



ONTAP® 9

High-Availability Configuration Guide

May 2021 | 215-11144_2021-05_en-us
doccomments@netapp.com

Updated for ONTAP 9.8

 **NetApp®**

Contents

Deciding whether to use the High-Availability Configuration Guide..... 4

Planning your HA pair configuration..... 5

- Best practices for HA pairs.....5
- Setup requirements and restrictions for standard HA pairs.....6
- Setup requirements and restrictions for mirrored HA pairs.....7
- Requirements for hardware-assisted takeover..... 7
- If your cluster consists of a single HA pair..... 8
- Storage configuration variations for HA pairs.....8
- HA pairs and storage system model types.....9
 - Single-chassis and dual-chassis HA pairs.....9
 - Interconnect cabling for systems with variable HA configurations.....9
 - HA configuration and the HA state PROM value.....10

Requirements for cabling HA pair..... 11

- System cabinet or equipment rack installation..... 11
 - HA pairs in an equipment rack..... 11
 - HA pairs in a system cabinet..... 11
- Required documentation..... 12
- Required tools..... 12
- Required equipment..... 13
- Preparing your equipment..... 14
 - Installing the nodes in equipment racks..... 14
 - Installing the nodes in a system cabinet..... 14
- Cabling a standard HA pair 14
 - Cabling the HA interconnect (all systems except 32xx or 80xx in separate chassis)..... 15
 - Cabling the HA interconnect (32xx systems in separate chassis)..... 16
 - Cabling the HA interconnect (80xx systems in separate chassis)..... 16
- Cabling a mirrored HA pair..... 17
 - Cabling the HA interconnect (all systems except 32xx or 80xx in separate chassis)..... 17
 - Cabling the HA interconnect (32xx systems in separate chassis)..... 18
 - Cabling the HA interconnect (80xx systems in separate chassis)..... 18
- Required connections for using uninterruptible power supplies with standard or mirrored HA pairs..... 19

Configuring an HA pair..... 20

- Verifying and setting the HA state on the controller modules and chassis.....20
- Setting the HA mode and enabling storage failover..... 21
 - Commands for setting the HA mode..... 22
 - Commands for enabling and disabling storage failover..... 22
- Enabling cluster HA and switchless-cluster in a two-node cluster..... 22
- Checking for common configuration errors using Config Advisor..... 23
- Configuring hardware-assisted takeover..... 23
 - Commands for configuring hardware-assisted takeover..... 23
 - System events that trigger hardware-assisted takeover..... 24

- Configuring automatic takeover..... 25
 - Commands for controlling automatic takeover..... 25
 - System events that result in an automatic takeover..... 25
- Configuring automatic giveback..... 25
 - Commands for configuring automatic giveback..... 25
 - How variations of the storage failover modify command affect automatic giveback..... 26
- Testing takeover and giveback..... 28

- Monitoring an HA pair.....30**
 - Commands for monitoring an HA pair..... 30
 - Description of node states displayed by storage failover show-type commands..... 30

- Halting or rebooting a node without initiating takeover.....39**
 - Commands for halting or rebooting a node without initiating takeover..... 39
 - Halting or rebooting a node without initiating takeover in a two-node cluster..... 39

- Performing a manual takeover..... 41**
 - Commands for performing and monitoring manual takeovers..... 41
 - Moving epsilon for certain manually initiated takeovers..... 42

- Performing a manual giveback..... 44**
 - If giveback is interrupted..... 44
 - If giveback is vetoed..... 44
 - Commands for performing a manual giveback..... 45

- Nondisruptive operations with HA pairs..... 47**
 - Where to find procedures for nondisruptive operations with HA pairs..... 47

- Relocating aggregate ownership within an HA pair..... 48**
 - Relocating aggregate ownership..... 48
 - Commands for aggregate relocation..... 49

- Appendix: Understanding takeover and giveback..... 51**
 - When takeovers occur..... 51
 - How hardware-assisted takeover speeds up takeover..... 51
 - How automatic takeover and giveback work..... 52
 - What happens during takeover..... 52
 - What happens during giveback..... 53
 - HA policy and how it affects takeover and giveback operations..... 54
 - Background disk firmware update and takeover, giveback, and aggregate relocation..... 55
 - Types of disk ownership..... 55

- Copyright, trademark, and machine translation..... 56**
 - Copyright..... 56
 - Trademark..... 56
 - Machine translation..... 56

Deciding whether to use the High-Availability Configuration Guide

This guide describes how to install, configure, and manage NetApp high-availability (HA) pairs.

