
cluster1::> storage tape show
```

```
Node: cluster1-01
```

| Device ID | Device Type | Description |
|-----------|-------------|-------------|
|-----------|-------------|-------------|

| Status | | |
|--------|--|--|
|--------|--|--|

| | | |
|-------|-------|-------|
| ----- | ----- | ----- |
| ----- | | |

| | | |
|-----------|------------|----------|
| sw4:10.11 | tape drive | HP LTO-3 |
|-----------|------------|----------|

| | | |
|-------|--|--|
| error | | |
|-------|--|--|

```
Node: cluster1-01
```

| Device ID | Device Type | Description |
|-----------|-------------|-------------|
|-----------|-------------|-------------|

| Status | | |
|--------|--|--|
|--------|--|--|

| | | |
|-------|-------|-------|
| ----- | ----- | ----- |
| ----- | | |

| | | |
|-------------|---------------|---------|
| sw4:10.11L1 | media changer | PX70-TL |
|-------------|---------------|---------|

| | | |
|--------|--|--|
| normal | | |
|--------|--|--|

The following example displays detailed information about a tape drive named sw4:10.11

```
cluster1::> storage tape show -device-id sw4:10.11
```

```
Node: cluster1-01
```

| Device ID | Device Type | Description |
|-----------|-------------|-------------|
|-----------|-------------|-------------|

| Status | | |
|--------|--|--|
|--------|--|--|

| | | |
|-------|-------|-------|
| ----- | ----- | ----- |
| ----- | | |

| | | |
|-----------|------------|----------|
| sw4:10.11 | tape drive | HP LTO-3 |
|-----------|------------|----------|

| | | |
|-------|--|--|
| error | | |
|-------|--|--|

storage tape trace

Enable/disable tape trace operations

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

Description

This command enables or disables diagnostic tape trace operations for all tape drives attached to the node you have specified.

Parameters

-node {<nodename>|local} - Node

Use this parameter to specify the node on which the tape trace feature is enabled or disabled.

[-is-trace-enabled {true|false}] - Tape Trace Enabled or Disabled

Use this parameter to enable or disable the tape trace feature. By default, the tape trace feature is enabled.

Examples

The following example enables tape trace operation on the node, cluster1-01.

```
cluster1::> storage tape trace -node cluster1-01 -is-trace-enabled true
```

storage tape alias clear

Clear alias names

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

Description

This command clears alias names for a tape drive or media changer.

Parameters

-node {<nodename>|local} - Node

Use this parameter to specify the node to which the tape drive is attached.

{ -name <text> - Alias Name That Is to Be Cleared

Use this parameter to specify the alias name that is to be cleared. You can use the 'storage tape show -alias' command to find more information about alias names of tape drives and media changers attached to a node. The `-clear-scope` and `-name` parameters are mutually exclusive. If you specify the `-name` parameter, a single alias name is cleared.

| -clear-scope {tape|media-changer|all} - Scope of Alias Clear Operation }

Use this parameter to specify the scope of the alias clear operation. The `-clear-scope` and `-name` parameters are mutually exclusive. If you specify the `-clear-scope` parameter, multiple aliases are cleared depending upon the value of the parameter.

The possible values for `-clear-scope` are as follows:

- `tape` - Clear all tape drive aliases

- media-changer - Clear all media-changer aliases
- all - Clear both tape drive and media-changer aliases

Examples

The following example clears an alias name 'st3' attached to the node, cluster1-01.

```
cluster1::> storage tape alias clear -node cluster1-01 -name st3
```

The following example clears all tape drive alias names attached to the node, cluster1-01.

```
cluster1::> storage tape alias clear -node cluster1-01 -clear-scope tape
```

The following example clears all media changer alias names attached to the node, cluster1-01.

```
cluster1::> storage tape alias clear -node cluster1-01 -clear-scope media-changer
```

The following example clears both tape and media changer alias names attached to the node, cluster1-01.

