



# **storage failover commands**

## **ONTAP 9.3 commands**

NetApp  
September 20, 2022

# Table of Contents

- storage failover commands ..... 1
  - storage failover giveback ..... 1
  - storage failover modify ..... 2
  - storage failover show-giveback ..... 5
  - storage failover show-takeover ..... 7
  - storage failover show ..... 11
  - storage failover takeover ..... 28
  - storage failover hwassist show ..... 30
  - storage failover hwassist test ..... 32
  - storage failover hwassist stats clear ..... 33
  - storage failover hwassist stats show ..... 33
  - storage failover internal-options show ..... 35
  - storage failover mailbox-disk show ..... 39
  - storage failover progress-table show ..... 41

# storage failover commands

## storage failover giveback

Return failed-over storage to its home node

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

### Description

The `storage failover giveback` command returns storage that has failed over to a node's partner back to the home node. This operation fails if other resource-intensive operations (for instance, system dumps) are running and make the giveback operation potentially dangerous or disruptive. Some options are available only at the advanced privilege level and higher. Run the [storage failover show-giveback](#) command to check the status of giveback operations.

NOTE:

- If the system ID of the partner has changed while the node is in takeover mode, the `storage failover giveback` command updates the ownership of the partner's disks to the new system ID while giving back.
- If the giveback operation fails due to the operation being vetoed by a subsystem, check the syslog or EMS output for a subsystem-specific reason for the abort. The corrective action is subsystem-specific and is detailed in the corrective action portion of the message. Follow the corrective action specified by the subsystem and then reissue the `storage failover giveback` command. If you cannot perform the corrective action, then use the `override-vetoes` option in the `storage failover giveback` command to force the giveback.
- If the giveback operation fails because the node cannot communicate with its partner, check the EMS output for the corrective action. Follow the corrective action and then reissue the `storage failover giveback` command. If you cannot perform the corrective action, then use the `-require-partner -waiting false` option in the `storage failover giveback` command to force the giveback. This parameter is available only at the advanced privilege level and higher.
- If the node does not receive notification that the partner has brought online the given-back aggregate and its volumes, the [storage failover show-giveback](#) command displays the giveback status for the aggregate as failed. A possible reason for this failure is that the partner is overloaded and slow in bringing the aggregate online. Run the [storage aggregate show](#) command to verify that the aggregate and its volumes are online on the partner node. The node will not attempt the giveback operation for remaining aggregates. To force the giveback, use the `-require-partner-waiting false` option in the `storage failover giveback` command. This parameter is available only at the advanced privilege level and higher.

### Parameters

**{ -ofnode {<nodename>|local} - Node to which Control is Givenback**

Specifies the node whose storage is currently taken over by its partner and will be given back by the giveback operation.

**| -fromnode {<nodename>|local} - Node Initiating Giveback }**

Specifies the node that currently holds the storage that is to be returned to the partner node.

### **[`-require-partner-waiting {true|false}`] - Require Partner in Waiting**

If this optional parameter is used and set to false, the storage is given back regardless of whether the partner node is available to take back the storage or not. If this parameter is used and set to true, the storage will not be given back if the partner node is not available to take back the storage. If this parameter is not used, the behavior defaults to the setting of the `-check-partner` option set with the [storage failover modify](#) command.

### **[`-override-vetoes <true>`] - Override All Vetoes**

If this optional parameter is used, the system overrides veto votes during a giveback operation. If this parameter is not used, the system does not proceed with a giveback if it is vetoed. This parameter, if used, can only be set to true.

### **[`-only-cfo-aggregates <true>`] - Giveback Only CFO Aggregates**

If this optional parameter is used, giveback of only the CFO aggregates (root aggregate and CFO style data aggregates) will be attempted. If this parameter is not used, giveback of all the aggregates (CFO and SFO aggregates) will be attempted. This parameter, if used, can only be set to true.

## **Examples**

The following example gives back storage that is currently held by a node named node1. The partner must be available for the giveback operation to occur.

```
node::> storage failover giveback -fromnode node1
```

The following example gives back only the CFO aggregates to a node named node2 (the aggregates are currently held by a node named node1). The partner must be available for the giveback operation to occur, and the veto-giveback process can be overridden.

```
node::> storage failover giveback -ofnode node2  
-override-vetoes true -only-cfo-aggregates true
```

## **Related Links**

- [storage failover show-giveback](#)
- [storage aggregate show](#)
- [storage failover modify](#)

# **storage failover modify**

Modify storage failover attributes

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

## **Description**

The `storage failover modify` command changes the storage-failover options for a node. Some options are available only at the advanced privilege level and higher.

## Parameters

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

This specifies the node whose storage-failover options are to be modified.

### **{ [-enabled {true|false}] - Takeover Enabled**

This optionally specifies whether storage failover is enabled. The default setting is `true`.

### **[-mode {ha|non\_ha}] - HA Mode }**

This specifies whether the node is set up in high-availability mode or stand-alone mode. If the node is a member of a high-availability configuration, set the value to `ha`. If the node is stand-alone, set the value to `non_ha`. Before setting the HA mode, you must complete the platform dependent steps to set up the system in a stand-alone or HA configuration as shown in the documentation for your platform.

### **[-auto-giveback {true|false}] - Auto Giveback Enabled**

This optionally specifies whether automatic giveback operations are enabled. An automatic giveback operation is invoked when one node of a failover pair is in takeover mode and the failed node is repaired and restarts. When the repaired node boots, the node in takeover mode detects this and initiates a giveback operation. The default setting is `false`, except for two-node clusters where the default setting is `true`.

### **[-detection-time <integer>] - Takeover Detection Time (secs)**

This optionally specifies the amount of time, in seconds, that a node remains unresponsive before its partner initiates a takeover operation. Possible values range from 10 to 180 seconds. The default setting is 15 seconds.

### **[-onfailure {true|false}] - Takeover on Failure Enabled**

This optionally specifies whether the node automatically takes over for its partner node if the partner node fails. The default setting is `true`. This parameter is available only at the advanced privilege level and higher.

### **[-onpanic {true|false}] - Takeover on Panic Enabled**

This optionally specifies whether the node automatically takes over for its partner node if the partner node panics. The default setting is `true`. Changing this parameter on one node automatically makes the same change on its partner node.

### **[-onshort-uptime {true|false}] - Takeover on Short Uptime Enabled**

This optionally specifies whether the node takes over for its partner node if the partner node fails within 60 seconds of starting up; the time period is modifiable by using the `-short-uptime` parameter. The default setting is `true`. This parameter is available only at the advanced privilege level and higher.

### **[-short-uptime <integer>] - Short Uptime (secs)**

This optionally specifies the time period used by the `-onshort-uptime` parameter. The default setting is 60 seconds. This parameter is available only at the advanced privilege level and higher.

### **[-attempts <integer>] - Number of Giveback Attempts**

This optionally specifies the number of times the node attempts an automatic giveback operation within 60 minutes; the time period is modifiable by using the `-attempts-time` parameter. The default setting is 2 attempts. This parameter is available only at the advanced privilege level and higher.

**[-attempts-time <integer>] - Giveback Attempts Period (minutes)**

This optionally specifies the time period used by the `-attempts` parameter. The default setting is 60 minutes. This parameter is available only at the advanced privilege level and higher.

**[-propagate {true|false}] - Propagate Status via Mailbox**

This optionally specifies whether storage-failover status is communicated via mailbox disks. The default setting is `true`. This parameter is available only at the advanced privilege level and higher.

**[-read-interval <integer>] - Node Status Read Interval (secs)**

This optionally specifies, in seconds, how frequently the node reads its partner node's status from the mailbox disks. The default setting is 5 seconds. This parameter is available only at the advanced privilege level and higher.

**[-write-interval <integer>] - Node Status Write Interval (secs)**

This optionally specifies, in seconds, how frequently the node writes its status to the mailbox disks. The default setting is 5 seconds. This parameter is available only at the advanced privilege level and higher.

**[-onreboot {true|false}] - Takeover on Reboot Enabled**

This optionally specifies whether the node automatically takes over for its partner if the partner reboots. The default setting is `true`. Takeover can occur if the partner exceeds the expected time to reboot even when this option is set to `false`. The expected time to reboot is different for different platforms. The minimum expected time to reboot is 180 seconds. The `-inhibit-takeover` option of the [system node reboot](#) command overrides this option: if a node is rebooted with `-inhibit-takeover` set to `true` then takeover does not occur, even if the `takeover on reboot` option is `true`. If a node does takeover due to the partner rebooting, then it will automatically giveback after the reboot, even if the `-auto-giveback` option is set to `false`. This is non-persistent behavior: if the node does takeover due to partner reboot and then itself reboots (prior to giveback) then it will not automatically giveback if the `-auto-giveback` option is set to `false`.