Note: This is archived content that includes some information specific only to ONTAP 9 versions 9.0 through 9.2.

You should use this guide under the following circumstances:

- You are setting up or managing a cluster running ONTAP 9.0 through ONTAP 9.2.
HA content for versions of ONTAP 9.3 and higher is available at <https://docs.netapp.com/us-en/ontap/>.
- You want to understand the range of HA pair capabilities.
- You want to understand the requirements and best practices for configuring HA pairs.
- You are using the ONTAP CLI to set up and monitor your HA pairs.

Planning your HA pair configuration

As you plan your HA pair, you must consider recommended best practices, the requirements, and the possible variations.

Best practices for HA pairs

To help your HA pair to be robust and operational, you must be familiar with configuration best practices.

- You must not use the root aggregate for storing data.
Storing user data in the root aggregate adversely affects system stability and increases the storage failover time between nodes in an HA pair.
- You must verify that each power supply unit in the storage system is on a different power grid so that a single power outage does not affect all power supply units.
- You must use LIFs (logical interfaces) with defined failover policies to provide redundancy and improve availability of network communication.
- Keep both nodes in the HA pair on the same version of ONTAP.
- Follow the documented procedures when upgrading your HA pair.
Upgrade, revert, or downgrade
- You must verify that you maintain a consistent configuration between the two nodes.
An inconsistent configuration is often the cause of failover problems.
- You must verify that you test the failover capability routinely (for example, during planned maintenance) to verify proper configuration.
- You must verify that each node has sufficient resources to adequately support the workload of both nodes during takeover mode.
- You must use the Config Advisor tool to help make failovers successful.
- If your system supports remote management (through a Service Processor), you must configure it properly.
System administration
- You must verify that you follow recommended limits for FlexVol volumes, dense volumes, Snapshot copies, and LUNs to reduce takeover or giveback time.
When adding FlexVol volumes to an HA pair, you should consider testing the takeover and giveback times to verify that they fall within your requirements.
- For systems using disks, ensure that you check for failed disks regularly and remove them as soon as possible.
Failed disks can extend the duration of takeover operations or prevent giveback operations.
Disk and aggregate management
- Multipath HA connection is required on all HA pairs except for some FAS22xx, FAS25xx, and FAS2600 series system configurations, which use single-path HA and lack the redundant standby connections.
- To receive prompt notification if the takeover capability becomes disabled, you should configure your system to enable automatic email notification for the `takeover impossible` EMS messages:
 - `ha.takeoverImpVersion`
 - `ha.takeoverImpLowMem`
 - `ha.takeoverImpDegraded`
 - `ha.takeoverImpUnsync`
 - `ha.takeoverImpIC`
 - `ha.takeoverImpHotShelf`

- `ha.takeoverImpNotDef`
- Avoid using the `-only-cfo-aggregates` parameter with the `storage failover giveback` command.

Setup requirements and restrictions for standard HA pairs

You must follow certain requirements and restrictions when setting up a new standard HA pair. These requirements help provide the data availability benefits of the HA pair design.

The following list specifies the requirements and restrictions you should be aware of when setting up a new standard HA pair:

- **Architecture compatibility**

Both nodes must have the same system model and be running the same ONTAP software and system firmware versions. The ONTAP release notes list the supported storage systems.

[ONTAP 9 Release Notes](#)
[NetApp Hardware Universe](#)
- **Nonvolatile memory (NVRAM or NVMEM) size and version compatibility**

The size and version of the system's nonvolatile memory must be identical on both nodes in an HA pair.
- **Storage capacity**
 - The number of disks or array LUNs must not exceed the maximum configuration capacity.
 - The total storage attached to each node must not exceed the capacity for a single node.
 - If your system uses native disks and array LUNs, the combined total of disks and array LUNs cannot exceed the maximum configuration capacity.
 - The total storage attached to each node must not exceed the capacity for a single node.
 - To determine the maximum capacity for a system using disks, array LUNs, or both, see the *Hardware Universe* at hwu.netapp.com.