```
cluster1::> storage tape alias clear -node cluster1-01 -clear-scope all
```

storage tape alias set

Set an alias name for tape drive or media changer

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

Description

This command sets an alias name for a tape drive or media changer.

Parameters

-node {<nodename>|local} - Node

Use this parameter to specify the node to which the tape drive is attached.

-name <text> - Alias Name for Tape Drive or Media Changer

Use this parameter to specify the alias name for tape drive or media changer. For a tape drive alias name, the format is 'st' followed by one or more digits. For a media changer alias name, the format is 'mc' followed by one or more digits.

-mapping <text> - Mapping for Alias Name

Use this parameter to specify the mapping for an alias name. Use the format 'SN[<serial-number>]'. Valid mapping for serial numbers are in the format 'SN[<serial-number>]' where the <serial-

number> is from 2 to 90 characters long and includes the following characters: 0-9, a-z, and A-Z.

Examples

The following example sets an alias name 'st3' for a tape drive with serial number SN[123456]L4 attached to the node, node1.

```
cluster1::storage tape alias> set -node node1 -name st3 -mapping  
SN[123456]L4.
```

The following example sets an alias name 'mc1' for a media changer with serial number SN[65432] attached to the node, node1.

```
cluster1::storage tape alias> set -node node1 -name mc1 -mapping  
SN[65432].
```

storage tape alias show

Displays aliases of all tape drives and media changers

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

Description

This command displays aliases of all tape drives and media changers attached to every node 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.

Examples

The following example shows the aliases of all tape drives and media changers attached to every node in the cluster:

```
cluster1::> storage tape alias show
```

Node: node1

| Alias | Mapping |
|-------|------------------------------|
| mc0 | SN[00FRU7800000_LL0]L1 |
| mc1 | SN[00FRU7800000_LL1]L1 |
| mc2 | SN[aa6a64c69360a0980248c8] |
| mc3 | SN[c940abe8b0c3a0980248c8] |
| mc4 | SN[fba082e6b335a0980248c8]L5 |
| st0 | SN[HU19487T7N] |
| st1 | SN[1068000230] |
| st10 | SN[fba0c508b335a0980248c8]L7 |

Node: node2

| Alias | Mapping |
|-------|----------------------|
| mc1 | SN[c940982fc48c8] |
| st3 | SN[ST456HT8N]L3 |
| st2 | SN[HG68000230]L2 |
| st11 | SN[aba673980248c8]L7 |

storage tape config-file delete

Delete a tape config file

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

Description

The `storage tape config-file delete` command deletes the specified tape drive configuration file from all nodes that are currently part of the cluster.

Parameters

-filename <text> - Config File Filename (privilege: advanced)

This parameter specifies the name of the tape configuration file that will be deleted from all nodes that are currently part of the cluster.

Examples

The following example deletes the specified tape drive configuration files on every node that is currently part of the cluster:

```
cluster1::> storage tape config-file delete -filename XYZ_LTO-6.TCF
```

storage tape config-file get

Get a tape drive configuration file

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

Description

The `storage tape config-file get` command uploads a specified tape drive configuration file to each node that is currently part of the cluster.

Parameters

-url <text> - Config File URL

This parameter specifies the URL that provides the location of the package to be fetched. Standard URL schemes, including HTTP and TFTP, are accepted.

Examples

The following example uploads the specified tape drive configuration file to each node that is currently part of the cluster:

```
cluster1::> storage tape config-file get -url  
http://example.com/~tapeconfigfile/XYZ_LTO-6.TCF
```

storage tape config-file show

Display the list of tape drive configuration files on the given node

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

Description

The `storage tape config-file show` command lists the tape drive configuration files loaded onto each node 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.

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

Selects information about tape drive configuration files for the specified node.

[*-config-file* <text>] - Tape Config File

Selects information about the tape drive configuration file specified.

Examples

The following example lists the tape drive config files loaded onto each node in the cluster:

```
cluster1::> storage tape config-file show

Node: node1

Tape Config Files
-----
CERTANCE_LTO2_ULTRIUM.TCF
CERTANCE_LTO3_ULTRIUM.TCF
HP_LT09.TCF
HP_LTO2.TCF
HP_LTO3_ULTRIUM.TCF
HP_LTO4_ULTRIUM.TCF
HP_LTO5_ULTRIUM.TCF
HP_LTO6_ULTRIUM.TCF
IBM_3592.TCF
IBM_3592E05.TCF
IBM_5038_sdfkjl.TCF
IBM_LTO2_ULT3580.TCF
IBM_LTO2_ULTRIUM.TCF
```

storage tape library config show

Display connectivity to back-end storage tape libraries.

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

Description

This command displays information such as how the storage tape libraries connect to the cluster, LUN groups, number of LUNs, WWPN, and switch port information. Use this command to verify the cluster's storage tape library configuration or to assist in troubleshooting.

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.

| **[-switch]**

If you specify this parameter, switch port information is shown.

| **[-instance] }**

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

[-node {<nodename>|local}] - Controller Name

The name of the clustered node for which information is being displayed.

[-group <integer>] - LUN Group

A LUN group is a set of LUNs that shares the same path set.

[-target-wwpn <text>] - Library Target Ports

The World Wide Port Name of a storage tape library port.

[-initiator <text>] - Initiator

The host bus adapter that the clustered node uses to connect to storage tape libraries.

[-array-name <array name>] - Library Name

Name of the storage tape library that is connected to the clustered node.

[-target-side-switch-port <text>] - Target Side Switch Port

This identifies the switch port that connects to the tape library's target port.

[-initiator-side-switch-port <text>] - Initiator Side Switch Port

This identifies the switch port that connects to the node's initiator port.

[-lun-count <integer>] - Number of LUNS

This is a command-line switch (`-lun-count`) used to restrict what LUN groups are displayed in the output.

Examples

The following example displays the storage tape library configuration information.

```
cluster1::> storage tape library config show
      LUN   LUN
Node   Group Count      Library Name      Library Target
Port Initiator
-----
cluster1-01
              0       2              TAPE_LIB_1
50050763124b4d6f      3d

cluster1::>
```

storage tape library path show-by-initiator

Display a list of LUNs on the given tape library

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

Description

This command displays path information for every initiator port connected to a tape library. The output is similar to the storage library path show command but the output is listed by initiator.

Parameters

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

fields used to be used in this display

| [-instance] }

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

[-node {<nodename>|local}] - Controller name

The name of the clustered node for which information is being displayed.

[-initiator <text>] - Initiator Port

Initiator port that the clustered node uses.

[-target-wwpn <text>] - Target Port

Target World Wide Port Name. Port on the storage tape library that is being used.

[-initiator-side-switch-port <text>] - Initiator Side Switch Port

Switch port connected to the clustered node.

[-target-side-switch-port <text>] - Target Side Switch Port

Switch port connected to the tape library.

[-array-name <array name>] - Library Name

Name of the storage tape library that is connected to the cluster.

[-tpgn <integer>] - Target Port Group Number

TPGN refers to the target port group to which the target port belongs. A target port group is a set of target ports which share the same LUN access characteristics and failover behaviors.

[-port-speed <text>] - Port Speed

Port Speed of the specified port.

[-path-io-kbps <integer>] - Kbytes of I/O per second on Path (Rolling Average)

Rolling average of Kbytes of I/O per second on the library path.

[-path-iops <integer>] - Number of I/O per second on Path (Rolling Average)

Rolling average of I/O per second on the library path.

[*-initiator-io-kbps* <integer>] - Kbytes of I/O per second on Initiator (Rolling Average)

Rolling average of Kbytes of I/O per second on the initiator port.

[*-initiator-iops* <integer>] - Number of I/O per second on Initiator (Rolling Average)

>Rolling average of I/O per second on the initiator port.

[*-target-io-kbps* <integer>] - Kbytes of I/O per second to Target (Rolling Average)

Rolling average of Kbytes of I/O per second on the target port.

[*-target-iops* <integer>] - Number of I/O per second to Target (Rolling Average)

Rolling average of I/O per second on the target port.

Examples

The following example displays the path information by initiator for a storage tape library.

```
cluster1::> storage tape library path show-by-initiator
Node: cluster1-01

Initiator I/O      Initiator Side      Path I/O      Target
Side  Target I/O
Initiator      (KB/s)      Switch Port      (KB/s)      Switch
Port      (KB/s)      Target Port Library Name
-----
0b      0 sw_tape:6      0
sw_tape:0      0 510a09800000412d TAPE_LIB_1

