



storage disk commands

Command reference

NetApp
July 17, 2026

Table of Contents

- storage disk commands 1
 - storage disk assign 1
 - Description 1
 - Parameters 1
 - Examples 3
 - storage disk fail 5
 - Description 5
 - Parameters 5
 - Examples 5
 - Related Links 5
 - storage disk reassign 6
 - Description 6
 - Parameters 6
 - Examples 6
 - storage disk refresh-ownership 7
 - Description 7
 - Parameters 7
 - Examples 7
 - storage disk remove-reservation 7
 - Description 7
 - Parameters 7
 - Examples 7
 - storage disk remove 8
 - Description 8
 - Parameters 8
 - Examples 8
 - Related Links 8
 - storage disk removeowner 8
 - Description 8
 - Parameters 9
 - Examples 9
 - storage disk replace 9
 - Description 9
 - Parameters 9
 - Examples 10
 - storage disk sanitize-spare 10
 - Description 10
 - Parameters 10
 - Examples 11
 - storage disk set-foreign-lun 11
 - Description 11
 - Parameters 11
 - Examples 11

storage disk set-led	11
Description	12
Parameters	12
Examples	13
storage disk show	13
Description	14
Parameters	14
Examples	30
Related Links	37
storage disk unfail	37
Description	38
Parameters	38
Examples	38
storage disk updatefirmware	38
Description	38
Parameters	39
Examples	39
Related Links	39
storage disk zerospares	39
Description	39
Parameters	39
Examples	39
storage disk error show	40
Description	40
Parameters	40
Examples	41
storage disk firmware revert	41
Description	41
Parameters	42
Examples	43
Related Links	44
storage disk firmware show-update-status	44
Description	44
Parameters	44
Examples	45
storage disk firmware update	45
Description	45
Parameters	48
Examples	48
Related Links	49
storage disk option modify	49
Description	49
Parameters	49
Examples	50
storage disk option show	51

Description	51
Parameters	51
Examples	52

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 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 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 | SSD-ZNS | 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. 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 "0" or "1".

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

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

[-s, -sysid <nvrleid>] - 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 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_``` 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 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.

Examples

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

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

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 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 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, 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 sanitize-spare

Sanitize a spare disk (non-cryptographically)

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

Description

The `storage disk sanitize-spare` command sanitizes the specified spare disk (non-cryptographically).

Parameters

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

This parameter specifies the spare disk to be sanitized (non-cryptographically).

Examples

The following example sanitizes a spare disk named 1.1.16 (non-cryptographically):

```
cluster1::> storage disk sanitize-spare -disk 1.1.16
```

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 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 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 1.0.* to turn on for 5 minutes:

```
Cluster1::> storage disk set-led -disk 1.0.* -action on -duration 5
```

The following example causes the LEDs on disks 1.0.0 and 1.0.1 to turn on for 2 minutes:

```
Cluster1::> storage disk set-led -disk 1.0.0,1.0.1 -action on -duration 2
```

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 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
- Disk type
- Container type (aggregate, broken, foreign, labelmaint, maintenance, mediator, remote, shared, spare, unassigned, unknown, volume, or unsupported)
- 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

| [**-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

| [**-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)

| **[-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

| **[-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

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

Displays the following information about disks:

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

| **[-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, 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 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.
- 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 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 ONTAP-v models:

- Disk name
- Backing disk vendor

- Backing disk model
- Backing disk serial number
- Backing disk device id

[`-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 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 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 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 <nvramid>] - 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 <nvramid>] - 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 ONTAP.
- toobig = The array LUN exceeds the maximum array LUN size that ONTAP supports.
- toosmall = The array LUN is less than the minimum array LUN size that 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 ONTAP.
- netapp = A SAN front-end LUN from one ONTAP system that is presented as external storage to another 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 <nvr amid>] - 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.

[-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.

`[-odelist {<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 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 <nvr amid>] - 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 <nvr amid>] - 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.

[-shelf-uid-hex <text>] - Shelf UUID Hex

Selects information about disks that are located within a shelf with the specified hexadecimal 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 | SSD-ZNS | 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.

[`-hw-minimum-os <ClusterVersion>`] - Hardware Minimum Supported ONTAP Version (x.y.z)

Selects information about disks that have the specified hardware minimum supported ONTAP version in x.y.z format, for example "9.8.0" or "9.9.1".

[`-bridge-serialno <text>`] - Bridge Serial Number

Selects information about disks that are specified by the bridge serial number.

[`-mailbox-disk-type <text>`] - Mailbox Disk Type

The type of mediator mailbox disk.