**[-delay-seconds <integer>] - Delay Before Auto Giveback (secs)**

This optionally specifies the minimum time that a node will stay in takeover state prior to performing an automatic giveback. If the taken over node recovers quickly (for example, if the takeover was due to a reboot), by delaying the giveback for a few minutes the outage during the takeover and giveback can be reduced to two short outages instead of one longer one. The allowed range is 0 to 600, inclusive. The default setting is 600 seconds. This option affects all types of auto-giveback. This parameter is available only at the advanced privilege level and higher.



This delay does not affect manual giveback.

**[-hwassist {true|false}] - Hardware Assist Enabled**

This optionally specifies whether the hardware assist feature is enabled. If set to `true` this feature helps in fast takeover detection times in certain cases.

**[-hwassist-partner-ip <IP Address>] - Partner's Hwassist IP**

This optionally specifies the IP address on which the partner node receives hardware assist alerts. For the hardware assist feature to be active, the value of this option should be equal to partner's node management IP address.

**[-hwassist-partner-port <integer>] - Partner's Hwassist Port**

This optionally specifies the port number on which partner node listens to hardware assist alerts. It is

recommended to have this value to be between 4000-4500. The default value is 4444.

**`[-hwassist-health-check-interval <integer>]` - Hwassist Health Check Interval (secs)**

This optionally specifies, in seconds, how frequently the hardware assist hardware on a node sends a heartbeat to its partner. The default value is 180.

**`[-hwassist-retry-count <integer>]` - Hwassist Retry Count**

This optionally specifies the number of times we repeat sending an hardware assist alert. The default value is 2.

**`[-auto-giveback-after-panic {true|false}]` - Auto Giveback After Takeover On Panic**

This optionally specifies whether a node should attempt automatic giveback operations if takeover was because of a disruption in the partner's operation. An automatic giveback operation is invoked when one node of a failover pair is in takeover mode and the failed node is repaired and restarts. When the repaired node boots, the node in takeover mode detects this and initiates a giveback operation automatically. The default setting is `true`.

**`[-aggregate-migration-timeout <integer>]` - Aggregate Migration Timeout (secs)**

This optionally specifies the amount of time, in seconds, the source node has to wait for the destination node to complete the aggregate migration before declaring the migration as failed. The default setting is 120 seconds.

## Examples

The following example enables the storage-failover service on a node named node0:

```
node::> storage failover modify -node node0 -enabled true
```

The following examples enable storage-failover takeover on a short uptime of 30 seconds on a node named node0:

```
node::*> storage failover modify -node node0 -onshort-uptime true -short-uptime 30
```

## Related Links

- [system node reboot](#)

## storage failover show-giveback

Display giveback status

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

## Description

The `storage failover show-giveback` command displays information about the giveback status of high-availability (HA) partner aggregates. The command displays the following information when no parameters are

specified:

- Node name
- Partner aggregate name
- Giveback Status

You can specify additional parameters to display only the information that matches those parameters. For example, to display information only about a particular aggregate, run the command with the ``-aggregate aggregate_name `` parameter.

## Parameters

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

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

**| [-instance ] }**

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

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

If this parameter is used, the command displays information about the giveback status of the aggregates belonging to the HA partner of the specified node.

**[-aggregate <text>] - Aggregate**

If this parameter is used, the command displays information about the giveback status of the specified aggregate.

**[-giveback-status <text>,...] - Aggregates Giveback State**

If this parameter is used, the command displays information about the aggregates with the specified giveback status.

**[-destination <text>] - Destination for Giveback**

If this parameter is used, the command displays information about the giveback status of the aggregates whose destination after the giveback is the specified node.

## Examples

The following example displays information about giveback status on all nodes:



```
node::> storage failover show-giveback
          Partner
Node      Aggregate      Giveback Status
-----
node0     -                No aggregates to give back
node1     -                No aggregates to give back
node2     -                No aggregates to give back
node3     -                No aggregates to give back
4 entries were displayed.
```

## storage failover show-takeover

### Display takeover status

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

### Description

The `storage failover show-takeover` command displays information about the takeover status of nodes in a cluster. The command also displays the takeover status of aggregates being taken over. During each phase of takeover, the takeover node and the node being taken over display their takeover status and the status of the aggregates being taken over. The command displays the following information when no parameters are specified:

- Node name
- Node takeover status - This contains a descriptive information about the phase of takeover.
- Aggregate
- Aggregate takeover status - This contains the following information:
  - Takeover status of the aggregate, such as "Done", "Failed", "In progress" and "Not attempted yet".
  - Reason for an aggregate takeover failure.
  - Corrective action, in case of an aggregate takeover failure.

You can specify additional parameters to display only the information that matches those parameters. For example, to display information only about a particular node, run the command with the `-node` ``_node_name_` ``` parameter.

## Parameters

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

If this parameter is specified, the command displays the specified fields for all nodes, in column style output.

**| [-instance ] }**

If this parameter is specified, the command displays the same detailed information as for the `-node` parameter, but for all nodes.

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

If this parameter is specified, the command displays information about the takeover status of the specified node, and the takeover status of the aggregates being taken over.

**[-node-takeover-status <text>] - Node's Takeover Status**

If this parameter is specified, the command displays information about the takeover status of the nodes with the specified `node-takeover-status`. The command also displays the takeover status of the aggregates belonging to the node being taken over.

**[-aggregate <text>] - Aggregate Being Taken Over**

If this parameter is specified, the command displays information about the takeover status of the specified aggregate, and the takeover status of the nodes containing the specified aggregate.

**[-aggregate-takeover-status <text>] - Aggregate's Takeover Status**

If this parameter is specified, the command displays information about the takeover status of the aggregates with the specified `aggregate takeover status`, and the takeover status of the nodes containing those aggregates.

## Examples

The following example shows the takeover status of two nodes, `nodeA` and `nodeB`, in an High Availability (HA) pair, when both are in normal mode; neither node has taken over its HA partner. In this case, there is no takeover status for the aggregates.

```
cluster1::> storage failover show-takeover
Node          Node Status          Aggregate          Takeover Status
-----
nodeA         Takeover not
               attempted.
               -
nodeB         Takeover not
               attempted.
               -
```

The following example shows the takeover status of two nodes, `nodeA` and `nodeB`, in an HA pair, when `nodeA` is in the SFO phase of an optimized takeover of `nodeB`. In this case, `nodeA` does not have information about the takeover status of `nodeB`'s aggregates.

```

cluster1::> storage failover show-takeover
Node          Node Status          Aggregate          Takeover Status
-----
nodeA         Optimized takeover
              of partner in
              progress. Partner,
              ("nodeB"), is
              relocating its SFO
              aggregates. Run the
              command "storage
              failover
              show-takeover -node
              nodeB" to display the
              relocation status of
              the partner.
              -
nodeB         Being taken over.
              aggr1          In progress, Module: backup.
              aggr2          Not attempted yet
              CFO aggregates Not attempted yet.

```

The following example shows the takeover status of two nodes, nodeA and nodeB, in an HA pair, when nodeA has completed the SFO phase of an optimized takeover of nodeB (but has not completed the CFO phase of the optimized takeover). In this case, nodeA has information about the takeover status of nodeB's aggregates.

```

cluster1::> storage failover show-takeover
ode          Node Status          Aggregate          Takeover Status
-----
odeA         Partner has
              relocated its
              aggregates. Takeover
              in progress.
              aggr1          Done
              aggr2          Done
              CFO aggregates In progress.
odeB         Relocated aggregates
              to partner. Waiting
              for partner to
              takeover.
              aggr1          Done
              aggr2          Done
              CFO aggregates Not attempted yet.

```

The following example shows the takeover status of two nodes, nodeA and nodeB, in an HA pair, when nodeA has completed the SFO and CFO phases of an optimized takeover of nodeB. In this case, nodeA has information about the takeover status of nodeB's aggregates. Since nodeB is not operational, an Remote Procedure Call(RPC) error is indicated in the command output.

```
cluster1::> storage failover show-takeover
ode      Node Status      Aggregate      Takeover Status
-----
odeA     Partner has
         relocated its
         aggregates. In
         takeover.
         aggr1           Done
         aggr2           Done
         CFO aggregates Done.
Warning: Unable to list entries on node nodeB. RPC: Port mapper failure -
RPC:
         Timed out
```

The following example shows the takeover status of two nodes, nodeA and nodeB, in an HA pair, when nodeA has aborted the SFO phase of an optimized takeover of nodeB. In this case, nodeA does not have information about the takeover status of nodeB's aggregates.

```

cluster1::> storage failover show-takeover
ode      Node Status      Aggregate      Takeover Status
-----
odeA     Optimized takeover
of partner aborted.
Run the command
"storage failover
show-takeover -node
nodeB" to display the
relocation status of
the partner.
-
odeB     Optimized takeover
by partner aborted.
aggr1    Failed: Destination node did
not online the aggregate on
time. To takeover the
remaining aggregates, run the
"storage failover takeover
-ofnode nodeB
-bypass-optimization true"
command. To giveback the
relocated aggregates, run the
"storage failover giveback
-ofnode nodeB" command.
aggr2    Not attempted yet
CFO aggregates Not attempted yet.