Note: After a failover, the takeover node temporarily serves data from all of the storage in the HA pair.
- **Disks and disk shelf compatibility**
 - FC, SATA, and SAS storage are supported in standard HA pairs.
 - FC disks cannot be mixed on the same loop as SATA or SAS disks.
 - AFF platforms support only SSD storage.
No other hard disk drives (HDDs) or LUNs are supported on these HA pairs.
 - Different connection types cannot be combined in the same stack.
 - Different types of storage can be used on separate stacks on the same node.
You can also dedicate a node to one type of storage and the partner node to a different type, if needed.

[NetApp Hardware Universe](#)
[Disk and aggregate management](#)

 - Multipath HA connection is required on all HA pairs except for some FAS22xx, FAS25xx, and FAS2600 series system configurations, which use single-path HA and lack the redundant standby connections.
- **Mailbox disks or array LUNs on the root volume**
 - Two disks are required if the root volume is on a disk shelf.
 - One array LUN is required if the root volume is on a storage array.
- **Interconnect adapters and cables**

The HA interconnect adapters and cables must be installed unless the system has two controllers in the chassis and an internal interconnect.

- **Network connectivity**

Both nodes must be attached to the same network and the Network Interface Cards (NICs) or onboard Ethernet ports must be configured correctly.

- **System software**

The same system software, such as SyncMirror, Server Message Block (SMB) or Common Internet File System (CIFS), or Network File System (NFS), must be licensed and enabled on both nodes.

Note: If a takeover occurs, the takeover node can provide only the functionality for the licenses installed on it. If the takeover node does not have a license that was being used by the partner node to serve data, your HA pair loses functionality after a takeover.

- **Systems using array LUNs**

For an HA pair using array LUNs, both nodes in the pair must be able to detect the same array LUNs.

Note: Only the node that is the configured owner of a LUN has read-and-write access to that LUN. During takeover operations, the emulated storage system maintains read-and-write access to the LUN.

Related reference

[Commands for performing and monitoring manual takeovers](#) on page 41

You can manually initiate the takeover of a node in an HA pair to perform maintenance on that node while it is still serving the data on its disks, array LUNs, or both to users.

Setup requirements and restrictions for mirrored HA pairs

The restrictions and requirements for mirrored HA pairs include those for a standard HA pair with these additional requirements for disk pool assignments and cabling.

- You must ensure that your pools are configured correctly:
 - Disks or array LUNs in the same plex must be from the same pool, with those in the opposite plex from the opposite pool.
 - There must be sufficient spares in each pool to account for a disk or array LUN failure.
 - Both plexes of a mirror should not reside on the same disk shelf because it might result in a single point of failure.
- The `storage failover` command's `-mode` option must be set to `ha`.
- If you are using array LUNs, paths to an array LUN must be redundant.

Related reference

[Commands for setting the HA mode](#) on page 22

There are specific ONTAP commands for setting the HA mode. The system must be physically configured for HA before HA mode is selected. A reboot is required to implement the mode change.

Requirements for hardware-assisted takeover

The hardware-assisted takeover feature is available on systems where the Service Processor module is configured for remote management. Remote management provides remote platform management capabilities, including remote access, monitoring, troubleshooting, logging, and alerting features.

Although a system with remote management on both nodes provides hardware-assisted takeover for both, hardware-assisted takeover is also supported on HA pairs in which only one of the two systems has remote management configured. Remote management does not have to be configured

on both nodes in the HA pair. Remote management can detect failures on the system in which it is installed and provide faster takeover times if a failure occurs on the system.

Related information

[System administration](#)

If your cluster consists of a single HA pair

A cluster with only two nodes presents unique challenges in maintaining a quorum, the state in which a majority of nodes in the cluster have good connectivity. In a two-node cluster, neither node holds *epsilon*, the value that designates one of the nodes as the master. *Epsilon* is required in clusters with more than two nodes. Instead, both nodes are polled continuously to ensure that if takeover occurs, the node that is still up and running has full read-write access to data as well as access to logical interfaces and management functions. This continuous polling function is referred to as *cluster high availability* or *cluster HA*.