- disk = disks from mediator node
- database = cloud database table

Examples

The following example displays information about all disks:

```
cluster1::> storage disk show
```

Disk	Usable Size	Shelf	Bay	Disk Type	Container Type	Container Name	Owner
1.1.1	10GB	1	1	SSD	spare	-	node1
1.1.4	78.59GB	1	4	SSD	spare	-	node1
1.1.12	10GB	1	12	SSD	spare	-	node1
1.2.12	10GB	2	12	SSD	broken	-	node1
1.3.7	78.59GB	3	7	SSD	aggregate	aggr0_u23	node1
1.1.6	78.59GB	1	6	SSD	broken	-	node1
1.2.10	78.59GB	2	10	SSD	aggregate	aggr0_u23	node1
1.4.9	78.59GB	4	9	SSD	aggregate	aggr0_u23	node1
1.1.0	10GB	1	0	SSD	aggregate	aggr0_u22	node2
1.4.1	10GB	4	1	SSD	aggregate	dp_degraded	node2
1.1.2	10GB	1	2	SSD	spare	-	node2
1.1.3	20GB	1	3	SSD	spare	-	node2
1.4.4	20GB	4	4	SSD	spare	-	node2
1.4.6	10GB	4	6	SSD	aggregate	dp_sdc	node2
1.1.5	268.0GB	1	5	SSD	maintenance	-	node2
1.3.0	10GB	3	0	SSD	aggregate	aggr0_u22	node2
1.4.11	10GB	4	11	SSD	spare	-	node2
1.4.13	20GB	4	13	SSD	broken	-	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
Mailbox Disk Type: disk
      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						
	data	2.11.2	2d	11	2	B	Pool0	SAS
15000	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						
	data	2.1.8	2a	1	8	B	Pool0	BSAS
7200	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						
	data	2.1.11	2a	1	11	B	Pool0	BSAS
7200	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    Pool10  BSAS
7200    9.77GB    9.91GB
Owner Node: node2
  Aggregate: aggr0
    Plex: plex0
      RAID Group: rg0

Usable Physical
      Position Disk                HA Shelf Bay Chan Pool  Type
RPM    Size      Size
-----
-----
          data      2.1.5                2a    1   5  B    Pool10  BSAS
7200    9.71GB   10.03GB
          dparity    2.1.2                2a    1   2  B    Pool10  BSAS
7200    9.71GB   10.03GB
          parity     2.1.4                2a    1   4  B    Pool10  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
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
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			Drawer							
Disk	Outage Reason	HA Shelf Bay	/Slot	Chan	Pool	Type				
RPM	Size	Size								

1.1.0	admin failed	0b	1	0	-/-	A	Pool0	FCAL		
10000	132.8GB	133.9GB								
1.2.6	admin removed	0b	2	6	-/-	A	Pool1	FCAL		
10000	132.8GB	134.2GB								

```
Original Owner: node2
```

```
Checksum Compatibility: block
```

Usable Physical			Drawer							
Disk	Outage Reason	HA Shelf Bay	/Slot	Chan	Pool	Type				
RPM	Size	Size								

1.1.0	admin failed	0a	1	0	-/-	B	Pool0	FCAL		
10000	132.8GB	133.9GB								
1.1.13	admin removed	0a	1	13	-/-	B	Pool0	FCAL		
10000	132.8GB	133.9GB								

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	Pool10	FCAL	10000
132.8GB 133.9GB							

1.2.11	admin testing 0b	2	11	A	Pool11	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	Pool11	FCAL	10000
132.8GB 133.9GB							

1.2.13	admin testing 0a	2	13	B	Pool11	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	cluster-2-01	cluster-2-01	4087518786	
4087518786					
VMw-1.14	Container	cluster-2-01	cluster-2-01	4087518786	
4087518786					
VMw-1.15	Container	cluster-2-01	cluster-2-01	4087518786	
4087518786					
	Root	cluster-2-01	cluster-2-01	4087518786	
4087518786					
	Data	cluster-2-01	cluster-2-01	4087518786	
4087518786					
VMw-1.16	Container	cluster-2-01	cluster-2-01	4087518786	
4087518786					
	Root	cluster-2-01	cluster-2-01	4087518786	
4087518786					
	Data1	cluster-2-01	cluster-2-01	4087518786	
4087518786					
	Data2	cluster-2-01	cluster-2-01	4087518786	
4087518786					
VMw-1.17	Container	cluster-2-01	cluster-2-01	4087518786	
4087518786					
VMw-1.18	Container	cluster-2-01	cluster-2-01	4087518786	
4087518786					
	Root	-	-	-	-
	Data	cluster-2-01	cluster-2-01	4087518786	
4087518786					
VMw-1.19	Container	cluster-2-01	cluster-2-01	4087518786	
4087518786					
	Root	cluster-2-01	cluster-2-01	4087518786	
4087518786					
	Data1	-	-	-	-
	Data2	cluster-2-01	cluster-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:

```
cluster1::*> storage disk unfail -disk 1.1.16
```

Warning: Failed disk "1.1.16" may have aggregate labels and file system data present. In that case, this command will attempt to bring this disk back into the aggregate with which this disk had formerly been associated and preserve file system data. 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 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`
```

[-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 ONTAP.
- toobig = The array LUN exceeds the maximum array LUN size that ONTAP supports.
- toosmall = The array LUN is less than the minimum array LUN size that 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 ONTAP.
- netapp = A SAN front-end LUN from one ONTAP system that is presented as external storage to another 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 ] (privilege: advanced) }`

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
```

	Update	Number	Average	Total
Node	State	Waiting	Duration	Completion
Update		Download	/Disk (Sec)	Est. (Min)
				Unable to
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. 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 will not occur until it is turned back on.

[-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 nodel -autoassign-policy bay
cluster1::> storage disk option show
```

Node	BKg. FW. Upd.	Auto Copy	Auto Assign	Auto Assign Policy
nodel	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.
```

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.