```

## storage failover show

Display storage failover status

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

### Description

The `storage failover show` command displays information about storage-failover configurations. By default, the command displays the following information:

- Node name.
- Partner node name.
- Whether storage failover is possible.
- The current state of storage failover. If the takeover is disabled the appropriate reason would be displayed.

To display detailed information about storage failover on a specific node, run the command with the `-node` parameter. The detailed view adds the following information:

- Node NVRAM ID.
- Partner NVRAM ID.
- Whether storage failover is enabled.
- Whether the storage-failover interconnect is available.
- Status of individual storage-failover interconnect links.
- Type and vendor of the storage-failover interconnect.
- Partner State
- Status codes from the takeover-by-partner process. Possible values include:
  - NVRAM\_DOWN
  - OPERATOR\_DISABLE\_NVRAM
  - PARTNER\_RESET
  - FM\_TAKEOVER
  - NVRAM\_MISMATCH
  - OPERATOR\_DENY
  - CLUSTER\_DISABLE
  - VERSION
  - SHELF\_HOT
  - REVERT\_IN\_PROGRESS
  - HALT\_NOTKOVER
  - TAKEOVER\_ON\_PANIC
- Reasons why takeover is not possible, if applicable. Possible values include:
  - NOT\_INIT
  - DISABLED
  - DEGRADED
  - MBX\_UNKNOWN
  - FM\_VERSION
  - PARTNER\_DISABLED
  - OPERATOR\_DENY
  - NVRAM\_MISMATCH
  - VERSION
  - IC\_ERROR
  - BOOTING
  - SHELF\_HOT
  - PARTNER\_REVERT\_IN\_PROGRESS
  - LOCAL\_REVERT\_IN\_PROGRESS

- PARTNER\_TAKEOVER
- LOCAL\_TAKEOVER
- HALT\_NOTKOVER
- LOG\_UNSYNC
- UNKNOWN
- WAITING\_FOR\_PARTNER
- LOW\_MEMORY
- HALTING
- MBX\_UNCERTAIN
- NO\_AUTO\_TKOVER
- Time until takeover, in seconds.
- Time until auto giveback, in seconds.
- Delay for auto giveback, in seconds.
- List of local mailbox disks.
- List of partner mailbox disks.
- Whether operator-initiated planned takeover will be optimized for performance by relocating SFO (non-root) aggregates serially to the partner prior to takeover.

You can specify additional parameters to select the displayed information. For example, to display information only about storage-failover configurations whose interconnect is down, run the command with `-interconnect-up false``.

## Parameters

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

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

**| [-options ]**

Displays the following information:

- Node name
- Whether automatic giveback operations are enabled
- Whether long-running operations are terminated when an automatic giveback operation is initiated
- Whether the node checks its partner's readiness before initiating a giveback operation
- The time, in seconds, that the node remains unresponsive before its partner initiates a takeover operation
- Whether the node automatically takes over for its partner if the partner fails
- Whether the node automatically takes over for its partner if the partner panics
- Whether the node automatically takes over for its partner if the partner reboots
- whether Hardware Assisted Takeover is enabled
- Ip address on which the partner node listens to the Hardware Assist alerts

- Port number on which the partner node listens to the Hardware Assist alerts
- Whether operator-initiated planned takeover will be optimized for performance by relocating SFO (non-root) aggregates serially to the partner prior to takeover

If this parameter is specified when the privilege level is set to advanced or higher, the command displays the information in the previous list and the following additional information:

- Whether the node takes over for its partner if its partner fails after a period of time, which is listed in the following field
- The number of seconds before the node takes over for its partner
- The number of times the node attempts an automatic giveback operation within a period of time
- The number of minutes in which the automatic giveback attempts can occur
- Whether storage-failover status is communicated via mailbox disks
- The interval at which the node reads its partner node's status from the mailbox disks
- The interval at which the node writes its status to the mailbox disks
- '
- The interval at which Hardware assist h/w sends a heartbeat
- The number of times the Hardware assist alert is sent

#### | [**-takeover-status**] |

Displays the following information:

- Node name
- Partner name
- Takeover enabled
- Takeover possible
- Interconnect up
- State
- Node NVRAM ID
- Partner NVRAM ID
- Reason Takeover Not Possible By Partner
- Reason Takeover Not Possible
- Time Until Takeover

#### | [**-advanced**] |

Displays the following information:

- Node name
- Whether kill messages are issued during a takeover operation
- Whether the node controls its partner's storage aggregates
- The time when firmware notification was received
- The time when booting notification was received



- The time at which the last takeover or giveback operation occurred, in microseconds
- The number of times the failover log was unsynchronized

#### | [**-iotime**]

Displays the following information:

- Node name
- Primary normal I/O time
- Primary transition I/O time
- Backup normal I/O time
- Backup transition I/O time

#### | [**-mailbox-status**]

Displays the following information:

- Node name
- Primary mailbox status
- Backup mailbox status

#### | [**-more-options**]

Displays the following information:

- Node name
- Whether takeover on short uptime is enabled
- Short uptime, in seconds
- Number of giveback attempts
- Interval of giveback attempts, in minutes
- Whether the primary mailbox is online
- Mailbox status read interval, in seconds
- Mailbox status write interval, in seconds

#### | [**-progress**]

Displays the following information:

- Node name
- Maximum resource-table index number
- Current resource-table index number
- Current resource-table entry

#### | [**-timeout**]

Displays the following information:

- Node name
- Fast timeout

- Slow timeout
- Mailbox timeout
- Connection timeout
- Operator timeout
- Firmware timeout
- Dump-core timeout
- Booting timeout
- Reboot timeout

**| [-transit ]**

Displays the following information:

- Node name
- Transit Timer Enabled
- Transit Timeout

**| [-instance ] }**

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

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

Selects the nodes whose name matches this parameter value.

**[-partner-name <text>] - Partner Name**

Selects the nodes that have the specified partner-name setting.

**[-nvramid <integer>] - Node NVRAM ID**

Selects the nodes that have the specified NVRAM ID setting.

**[-partner-nvramid <integer>] - Partner NVRAM ID**

Selects the nodes that have the specified partner NVRAM ID setting.

**[-enabled {true|false}] - Takeover Enabled**

Selects the nodes that have the specified takeover-enabled setting.

**[-mode {ha|non\_ha}] - HA Mode**

Selects the nodes that have the specified HA-mode setting. If the value is set to `ha` then the node is a member of a storage-failover configuration. If it is set to `non-ha` then it is in a stand alone configuration.

**[-possible {true|false}] - Takeover Possible**

Selects the nodes that have the specified failover-possible setting.

**[-reason <text>, ...] - Reason Takeover not Possible**

Selects the nodes that have the specified reason-not-possible setting. Possible values include:

- NOT\_INIT
- DISABLED

- DEGRADED
- MBX\_UNKNOWN
- FM\_VERSION
- PARTNER\_DISABLED
- OPERATOR\_DENY
- NVRAM\_MISMATCH
- VERSION
- IC\_ERROR
- BOOTING
- SHELF\_HOT
- PARTNER\_REVERT\_IN\_PROGRESS
- LOCAL\_REVERT\_IN\_PROGRESS
- PARTNER\_TAKEOVER
- LOCAL\_TAKEOVER
- HALT\_NOTKOVER
- LOG\_UNSYNC
- UNKNOWN
- WAITING\_FOR\_PARTNER
- LOW\_MEMORY
- HALTING
- MBX\_UNCERTAIN
- NO\_AUTO\_TKOVER

**[-interconnect-up {true|false}] - Interconnect Up**

Selects the nodes that have the specified interconnect-up setting.

**[-interconnect-links <text>] - Interconnect Links**

Selects the nodes that have the specified interconnect-links setting.

**[-interconnect-type <text>] - Interconnect Type**

Selects the nodes that have the specified interconnect-type setting.

**[-state-description <text>] - State Description**

Selects the nodes that have the specified state-description setting.