Cluster HA is different and separate from the high availability provided by HA pairs and the `storage failover` commands. While crucial to full functional operation of the cluster after a failover, cluster HA does not provide the failover capability of the storage failover functionality.

Related tasks

[Enabling cluster HA and switchless-cluster in a two-node cluster](#) on page 22

A cluster consisting of only two nodes requires special configuration settings. Cluster high availability (HA) differs from the HA provided by storage failover, and is required in a cluster if it contains only two nodes. If you have a switchless configuration, the `switchless-cluster` option must be enabled. Starting in ONTAP 9.2, detection of a switchless cluster is automatic.

Related reference

[Halting or rebooting a node without initiating takeover](#) on page 39

You can prevent an automatic storage failover takeover when you halt or reboot a node. This ability enables specific maintenance and reconfiguration operations.

Related information

[System administration](#)

Storage configuration variations for HA pairs

Because your storage management and performance requirements can vary, you can configure HA pairs symmetrically, asymmetrically, as an active/passive pair, or with shared disk shelf stacks.

Symmetrical (active/active) configurations

In a symmetrical configuration, each node has the same amount of storage.

Asymmetrical configurations

In an asymmetrical standard configuration, one node has more storage than the other node. This configuration is supported as long as the nodes do not exceed the maximum storage capacity.

Active/passive configurations

In active or passive configurations, the passive node has a root volume, and the active node has all of the remaining storage in addition to serving all of data requests during normal operation. The passive node responds to data requests only if it has taken over the active node.

Shared stacks

In this configuration, shared stacks between the nodes are particularly useful for active/passive configurations, as described in the Active/passive configurations bullet.

HA pairs and storage system model types

Different model storage systems support different HA configurations. This includes the physical configuration of the HA pair and the manner in which the system recognizes that it is in an HA pair.

Note: The physical configuration of the HA pair does not affect the cluster cabling of the nodes in the HA pair.

You can find more information about HA configurations supported by storage system models in the *Hardware Universe*.

Related information

[NetApp Hardware Universe](#)

Single-chassis and dual-chassis HA pairs

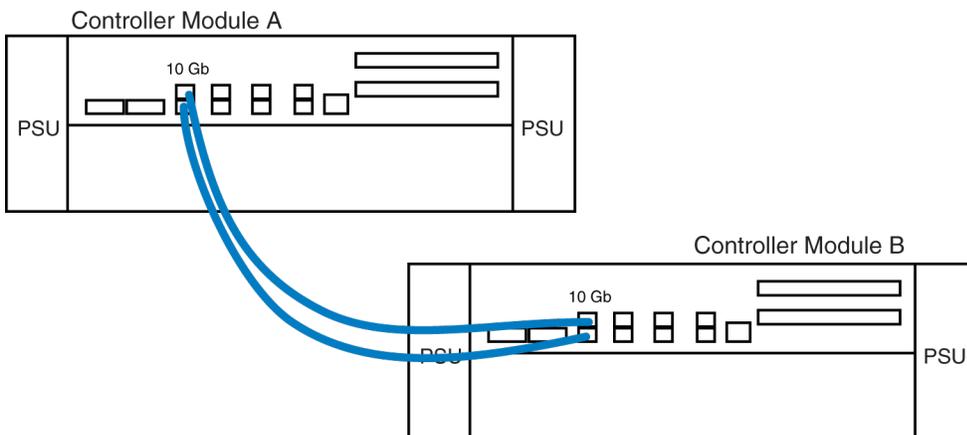
Depending on the model of the storage system, an HA pair can consist of two controllers in a single chassis, or two controllers in two separate chassis. Some models can be configured either way, while other models can be configured only as a single-chassis HA pair or dual-chassis HA pair.

The following example shows a single-chassis HA pair:



In a single-chassis HA pair, both controllers are in the same chassis. The HA interconnect is provided by the internal backplane. No external HA interconnect cabling is required.