sw_tape:1      0 510a09820000412d TAPE_LIB_1
3d      0 N/A      0
N/A      0 50050763124b4d6f TAPE_LIB_2
3 entries were displayed.
```

storage tape library path show

Display a list of Tape Libraries on the given path

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

Description

This command displays path information for a tape library and has the following parameters by default:

- Node name
- Initiator port
- Target port
- TPGN (Target Port Group Number)

- Port speeds
- Path I/O in Kbytes/sec
- IOPs

Parameters

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

fields used to be used in this display

| [-detail]

Using this option displays the following:

- Target IOPs
- Target LUNs
- Path IOPs
- Path errors
- Path quality
- Path LUNs
- Initiator IOPs
- Initiator LUNs

| [-instance] }

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

[-node {<nodename>|local}] - Controller name

The name of the clustered node for which information is being displayed.

[-array-name <array name>] - Library Name

Name of the storage tape library that is connected to the cluster.

[-target-wwpn <text>] - Target Port

Target World Wide Port Name. Port on the storage tape library that is being used.

[-initiator <text>] - Initiator Port

Initiator port that the clustered node uses.

[-initiator-side-switch-port <text>] - Initiator Side Switch Port

Switch port connected to the clustered node.

[-tpgn <integer>] - Target Port Group Number

TPGN refers to the target port group to which the target port belongs. A target port group is a set of target ports which share the same LUN access characteristics and failover behaviors.

[-port-speed <text>] - Port Speed

Port Speed of the specified port.

[-path-io-kbps <integer>] - Kbytes of I/O per second on Path(Rolling Average)

Rolling average of Kbytes of I/O per second on the library path.

[-path-iops <integer>] - Number of I/O per second on Path(Rolling Average)

Rolling average of I/O per second on the library path.

[-initiator-io-kbps <integer>] - Kbytes of I/O per second on Initiator(Rolling Average)

Rolling average of Kbytes of I/O per second on the initiator port.

[-initiator-iops <integer>] - Number of I/O per second on Initiator(Rolling Average)

>Rolling average of I/O per second on the initiator port.

[-target-io-kbps <integer>] - Kbytes of I/O per second to Target(Rolling Average)

Rolling average of Kbytes of I/O per second on the target port.

[-target-iops <integer>] - Number of I/O per second to Target(Rolling Average)

Rolling average of I/O per second on the target port.

[-target-side-switch-port <text>] - Target Side Switch Port

Switch port connected to the tape library.

[-path-link-errors <integer>] - Link Error count on path

Fibre Channel link error count.

[-path-quality <integer>] - Percentage of weighted error threshold

A number representing the threshold of errors that is allowed on the path. Path quality is a weighted error value. When the error weight of a path exceeds the threshold, I/O is routed to a different path.

[-path-lun-in-use-count <integer>] - Number of LUNs in the in-use state on this path

Number of LUNs on this path.

[-initiator-lun-in-use-count <integer>] - Number of LUNs in the in-use state on this initiator

Number of LUNs on this initiator.

[-target-lun-in-use-count <integer>] - Number of LUNs in the in-use state on this target

Number of LUNs on this target.

Examples

The following example displays the path information for a storage tape library

```
cluster1::> storage tape library path show
```

| Node | Initiator | Target | Port | TPGN | Speed |
|-------------|-----------|------------------|------|------|--------|
| (KB/s) | IOPs | | | | |
| cluster1-01 | 3d | 50050763124b4d6f | | 61 | 4 Gb/S |
| 0 | 0 | | | | |
| cluster1-01 | 0b | 510a09800000412d | | 35 | 4 Gb/S |
| 0 | 0 | | | | |
| cluster1-01 | 0b | 510a09820000412d | | 1 | 4 Gb/S |
| 0 | 0 | | | | |

3 entries were displayed.

storage tape load-balance modify

Modify the tape load balance configuration

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

Description

The `storage tape load-balance modify` command modifies the tape load balance setting for a specified node in the cluster.

Parameters

-node {<nodename>|local} - Node

This parameter specifies the node on which the tape load balance setting is to be modified.

[-is-enabled {true|false}] - Is Tape Load Balance Enabled

This parameter specifies whether tape load balancing is enabled on the node. The default setting is false.

Examples

The following example modifies the tape load balance setting on node1 in the cluster:

```
cluster1::> storage tape load-balance modify -node node1 -is-enabled true
```

storage tape load-balance show

Displays the tape load balance configuration

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

Description

The `storage tape load-balance show` command displays tape load balance settings for each node 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.

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

Selects information about tape load balancing for the specified node.

[-is-enabled {true|false}] - Is Tape Load Balance Enabled

Selects information about load balance configuration as specified by enabled or disabled setting.

Examples

The following example shows the load balance setting for each node in the cluster:

```
cluster1::> storage tape load-balance show
```

| Node | Enabled |
|-------|---------|
| ----- | ----- |
| node1 | false |
| node2 | false |

2 entries were displayed.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.