**[-partner-state <text>] - Partner State**

Selects the nodes that have the specified partner-state setting. Possible values include:

- OPERATOR COMPLETED
- DEBUGGUER COMPLETED
- PROGRESS COUNTER

- I/O ERROR
- BAD CHECKSUM
- RESERVED
- UNKNOWN
- INITIALIZING
- IN POWER-ON SELF TEST
- BOOTING
- BOOT FAILED
- WAITING
- KERNEL LOADED
- UP
- IN DEBUGGER
- WAITING FOR OPERATOR INPUT
- DUMPING CORE
- HALTED
- REBOOTING
- WAITING FOR GIVEBACK (DISK RESERVATIONS)
- WAITING FOR GIVEBACK (HA MAILBOXES)
- DUMPING SPARECORE
- MULTI-DISK PANIC
- IN TAKEOVER

**[-time-until-takeover <integer>] - Time Until Takeover**

Selects the nodes that have the specified time-until-takeover setting.

**[-partner-reason <text>,...] - Reason Takeover not Possible by Partner**

Selects the nodes that have the specified partner-reason text setting.

**[-killpackets {true|false}] - Issue Kill Packets**

Selects the nodes that have the specified kill packets setting.

**[-partner-aggregates {true|false}] - Control Partner Aggregates**

Selects the nodes that have the specified partner aggregates setting.

**[-current-index <integer>] - Current Progress Index**

Selects the nodes that have the specified current-progress index setting.

**[-current-entry <text>] - Current Progress Entry**

Selects the nodes that have the specified current-progress entry setting.

**[-maximum-index <integer>] - Maximum Progress Index**

Selects the nodes that have the specified maximum-progress index setting.

**[-pmbx-status <text>,...] - Primary Mailbox Status**

Selects the nodes that have the specified primary mailbox status setting. Possible values include:

- MBX\_STATUS\_NODISKS
- MBX\_STATUS\_UNCERTAIN
- MBX\_STATUS\_STALE
- MBX\_STATUS\_CONFLICTED
- MBX\_STATUS\_OLD\_VERSION
- MBX\_STATUS\_NOT\_FOUND
- MBX\_STATUS\_WRONG\_STATE
- MBX\_STATUS\_BACKUP

**[-bmbx-status <text>,...] - Backup Mailbox Status**

Selects the nodes that have the specified backup-mailbox status setting. See the description of the `-pmbx-status` parameter for a list of possible values.

**[-major-seq-num-local <integer>] - Local Major Sequence Number**

Selects the nodes that have the specified mailbox heartbeat major sequence number on the local node.

**[-minor-seq-num-local <integer>] - Local Minor Sequence Number**

Selects the nodes that have the specified mailbox heartbeat minor sequence number on the local node.

**[-major-seq-num-partner <integer>] - Partner Major Sequence Number**

Selects the nodes that have the specified mailbox heartbeat major sequence number on the partner node.

**[-minor-seq-num-partner <integer>] - Partner Minor Sequence Number**

Selects the nodes that have the specified mailbox heartbeat minor sequence number on the partner node.

**[-local-mbx-node-status <Mailbox Status>] - Local Mailbox Node Status**

Selects the nodes that have the specified local mailbox node status. Possible values include:

- MBX\_UNKNOWN - Local node is up, mailbox uninitialized
- MBX\_TAKEOVER\_DISABLED - Local node is up but takeover is disallowed
- MBX\_TAKEOVER\_ENABLED - Local node is up and takeover is allowed
- MBX\_TAKEOVER\_ACTIVE - Partner node has taken over
- MBX\_GIVEBACK\_DONE - Giveback completed, but local node has not yet restarted

**[-mbx-abs-time-local <integer>] - Local Mailbox Absolute Time**

Selects the nodes that have the specified local mailbox channel absolute time. This time is measured in msec since 1/1/1970 (epoch).

**[-mbx-sk-time-local <integer>] - Local Mailbox Kernel Time**

Selects the nodes that have the specified local mailbox channel Kernel Time.

**[-mbx-sk-cycles-local <integer>] - Local Mailbox CPU Cycles**

Selects the nodes that have the specified local mailbox channel CPU Cycle count.

**[-ic-abs-time-local <integer>] - Local IC Absolute Time**

Selects the nodes that have the specified local Interconnect channel absolute time. This time is measured in msec since 1/1/1970 (epoch).

**[-ic-sk-time-local <integer>] - Local IC Kernel Time**

Selects the nodes that have the specified local Interconnect channel Kernel Time.

**[-ic-sk-cycles-local <integer>] - Local IC CPU Cycles**

Selects the nodes that have the specified local Interconnect channel CPU Cycle count.

**[-partner-mbx-node-status <Mailbox Status>] - Partner Mailbox Node Status**

Selects the nodes that have the specified partner mailbox node status. Possible values include:

- MBX\_UNKNOWN
- MBX\_TAKEOVER\_DISABLED
- MBX\_TAKEOVER\_ENABLED
- MBX\_TAKEOVER\_ACTIVE
- MBX\_GIVEBACK\_DONE

**[-mbx-abs-time-partner <integer>] - Partner Mailbox Absolute Time**

Selects the nodes that have the specified partner mailbox channel absolute time. This time is measured in msec since 1/1/1970 (epoch).

**[-mbx-sk-time-partner <integer>] - Partner Mailbox Kernel Time**

Selects the nodes that have the specified partner mailbox channel Kernel Time.

**[-mbx-sk-cycles-partner <integer>] - Partner Mailbox CPU Cycles**

Selects the nodes that have the specified partner mailbox channel CPU Cycle count.

**[-mbx-major-seq-num-partner <integer>] - Partner Mailbox Major Sequence Number**

Selects the nodes that have the specified partner mailbox channel major sequence number.

**[-mbx-minor-seq-num-partner <integer>] - Partner Mailbox Minor Sequence Number**

Selects the nodes that have the specified partner mailbox channel minor sequence number.

**[-ic-abs-time-partner <integer>] - Partner IC Absolute Time**

Selects the nodes that have the specified partner Interconnect channel absolute time. This time is measured in msec since 1/1/1970 (epoch).

**[-ic-sk-time-partner <integer>] - Partner IC Kernel Time**

Selects the nodes that have the specified partner Interconnect channel Kernel Time.

**[-ic-sk-cycles-partner <integer>] - Partner IC CPU Cycles**

Selects the nodes that have the specified partner Interconnect channel CPU Cycle count.

**[-ic-major-seq-num-partner <integer>] - Partner IC Major Sequence Number**

Selects the nodes that have the specified partner Interconnect channel major sequence number.

**[-ic-minor-seq-num-partner <integer>] - Partner IC Minor Sequence Number**

Selects the nodes that have the specified partner Interconnect channel minor sequence number.

**[-local-takeover-info <text>] - Local Takeover Info**

Selects the nodes that have the specified local node takeover information. This includes the type of negotiated failover request, or if takeover is not possible, the reason why takeover is disabled. Possible values include:

- NOTKOVER\_NVRAM\_DOWN - NVRAM mirror is down
- NOTKOVER\_OPERATOR\_DISABLE\_NVRAM - Operator disabled
- NOTKOVER\_PARTNER\_RESET - A link reset is in progress
- NOTKOVER\_FM\_TAKEOVER - The failover monitor has declared takeover
- NOTKOVER\_NVRAM\_MISMATCH - NVRAM sizes mismatch
- NOTKOVER\_OPERATOR\_DENY - Operator denies takeover
- NOTKOVER\_CLUSTER\_DISABLE - Cluster is disabled
- NOTKOVER\_VERSION - Version mismatch
- NOTKOVER\_SHELF\_HOT - Disk shelf is too hot
- NOTKOVER\_REVERT\_IN\_PROGRESS - Revert is in progress
- NOTKOVER\_HALT\_NOTKOVER - Node halted in no-takeover mode
- TKOVER\_ON\_REBOOT - Enable takeover on reboot
- TKOVER\_ON\_PANIC - Enabled takeover on panic
- TKOVER\_ON\_STUTTER\_DISABLED - Disable takeover on short uptime
- NFO\_DISK\_SHELF\_ENABLED - Negotiated failover for disk shelf module is enabled
- NFO\_ISCSI\_ENABLED - Negotiated failover for network interfaces module is enabled
- NFO\_FCP\_TARGET\_ENABLED - Negotiated failover for fcp target module is enabled

**[-partner-takeover-info <text>] - Partner Takeover Info**

Selects the nodes that have the specified partner node takeover information. This includes the type of negotiated failover request, or if takeover is not possible, the reason why takeover is disabled. Possible values include:

- NOTKOVER\_NVRAM\_DOWN - NVRAM mirror is down
- NOTKOVER\_OPERATOR\_DISABLE\_NVRAM - Operator disabled
- NOTKOVER\_PARTNER\_RESET - A link reset is in progress
- NOTKOVER\_FM\_TAKEOVER - The failover monitor has declared takeover
- NOTKOVER\_NVRAM\_MISMATCH - NVRAM sizes mismatch
- NOTKOVER\_OPERATOR\_DENY - Operator denies takeover
- NOTKOVER\_CLUSTER\_DISABLE - Cluster is disabled
- NOTKOVER\_VERSION - Version mismatch