The following example shows a dual-chassis HA pair and the HA interconnect cables:



In a dual-chassis HA pair, the controllers are in separate chassis. The HA interconnect is provided by external cabling.

Interconnect cabling for systems with variable HA configurations

In systems that can be configured either as a single-chassis or dual-chassis HA pair, the interconnect cabling is different depending on the configuration.

The following table describes the interconnect cabling for 32xx and 62xx systems:

If the controller modules in the HA pair are...	The HA interconnect cabling is...
Both in the same chassis	Not required, since an internal interconnect is used
Each in a separate chassis	Required

HA configuration and the HA state PROM value

Some controller modules and chassis automatically record in a PROM whether they are in an HA pair or stand-alone. This record is the *HA state* and must be the same on all components within the stand-alone system or HA pair. The HA state can be manually configured if necessary.

Related tasks

[Verifying and setting the HA state on the controller modules and chassis](#) on page 20

For systems that use the HA state value, the value must be consistent in all components in the HA pair. You can use the Maintenance mode `ha-conf ig` command to verify and, if necessary, set the HA state.

Requirements for cabling HA pair

If you want to install and cable a new standard or mirrored HA pair, you must have the correct tools and equipment and you must connect the controllers to the disk shelves. If it is a dual-chassis HA pair, you must also cable the HA interconnect between the nodes. HA pairs can be installed in either NetApp system cabinets or in equipment racks.

The term V-Series system refers to the storage systems released prior to Data ONTAP 8.2.1 that can use array LUNs. The FAS systems released in Data ONTAP 8.2.1 and later can use array LUNs if the proper license is installed.

The specific cabling procedure you use depends on whether you have a standard or mirrored HA pair.

[NetApp Documentation: Disk Shelves](#)

Multipath HA connection is required on all HA pairs except for some FAS22xx, FAS25xx, and FAS2600 series system configurations, which use single-path HA and lack the redundant standby connections.

System cabinet or equipment rack installation

You need to install your HA pair in one or more NetApp system cabinets or in standard telco equipment racks. Each of these options has different requirements.

HA pairs in an equipment rack

Depending on the amount of storage you ordered, you need to install the equipment in one or more telco-style equipment racks.

The equipment racks can hold one or two nodes on the bottom as well as eight or more disk shelves. For information about how to install the disk shelves and nodes into equipment racks, see the appropriate documentation that came with your equipment.

[NetApp Documentation: Product Library A-Z](#)

HA pairs in a system cabinet

Depending on the number of disk shelves, the HA pair you ordered arrives in a single system cabinet or multiple system cabinets.

The number of system cabinets you receive depends on how much storage you ordered. All internal adapters such as networking adapters, Fibre Channel adapters, and other adapters arrive preinstalled in the nodes.

If it comes in a single system cabinet, both the Channel A and Channel B disk shelves are cabled, and the HA adapters are also precabled.

If the HA pair you ordered has more than one cabinet, you must complete the cabling by cabling the local node to the partner node's disk shelves and the partner node to the local node's disk shelves. You must also cable the nodes together by cabling the NVRAM HA interconnects. If the HA pair uses switches, you must install the switches as described in the accompanying switch documentation. The system cabinets might also need to be connected to each other. See your *System Cabinet Guide* for information about connecting your system cabinets together.

Required documentation

Installing an HA pair requires that you have the correct documentation.

The following table lists and briefly describes the documentation you might need to refer to when preparing a new HA pair, or converting two stand-alone systems into an HA pair:

Manual name	Description
NetApp Hardware Universe	This utility describes the physical requirements that your site must meet to install NetApp equipment.
The appropriate system cabinet guide	This guide describes how to install NetApp equipment into a system cabinet.
The appropriate disk shelf guide NetApp Documentation: Disk Shelves	These guides describe how to cable a disk shelf to a storage system.
The appropriate hardware documentation for your storage system model	These guides describe how to install the storage system, connect it to a network, and bring it up for the first time.
Diagnostics Guide	This guide describes the diagnostics tests that you can run on the storage system.
Network and LIF management	This guide describes how to perform network configuration for the storage system.
Upgrade, revert, or downgrade	This guide describes how to upgrade storage system and disk firmware, and how to upgrade storage system software.
System administration	This guide describes general storage system administration, including tasks such as adding nodes to a cluster.
FlexArray virtualization installation requirements and reference	If you are installing an ONTAP system HA pair, this guide provides information about cabling ONTAP systems to storage arrays. You can also refer to the ONTAP system implementation guides for information about configuring storage arrays to work with ONTAP systems.
FlexArray virtualization implementation for third-party storage	If you are installing an ONTAP system HA pair, this guide provides information about configuring storage arrays to work with ONTAP systems.