- NOTKOVER\_SHELF\_HOT - Disk shelf is too hot
- NOTKOVER\_REVERT\_IN\_PROGRESS - Revert is in progress
- NOTKOVER\_HALT\_NOTKOVER - Node halted in no-takeover mode
- TKOVER\_ON\_REBOOT - Takeover on reboot is enabled
- TKOVER\_ON\_PANIC - Takeover on panic is enabled
- TKOVER\_ON\_STUTTER\_DISABLED - Disable takeover on short uptime
- NFO\_DISK\_SHELF\_ENABLED - Negotiated failover for disk shelf module is enabled
- NFO\_ISCSI\_ENABLED - Negotiated failover for network interfaces module is enabled
- NFO\_FCP\_TARGET\_ENABLED - Negotiated failover for fcp target module is enabled

**[-local-headswap-state <Headswap State>] - Local Head Swap State**

Selects the nodes that have the specified local node headswap state. Possible values are:

- HEADSWAP\_NONE - head swap not in progress
- HEADSWAP\_START - head swap started
- HEADSWAP\_CFO\_START - CFO phase of head swap started
- HEADSWAP\_CFO\_END - CFO phase of head swap completed
- HEADSWAP\_SFO\_START - SFO phase of head swap started

**[-partner-headswap-state <Headswap State>] - Partner Head Swap State**

Selects the nodes that have the specified partner node headswap state. Possible values are:

- HEADSWAP\_NONE - head swap not in progress
- HEADSWAP\_START - head swap started
- HEADSWAP\_CFO\_START - CFO phase of head swap started
- HEADSWAP\_CFO\_END - CFO phase of head swap completed
- HEADSWAP\_SFO\_START - SFO phase of head swap started

**[-fast-timeout <integer>] - Fast Timeout**

Selects the nodes that have the specified fast-timeout configuration setting.

**[-slow-timeout <integer>] - Slow Timeout**

Selects the nodes that have the specified slow-timeout setting.

**[-mailbox-timeout <integer>] - Mailbox Timeout**

Selects the nodes that have the specified mailbox-timeout setting.

**[-connect-timeout <integer>] - Connect Timeout**

Selects the nodes that have the specified connect-timeout setting.

**[-operator-timeout <integer>] - Operator Timeout**

Selects the nodes that have the specified operator-timeout setting.



**[-firmware-timeout <integer>] - Firmware Timeout**

Selects the nodes that have the specified firmware-timeout setting.

**[-dumpcore-timeout <integer>] - Dumpcore Timeout**

Selects the nodes that have the specified dump-core timeout setting.

**[-booting-timeout <integer>] - Booting Timeout**

Selects the nodes that have the specified booting-timeout setting.

**[-transit-timer {true|false}] - Transit Timer Enabled**

Selects the nodes that have the specified transit-timer setting.

**[-transit-timeout <integer>] - Transit Timeout**

Selects the nodes that have the specified transit timeout.

**[-firmware-received <integer>] - Firmware Received**

Selects the nodes that have the specified firmware-reception time.

**[-firmware-received-cycles <integer>] - Firmware Received in CPU Cycles**

Selects the nodes that have the specified firmware-reception time in CPU Cycles.

**[-booting-received <integer>] - Booting Received**

Selects the nodes that have the specified booting-reception time.

**[-transit-time <integer>] - Transit Event Time**

Selects the nodes whose last failover event occurred at the specified time.

**[-pnormal <integer>] - Primary Normal IO Time**

Selects the nodes that have the specified normal primary-mailbox I/O time.

**[-ptransition <integer>] - Primary Transition IO Time**

Selects the nodes that have the specified transitional primary-mailbox I/O time.

**[-bnormal <integer>] - Backup Normal IO Time**

Selects the nodes that have the specified normal backup-mailbox I/O time.

**[-btransition <integer>] - Backup Transition IO Time**

Selects the nodes that have the specified transitional backup-mailbox I/O time.

**[-logs-unsynced <integer>] - Logs Unsynced Count**

Selects the nodes that have the specified count of unsynchronized logs.

**[-auto-giveback {true|false}] - Auto Giveback Enabled**

Selects the nodes that have the specified auto-giveback setting.

**[-detection-time <integer>] - Takeover Detection Time (secs)**

Selects the nodes that have the specified detection-time setting.

- [-onfailure {true|false}] - Takeover on Failure Enabled**  
Selects the nodes that have the specified takeover-on-failure setting.
- [-onpanic {true|false}] - Takeover on Panic Enabled**  
Selects the nodes that have the specified takeover-on-panic setting.
- [-onshort-uptime {true|false}] - Takeover on Short Uptime Enabled**  
Selects the storage-failover configurations that match this parameter value.
- [-short-uptime <integer>] - Short Uptime (secs)**  
Selects the nodes that have the specified short-uptime value.
- [-attempts <integer>] - Number of Giveback Attempts**  
Selects the nodes that have the specified number of giveback attempts.
- [-attempts-time <integer>] - Giveback Attempts Period (minutes)**  
Selects the nodes that have the specified time setting for giveback attempts.
- [-propagate {true|false}] - Propagate Status via Mailbox**  
Selects the nodes that have the specified propagate-status-via-mailbox setting.
- [-read-interval <integer>] - Node Status Read Interval (secs)**  
Selects the nodes that have the specified read interval.
- [-write-interval <integer>] - Node Status Write Interval (secs)**  
Selects the nodes that have the specified write interval.
- [-onreboot {true|false}] - Takeover on Reboot Enabled**  
Selects the nodes that have the specified takeover-on-reboot setting.
- [-delay-seconds <integer>] - Delay Before Auto Giveback (secs)**  
Selects the nodes that have the specified delay (in seconds) for the auto giveback.
- [-hwassist {true|false}] - Hardware Assist Enabled**  
Selects the nodes that have the specified hwassist setting.
- [-hwassist-partner-ip <IP Address>] - Partner's Hwassist IP**  
Selects the nodes that have the specified hwassist-partner-ip setting.
- [-hwassist-partner-port <integer>] - Partner's Hwassist Port**  
Selects the nodes that have the specified hwassist-partner-port setting.
- [-hwassist-health-check-interval <integer>] - Hwassist Health Check Interval (secs)**  
Selects the nodes that have the specified hwassist health check interval, in seconds.
- [-hwassist-retry-count <integer>] - Hwassist Retry Count**  
Selects the nodes that have the specified hwassist retry count, in seconds.

**[-hwassist-status <text>] - Hwassist Status**

Selects the nodes that have the specified hwassist-status setting.

**[-time-until-autogiveback <integer>] - Time Until Auto Giveback (secs)**

Selects the nodes that have the specified time(in seconds) until auto giveback.

**[-local-mailbox-disks <text>] - Local Mailbox Disks**

Selects the nodes that have the specified mailbox disks on the local node.

**[-partner-mailbox-disks <text>] - Partner Mailbox Disks**

Selects the nodes that have the specified mailbox disks on the partner node.

**[-local-firmware-state <text>] - Local Firmware State**

Selects the nodes that have the specified firmware state on the local node.

**[-local-firmware-progress <integer>] - Local Firmware Progress Counter**

Selects the nodes that have the specified firmware progress counter for the local node.

**[-partner-firmware-state <text>] - Partner Firmware State**

Selects the nodes that have the specified firmware state of the partner node.

**[-partner-firmware-progress <integer>] - Partner Firmware Progress Counter**

Selects the nodes that have the specified firmware progress counter for the partner node.

**[-local-missing-disks <text>] - Missing Disks on Local Node**

Selects the nodes that have the specified missing disks on the local node.

**[-partner-missing-disks <text>] - Missing Disks on Partner Node**

Selects the nodes that have the specified missing disks on the partner node.

**[-reboot-timeout <integer>] - Reboot Timeout**

Selects the nodes that have the specified reboot timeout.

**[-time-since-takeover <text>] - Time Since Takeover**

Selects the nodes that have been in takeover mode for the specified amount of time.

**[-auto-giveback-after-panic {true|false}] - Auto Giveback After Takeover On Panic**

Selects the nodes that have the specified auto-giveback-after-panic setting. If *true* then an automatic giveback operation is invoked when the failover node of an HA pair is repaired and rebooted. The takeover node of the HA pair detects this and initiates a giveback operation automatically.

**[-is-giveback-requested {true|false}] - Giveback Requested**

Selects the nodes that have the specified is-giveback-requested setting. If *true*, a deferred giveback request has been made by the local node.

**[-auto-giveback-last-veto-check <integer>] - Auto Giveback Last Veto Check**

Selects the nodes that have the specified auto-giveback-last-veto-check time. This setting indicates the time, in milliseconds, when the local node made the most recent giveback veto check.

**[`-is-auto-giveback-attempts-exceeded {true|false}`] - Auto Giveback Attempts Exceeded**

Selects the nodes that have the specified `is-auto-giveback-attempts-exceeded` setting. If `true`, the local node has exceeded the maximum number of allowed auto giveback attempts.

**[`-was-auto-giveback-done {true|false}`] - Was Auto Giveback Done**

Selects the nodes that have the specified `was-auto-giveback-done` setting. If `true`, the last giveback was automatic (as opposed to a manual giveback).

**[`-is-cifs-auto-giveback-stopping {true|false}`] - Is CIFS Auto Giveback Stopping**

Selects the nodes that have the specified `is-cifs-auto-giveback-stopping` setting. If `true`, the local node has initiated CIFS termination as part of an automatic giveback.

**[`-aggregate-migration-timeout <integer>`] - Aggregate Migration Timeout (secs)**

Selects the nodes that have the specified aggregate migration timeout.

**[`-is-mirror-enabled {true|false}`] - Is NVRAM Mirroring Enabled**

Selects the nodes that have the specified `is-mirror-enabled` setting. If `true`, then NVRAM mirroring is enabled.

**[`-is-mirror-consistency-required {true|false}`] - Is Mirror Consistency Required**

Selects the nodes that have the specified `is-mirror-consistency-required` setting. If `true`, then NVRAM mirror consistency is required.

**[`-is-memory-insufficient {true|false}`] - Is Memory Insufficient To Takeover**

Selects the nodes that have the specified `is-memory-insufficient` setting. If `true`, the local node does not have enough memory to perform a takeover.

**[`-memio-state <memio status>`] - Current State of Memio Link**

Selects the nodes that have the specified memio layer link current state. Possible values are:

- UNINIT - Uninitialized
- CLOSED - Closed
- HB\_LISTEN - Listening for connect
- SYN\_SENT - Sent generation information
- ESTABLISHED - Connection established

**[`-is-degraded {true|false}`] - Are Partner Mailbox Disks Not Known**

Selects the nodes that have the specified `is-degraded` setting. If `true`, takeovers are deferred because partner mailbox disks are not known.

**[`-reserve-policy <reserve policy>`] - FM Reservation Policy**

Selects the nodes that have the specified disk reservation policy. Possible values are:

- RESERVE\_NO\_DISKS - no disk reservations made during takeover, nor are disk reservations released during giveback
- RESERVE\_LOCK\_DISKS\_ONLY - only mailbox disks are released during takeover and released during giveback
- RESERVE\_ONLY\_AT\_TAKEOVER - reservations are issued only at takeover time. All disks are

reserved. All reservations are released at giveback

- RESERVE\_ALWAYS\_AFTER\_TAKEOVER - reservations are issued at takeover. When disks are subsequently added, they are also reserved. All disks are released at giveback

**[-reset-disks {true|false}] - Issue Disk Resets during Failover**

Selects the nodes that have the specified reset-disks setting. If *true*, disks are reset during takeover/giveback.

**[-total-system-uptime <integer>] - Total System Uptime**

Selects the nodes that have the specified total system uptime, in milliseconds.

**[-current-time <integer>] - Current System Time**

Selects the nodes that have the specified current time on the node.

**[-fm-takeover-state <FM Takeover/Giveback Transition>] - FM Takeover State**

Selects the nodes that have the specified takeover state. Possible values are:

- FT\_NONE - Not in takeover
- FT\_TAKEOVER\_STARTED - Local node has initiated takeover
- FT\_TAKEOVER\_COMMITTED - Takeover has been committed
- FT\_TAKEOVER\_DONE\_OK - Local node successfully completed takeover
- FT\_TAKEOVER\_DONE\_FAILED - Takeover failed

**[-fm-giveback-state <FM Takeover/Giveback Transition>] - FM Giveback State**

Selects the nodes that have the specified giveback state. Possible values are:

- FT\_NONE - Not in giveback
- FT\_GIVEBACK\_READY - Partner node is ready for giveback
- FT\_GIVEBACK\_STARTED - Local node has initiated giveback
- FT\_GIVEBACK\_COMMITTED - Giveback has been committed
- FT\_GIVEBACK\_DONE\_OK - Giveback completed successfully

**[-takeover-reason <FM Takeover Reason>] - Reason why takeover triggered**

Selects the nodes that have the specified takeover reason. Possible values are:

- TAKEOVER\_NONE - Not in takeover
- TAKEOVER\_IMMEDIATE - Operator initiated forced takeover
- TAKEOVER\_NDU - Takeover initiated as part of NDU
- TAKEOVER\_FORCED - Operator initiated forced takeover, possible data loss
- TAKEOVER\_EARLY - Takeover occurred during the boot process
- TAKEOVER\_OPERATOR\_EXP - Takeover occurred after the operator timeout expired
- TAKEOVER\_POST\_FAILED - Takeover occurred on POST failure
- TAKEOVER\_PANIC - Takeover on panic
- TAKEOVER\_SHORTUPTIME - Takeover after rapid toggling between up and down states

- TAKEOVER\_SPARECORE\_EXP - Takeover on panic timeout expiration
- TAKEOVER\_REBOOT\_EXP - Takeover on reboot timer expiration
- TAKEOVER\_BOOTING\_EXP - Takeover on booting timer expiration
- TAKEOVER\_FIRMWARE\_EXP - Takeover on firmware timer expiration
- TAKEOVER\_NFO\_SHUTDOWN - Takeover on negotiated failover shutdown
- TAKEOVER\_NFO\_TIMER - Takeover on negotiated failover timer expiration
- TAKEOVER\_MDP - Takeover on multi-disk panic
- TAKEOVER\_REBOOT - Takeover on reboot
- TAKEOVER\_HALT - Takeover on halt
- TAKEOVER\_CLAM - CLAM-triggered takeover
- TAKEOVER\_HWASSIST - Hardware-assisted takeover
- TAKEOVER\_NORMAL - Operator initiated takeover

**[`-ha-type {none|shared_storage|non_shared_storage}`] - HA Type**

If this parameter is specified, the command selects the nodes that have the specified HA-type setting. If the value is set to `shared_storage`, then the node is in a storage-failover configuration using the shared storage. If it is set to `non_shared_storage`, then the node is in a storage-failover configuration using the unshared storage. If it is set to `none`, then the node is not part of a storage-failover configuration.

## Examples

The following example displays information about all storage-failover configurations:

```
cluster1::> storage failover show
                Takeover
Node      Partner  Possible State
-----
node0     node1     true      Connected to node1
node2     node3     true      Connected to node3
node1     node0     true      Connected to node0
node3     node2     true      Connected to node2
4 entries were displayed.
```

## storage failover takeover

Take over the storage of a node's partner

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

### Description

The `storage failover takeover` command initiates a takeover of the partner node's storage.

## Parameters

### { **-ofnode** {<nodename>|local} - Node to Takeover

This specifies the node that is taken over. It is shut down and its partner takes over its storage.

### | **-bynode** {<nodename>|local} - Node Initiating Takeover }

This specifies the node that is to take over its partner's storage.

### [**-option** <takeover option>] - Takeover Option

This optionally specifies the style of takeover operation. Possible values include the following:

- **normal** - Specifies a normal takeover operation; that is, the partner is given the time to close its storage resources gracefully before the takeover operation proceeds. This is the default value.
- **immediate** - Specifies an immediate takeover. In an immediate takeover, the takeover operation is initiated before the partner is given the time to close its storage resources gracefully. The use of this option results in an immediate takeover which does not do a clean shutdown. In case of NDU this can result in a NDU failure.



If this option is specified, negotiated takeover optimization is bypassed even if the `-bypass -optimization` option is set to false.



If this option is specified, migration of data LIFs from the partner will be delayed even if the `-skip-lif-migration-before-takeover` option is not specified. If possible, migrate the data LIFs to another node prior to specifying this option.

- **allow-version-mismatch** - If this value is specified, the takeover operation is initiated even if the partner is running a version of software that is incompatible with the version running on the node. In this case, the partner is given the time to close its storage resources gracefully before the takeover operation proceeds. However, the takeover operation will not be allowed if the partner has higher WAFL or RAID label versions. Use this value as part of a nondisruptive upgrade or downgrade procedure.
- **force** - If this value is specified, the takeover operation is initiated even if the node detects an error that normally prevents a takeover operation from occurring. This value is available only at the advanced privilege level and higher.



If this option is specified, negotiated takeover optimization is bypassed even if the `-bypass -optimization` option is set to false.



The use of this option can potentially result in data loss. If the HA interconnect is detached or inactive, or the contents of the failover partner's NVRAM cards are unsynchronized, takeover is normally disabled. Using the `-force` option enables a node to take over its partner's storage despite the unsynchronized NVRAM, which can contain client data that can be lost upon storage takeover.