Related information

[NetApp Documentation: Product Library A-Z](#)

Required tools

You must have the correct tools to install the HA pair.

You need the following tools to install the HA pair:

- #1 and #2 Phillips screwdrivers
- Hand level
- Marker

Required equipment

When you receive your HA pair, you should receive a list of required equipment.

For more information, see the *Hardware Universe* to confirm your storage system type, storage capacity, and so on.

hwu.netapp.com

Required equipment	Details
Storage system	Two of the same type of storage system
Storage	For more information, see the <i>Hardware Universe</i> . hwu.netapp.com
HA interconnect adapter card (for applicable controller modules that do not share a chassis) Note: When 32xx systems are in a dual-chassis HA pair, the c0a and c0b 10 GbE ports are the HA interconnect ports. They do not require an HA interconnect adapter. Regardless of configuration, the 32xx system's c0a and c0b ports cannot be used for data. They are used only for the HA interconnect.	InfiniBand (IB) HA adapter (The NVRAM adapter card functions as the HA interconnect adapter on applicable storage systems.) For more information, see the <i>Hardware Universe</i> hwu.netapp.com
SAS HBAs, if applicable	Minimum of two SAS HBAs or their equivalent in onboard ports
Fibre Channel switches, if applicable	Not applicable
Small form-factor pluggable (SFP) modules, if applicable	Not applicable
NVRAM HA adapter media converter	Only if using fiber cabling
Cables (provided with shipment unless otherwise noted)	<ul style="list-style-type: none"> • Two SAS controller-to-disk shelf cables per stack • Multiple disk shelf-to-disk shelf cables, if applicable • For systems using the IB HA interconnect adapter, two 4xIB copper cables, two 4xIB optical cables, or two optical cables with media converters <p style="margin-left: 40px;">Note: You must purchase longer optical cables separately for cabling distances greater than 30 meters.</p> <ul style="list-style-type: none"> • For 32xx systems that are in a dual-chassis HA pair, 10 GbE cables (Twin axial cabling or SR optical cables) for the HA interconnect

Preparing your equipment

You must install your nodes in your system cabinets or equipment racks, depending on your installation type.

Installing the nodes in equipment racks

Before you cable your nodes together, you must install the nodes and disk shelves in the equipment rack, label the disk shelves, and connect the nodes to the network.

Steps

1. Install the nodes in the equipment rack as described in the guide for your disk shelf, hardware documentation, or the *Installation and Setup Instructions* that came with your equipment.
2. Install the disk shelves in the equipment rack as described in the appropriate disk shelf guide.
3. Label the interfaces, where appropriate.
4. Connect the nodes to the network as described in the setup instructions for your system.

Result

The nodes are now in place and connected to the network; power is available.

After you finish

Cable the HA pair.

Installing the nodes in a system cabinet

Before you cable your nodes together, you must install the system cabinet, nodes, and any disk shelves, and connect the nodes to the network. If you have two cabinets, the cabinets must be connected together.

Steps

1. Install the system cabinets, nodes, and disk shelves as described in the *System Cabinet Guide*.
If you have multiple system cabinets, remove the front and rear doors and any side panels that need to be removed, and connect the system cabinets together.
2. Connect the nodes to the network, as described in the *Installation and Setup Instructions* for your system.
3. Connect the system cabinets to an appropriate power source and apply power to the cabinets.

Result

The nodes are now in place and connected to the network, and power is available.

After you finish

Proceed to cable the HA pair.

Cabling a standard HA pair

To cable a standard HA pair you should cable the SAS disk shelves and then cable the HA interconnect.