### [**-bypass-optimization** {true|false}] - Bypass Takeover Optimization

If this is an operator-initiated planned takeover, this parameter specifies whether the takeover optimization is bypassed. This parameter defaults to false.



This parameter is ignored and negotiated takeover optimization automatically bypassed if the `-immediate` option, the `-force` option, or the `-allow-disk-inventory-mismatch` parameter is specified as part of the same `storage failover takeover` command.

#### **[`-allow-disk-inventory-mismatch {true|false}`] - Disk inventory**

If this parameter is specified, the takeover operation is initiated even if the local node cannot see the partner's filesystem disks.



If this parameter is specified, negotiated takeover optimization is bypassed even if the `-bypass-optimization` parameter is set to `false`.



The use of this parameter can potentially result in client outage.

#### **[`-skip-lif-migration-before-takeover <true>`] - Skip Migrating LIFs Away from Node Prior to Takeover**

This parameter specifies that LIF migration prior to takeover is skipped. However if LIFs on this node are configured for failover, those LIFs may still failover after the takeover has occurred. Without this parameter, the command attempts to synchronously migrate data and cluster management LIFs away from the node prior to its takeover. If the migration fails or times out, the takeover is aborted.

#### **[`-ignore-quorum-warnings <true>`] - Skip Quorum Check Before Takeover**

If this parameter is specified, quorum checks will be skipped prior to the takeover. The operation will continue even if there is a possible data outage due to a quorum issue.

## Examples

The following example causes a node named `node0` to initiate a negotiated optimized takeover of its partner's storage:

```
cluster1::> storage failover takeover -bynode node0
```

The following example causes a node named `node0` to initiate an immediate takeover of its partner's storage:

```
cluster1::> storage failover takeover -bynode node0 -option immediate
```

## storage failover hwassist show

Display hwassist status

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

### Description

The `storage failover hwassist show` command displays information about hardware assisted takeover configurations. By default, the command displays the following information:



- Node name.
- Partner node name.
- Whether hardware assisted takeover is enabled.
- IP address on which the local node receives hardware assist alerts.
- Port on which local node receives hardware assist alerts.
- Hardware assist monitor status.
- If the monitor is inactive, the reason it is inactive.
- If the monitor is inactive, the corrective action to make it active.

## Parameters

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

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

**| [-instance ] }**

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

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

Selects the hwassist configurations that match this parameter value.

**[-partner-name {<nodename>|local}] - Name of the Partner Node**

Selects the hwassist configurations that match this parameter value.

**[-enabled {true|false}] - Local Hardware Assist Enabled**

Selects the hwassist configurations that match this parameter value.

**[-local-status <text>] - Local Node's Hwassist Status**

Selects the hwassist configurations that match this parameter value (*active* or *inactive*).

**[-local-ip <text>] - IP Address on Which Local Node is Listening**

Selects the hwassist configurations that match this parameter value.

**[-local-port <integer>] - Port on Which Local Node is Listening**

Selects the hwassist configurations that match this parameter value.

**[-local-inactive <text>] - Local Node's Hwassist Inactive Status Reason**

Selects the hwassist configurations that match this parameter value.

**[-local-action <text>] - Corrective Action on Local Node**

Selects the hwassist configurations that match this parameter value.

## Examples

The following example displays the hardware assist information for the local node and its partner:

```

cluster1::> storage failover hwassist show
Node
-----
ha1
                Partner : ha2
        Hwassist Enabled : true
                Hwassist IP : 10.225.248.19
                Hwassist Port : 4444
                Monitor Status : active
                Inactive Reason : -
                Corrective Action : -

ha2
                Partner : ha1
        Hwassist Enabled : true
                Hwassist IP : 10.225.248.21
                Hwassist Port : 4444
                Monitor Status : active
                Inactive Reason : -
                Corrective Action : -

```

## storage failover hwassist test

Test the hwassist functionality

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

### Description

The `storage failover hwassist test` command tests the Hardware Assist h/w connectivity between the two nodes in a HA pair. The test result can be one of the following.

- Hardware Assist is not initialized.
- HW assist is not supported.
- Partner is throttling alerts.
- Resource is busy.
- Hardware Assist h/w returned an error.
- No response from partner. Timed out.
- Unexpected abort.
- Partner has taken over.
- Interconnect is down between nodes.
- Partner is not booted up yet.

## Parameters

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

This specifies the node from which a test alert is initiated.

## Examples

The following command issues a test alert from the node cluster1-01:

```
cluster1::> storage failover hwassist test -node cluster1-01
Info: Operation successful.
```

## storage failover hwassist stats clear

Clear the hwassist statistics

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

### Description

The `storage failover hwassist stats clear` command clears the statistics information maintained by Hardware Assist functionality.

## Parameters

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

This specifies the node on which the statistics are to be cleared.

## Examples

The following example clears the hwassist statistics on the node cluster1-01:

```
cluster1::> storage failover hwassist stats clear -node cluster1-01
```

## storage failover hwassist stats show

Display hwassist statistics

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

### Description

The `storage failover hwassist stats show` command displays statistics about the hardware assist alerts processed by a node. The command displays the following information for each alert:

- Locally enabled.

- Partner Inactive Reason.
- Alert type.
- Event that triggered the alert.
- The number of times the alert has been received.
- Whether takeover was possible on receiving the alert.
- The last time at which the alert was received.

## Parameters

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

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

**| [-instance ] }**

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

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

Selects the hwassist statistics for the specified node.

## Examples

The following example displays the hwassist statistics for the node ha1:

```

cluster1::> storage failover hwassist stats show -node ha1
Node: ha1
      Local Enabled: true
      Partner Inactive Reason: -
Alert Type  Alert Event                Count Takeover  Last Received
-----
-----
system_down power_loss                      0 Yes         ---
system_down l2_watchdog_reset          0 Yes         ---
system_down power_off_via_rlm       0 Yes         ---
system_down power_cycle_via_rlm     0 Yes         ---
system_down reset_via_rlm           0 Yes         ---
system_down power_off_via_sp        0 Yes         ---
system_down power_cycle_via_sp      0 Yes         ---
system_down reset_via_sp             0 Yes         ---
system_down post_error               0 No          ---
system_down abnormal_reboot         0 No          ---
system_down loss_of_heartbeat        0 No          ---
keep_alive  periodic_message           10 No          Wed Mar  9 13:41:28
EST 2016
test        test                    0 No          ---
ID_mismatch ---                      0 ---        ---
Key_mismatch ---                    0 ---        ---
Unknown     ---                      0 ---        ---
            alerts_throttled       0 ---        ---

```

The following example displays the hwassist statistics for the node ha1 where hardware assist hardware is not supported.

```

cluster1::> storage failover hwassist stats show -node ha1
Node: ha1
      Local Enabled: false
      Partner Inactive Reason: HW assist is not supported on partner.
Alert Type  Alert Event                Count Takeover  Last Received
-----
-----
-

```

## storage failover internal-options show

Display the internal options for storage failover

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

## Description

The `storage failover internal-options show` command displays the following information about the storage failover configuration:

- Node name
- Whether automatic giveback is enabled
- Whether partner checking is enabled
- Takeover detection time, in seconds
- Whether takeover on failover is enabled
- Whether takeover on panic is enabled
- Whether takeover on reboot is enabled
- Whether hardware-assisted takeover is enabled
- IP address on which the partner node listens to the hardware-assisted takeover alerts
- Port on which the partner node listens to the hardware-assisted takeover alerts
- Whether takeover on short uptime is enabled (detailed view only)
- Short uptime interval, in seconds (detailed view only)
- Number of giveback attempts (detailed view only)
- Giveback attempt interval, in minutes (detailed view only)
- Whether status is propagated through SFO mailboxes (detailed view only)
- Status read interval, in seconds (detailed view only)
- Status write interval, in seconds (detailed view only)
- Hardware-assisted takeover retry count (detailed view only)
- Hardware-assisted takeover heartbeat period (detailed view only)
- Whether operator-initiated planned takeover is optimized

## Parameters

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

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

**| [-more ]**

This parameter displays the following additional information: :

- Node name
- Whether takeover on short uptime is enabled
- Short uptime interval, in seconds
- Number of giveback attempts
- Giveback attempt interval, in minutes
- Whether status is propagated through SFO mailboxes
- Status read interval, in seconds

- Status write interval, in seconds
- Hardware-assisted takeover retry count
- Hardware-assisted takeover heartbeat period

**[`-instance` ] }**

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

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

Selects configuration information for the specified node.

**[`-auto-giveback` {true|false}] - Auto Giveback Enabled**

Selects configuration information for nodes that have the specified automatic giveback setting.

**[`-check-partner` {true|false}] - Check Partner Enabled**

Selects configuration information for nodes that have the specified partner-checking setting.

**[`-detection-time` <integer>] - Takeover Detection Time (secs)**

Selects configuration information for nodes that have the specified takeover detection time setting.

**[`-onfailure` {true|false}] - Takeover on Failure Enabled**

Selects configuration information for nodes that have the specified takeover-on-failure setting.

**[`-onpanic` {true|false}] - Takeover on Panic Enabled**

Selects configuration information for nodes that have the specified takeover-on-panic setting.

**[`-onshort-uptime` {true|false}] - Takeover on Short Uptime Enabled**

Selects configuration information for nodes that have the specified takeover-on-short-uptime setting.

**[`-short-uptime` <integer>] - Short Uptime (secs)**

Selects configuration information for nodes that have the specified takeover-on-short-uptime time setting.

**[`-attempts` <integer>] - Number of Giveback Attempts**

Selects configuration information for nodes that have the specified number of giveback attempts setting.

**[`-attempts-time` <integer>] - Giveback Attempts Minutes**

Selects configuration information for nodes that have the specified giveback attempt time setting.

**[`-propagate` {true|false}] - Propagate Status via Mailbox**

Selects configuration information for nodes that have the specified setting for propagation of status through Storage Failover mailboxes.

**[`-read-interval` <integer>] - Node Status Read Interval (secs)**

Selects configuration information for nodes that have the specified status read interval setting.

**[`-write-interval` <integer>] - Node Status Write Interval (secs)**

Selects configuration information for nodes that have the specified status write interval setting.

**[-onreboot {true|false}] - Takeover on Reboot Enabled**

Selects configuration information for nodes that have the specified takeover-on-reboot setting.

**[-delay-seconds <integer>] - Delay Before Auto Giveback (secs)**

If this parameter is specified, the command displays information only about the node or nodes that have the specified delay for auto giveback.

**[-hwassist {true|false}] - Hwassist Enabled**

Selects configuration information for nodes that have the specified hardware-assisted takeover setting.

**[-hwassist-partner-ip <text>] - Partner's Hwassist IP**

Selects configuration information for nodes that have the specified partner IP setting for hardware-assisted takeovers.

**[-hwassist-partner-port <integer>] - Partner's Hwassist Port**

Selects configuration information for nodes that have the specified partner port setting for hardware-assisted takeovers.

**[-hwassist-health-check-interval <integer>] - Hwassist Health Check Interval (secs)**

Selects configuration information for nodes that have the specified health check interval setting for hardware-assisted takeovers

**[-hwassist-retry-count <integer>] - Hwassist Retry Count**

Selects configuration information for nodes that have the specified retry count (in seconds) for hardware-assisted takeovers.

**[-mode {ha|non\_ha}] - HA Mode**

If this parameter is specified, the command displays information only about the node or nodes that have the specified HA mode.

**[-bypass-takeover-optimization {true|false}] - Bypass Takeover Optimization Enabled**

Selects configuration information for nodes that have the specified setting for bypass takeover optimization ( `_ true_` means that optimized operator-initiated planned takeover is bypassed, `false` means that it is enabled). Operator-initiated planned takeover is optimized when SFO aggregates are relocated serially to the partner prior to takeover. This reduces client outage.

## Examples

The following example displays detailed information about the internal options for storage failover on a node named node2:



```

cluster1::*> storage failover internal-options show -node node2
Node: node2

                Auto Giveback Enabled: false
                Check Partner Enabled: true
                Takeover Detection Time (secs): 15
                Takeover On Failure Enabled: true
                Takeover On Panic Enabled: false
                Takeover On Short Uptime Enabled: true
                Short Uptime (secs): -
                Number of Giveback Attempts: 3
                Giveback Attempts Minutes: 10
                Propagate Status Via Mailbox: true
                Node Status Read Interval (secs): 5
                Node Status Write Interval (secs): 5
                Failover the Storage when Cluster Ports Are Down: -
                Failover Interval when Cluster Ports Are Down (secs): -
                Takeover on Reboot Enabled: true
                Delay Before Auto Giveback (secs): 300
                Hardware Assist Enabled: true
                Partner's Hw-assist IP:
                Partner's Hw-assist Port: 4444
                Hw-assist Health Check Interval (secs): 180
                Hw-assist Retry count: 2
                HA mode: ha
                Bypass Takeover Optimization Enabled: true

```

## storage failover mailbox-disk show

Display information about storage failover mailbox disks

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

### Description

The `storage failover mailbox-disk show` command lists the mailbox disks that are used by storage failover. The command displays the following information:

- Node name
- Whether the mailbox disk is owned by the local node or by its partner
- Disk name
- Disk universal unique identifier (UUID)

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

## Parameters

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

If -fields <fieldname>,... is used, the command displays only the specified fields.

**| [-instance ] }**

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

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

Selects the mailbox disks that are associated with the specified node.

**[-location {local|partner|tertiary}] - Mailbox Owner**

Selects the mailbox disks that have the specified relationship to the node.

**[-diskindex <integer>] - Mailbox Disk Index**

Selects the mailbox disk that has the specified index number.

**[-diskname <text>] - Mailbox Disk Name**

Selects the mailbox disks that match the specified disk name.

**[-diskuuid <text>] - Mailbox Disk UUID**

Selects the mailbox disks that match the specified UUID.

**[-physical-location {local|partner|mediator}] - Mailbox Disk Physical Location**

Selects the mailbox disks that match the specified physical location.

**[-location-id <nvramid>] - System ID of the Node where the Disk is Attached**

Selects the mailbox disks that match the specified location-id.

**[-location-name <text>] - Mailbox Disk Location**

Selects the mailbox disks that match the specified location-name.

## Examples

The following example displays information about the mailbox disks on a node named node1:

```
cluster1::*> storage failover mailbox-disk show -node node1
Node      Location  Index Disk Name      Physical Location  Disk UUID
-----
node1
  local    0 1.0.4      local    20000000:8777E9D6:[...]
  local    1 1.0.6      partner  20000000:8777E9DE:[...]
  partner  0 1.0.1      local    20000000:877BA634:[...]
  partner  1 1.0.2      partner  20000000:8777C1F2:[...]
```

# storage failover progress-table show

Display status information about storage failover operations

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

## Description

The `storage failover progress-table show` displays status information about storage-failover operations. This information is organized in a resource table. The command displays the following information:

- Node name
- Resource-entry index number
- Resource-entry name
- Resource-entry state
- Resource-entry failure code
- Resource-entry time delta

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

## Parameters

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

If `-fields <fieldname>, ...` is used, the command will only displays only the specified fields.

**| [-instance ] }**

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

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

Selects the status information for the specified node.

**[-index <integer>] - Resource Table Index**

Selects the status information for the specified index number.

**[-entryname <text>] - Resource Table Entry Name**

Selects the status information for the specified entry name.

**[-state <text>] - Resource Table Entry State**

Selects the status information for the specified state. Possible values include UP, START\_RUNNING, START\_DONE, START\_FAILED, STOP\_RUNNING, STOP\_FAILED, TAKEOVER\_BARRIER, and ONLY\_WHEN\_INITD.

**[-failurecode <text>] - Entry Failure Code**

Selects the status information for the specified failure code. Possible values include OK, FAIL, FAIL\_ALWAYS, HANG, PANIC, and VETO.

**[-timedelta <integer>] - Entry Time Delta**

Selects the status information for the specified time delta.

## Examples

The following example displays the entire storage-failover resource table:

```
cluster1::*> storage failover progress-table show
Node   Entry Name                               State   Time Delta
-----
node0
  Pre-rsrctbl: fmdisk_resumePartnerDi    start_done    6
  Pre-rsrctbl: coredump_get_busy_spar     start_done    107
  Pre-rsrctbl: raid_preread_labels_be     start_done    1
  Pre-rsrctbl: fmdisk_reserve_all         start_done    84
  rsrctbl: fmrsrc_giveback_done           start_done    0
  rsrctbl: fmic                           start_done    0
  rsrctbl: fmdisk_reserve                 start_done    171
  rsrctbl: fm_partnerSlowTimeout          start_done    1
  rsrctbl: fmdisk_inventory               start_done    0
  rsrctbl: fmfsm_reserve                  start_done    0
Press <space> to page down, <return> for next line, or 'q' to quit...
Node   Entry Name                               State   Time Delta
-----
node0
  rsrctbl: rdb-ha                         start_done    36
  rsrctbl: giveback_cleanup_wait          start_done    0
  rsrctbl: priority_ha                    start_done    0
  rsrctbl: raid                           start_done    113
  rsrctbl: raid_disaster_early            start_done    0
  rsrctbl: wafl_nvram_replay              start_done    0
  rsrctbl: takeover_test_1                start_done    0
```

## Copyright Information

Copyright © 2022 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means-graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system-without prior written permission of the copyright owner.

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

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

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

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

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

## 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.