Before you begin

You must have cabled the SAS disk shelves. The disk shelf product library has instructions for cabling SAS disk shelves.

Related information

[NetApp Documentation: Disk Shelves](#)

Steps

1. *Cabling the HA interconnect (all systems except 32xx or 80xx in separate chassis)* on page 15
To cable the HA interconnect between the HA pair nodes, you must make sure that your interconnect adapter is in the correct slot. You must also connect the adapters on each node with the optical cable.
2. *Cabling the HA interconnect (32xx systems in separate chassis)* on page 16
To enable the HA interconnect between 32xx controller modules that reside in separate chassis, you must cable the onboard 10-GbE ports on one controller module to the onboard 10-GbE ports on the partner.
3. *Cabling the HA interconnect (80xx systems in separate chassis)* on page 16
To enable the HA interconnect between 80xx controller modules that reside in separate chassis, you must cable the QSFP InfiniBand ports on one I/O expansion module to the QSFP InfiniBand ports on the partner's I/O expansion module.

Cabling the HA interconnect (all systems except 32xx or 80xx in separate chassis)

To cable the HA interconnect between the HA pair nodes, you must make sure that your interconnect adapter is in the correct slot. You must also connect the adapters on each node with the optical cable.

About this task

This procedure applies to all dual-chassis HA pairs (HA pairs in which the two controller modules reside in separate chassis) except the 32xx or 80xx in separate chassis, regardless of disk shelf type.

Steps

1. Verify that your interconnect adapter is in the correct slot for your system in an HA pair.

hwu.netapp.com

For systems that use an NVRAM adapter, the NVRAM adapter functions as the HA interconnect adapter.

2. Plug one end of the optical cable into one of the local node's HA adapter ports, then plug the other end into the partner node's corresponding adapter port.

You must not cross-cable the HA interconnect adapter. Cable the local node ports only to the identical ports on the partner node.

If the system detects a cross-cabled HA interconnect, the following message appears on the system console and in the event log (accessible using the `event log show` command):

```
HA interconnect port <port> of this appliance seems to be connected to port <port> on the partner appliance.
```

3. Repeat Step 2 for the two remaining ports on the HA adapters.

Result

The nodes are connected to each other.

After you finish

Configure the system.

Cabling the HA interconnect (32xx systems in separate chassis)

To enable the HA interconnect between 32xx controller modules that reside in separate chassis, you must cable the onboard 10-GbE ports on one controller module to the onboard 10-GbE ports on the partner.

About this task

This procedure applies to 32xx systems regardless of the type of attached disk shelves.

Steps

1. Plug one end of the 10 GbE cable to the c0a port on one controller module.
2. Plug the other end of the 10 GbE cable to the c0a port on the partner controller module.
3. Repeat the preceding steps to connect the c0b ports.

You must not cross-cable the HA interconnect adapter; you must cable the local node ports only to the identical ports on the partner node.

If the system detects a cross-cabled HA interconnect, the following message appears on the system console and in the event log (accessible using the `event log show` command):

```
HA interconnect port <port> of this appliance seems to be connected to port <port> on  
the partner appliance.
```

Result

The nodes are connected to each other.

After you finish

You should configure the system.

Cabling the HA interconnect (80xx systems in separate chassis)

To enable the HA interconnect between 80xx controller modules that reside in separate chassis, you must cable the QSFP InfiniBand ports on one I/O expansion module to the QSFP InfiniBand ports on the partner's I/O expansion module.

About this task

Because the 80xx storage controller modules do not include external HA interconnect ports, you must use the HA interconnect ports on the I/O expansion modules to deploy these controller models in separate chassis.

This procedure applies to 80xx systems, regardless of the type of attached disk shelves.

Steps

1. Plug one end of the QSFP InfiniBand cable to the ib0a port on one I/O expansion module.
2. Plug the other end of the QSFP InfiniBand cable to the ib0a port on the partner's I/O expansion module.
3. Repeat the preceding steps to connect the ib0b ports.

Do not cross-cable the HA interconnect ports; cable the local node ports only to the identical ports on the partner node.

If the system detects a cross-cabled HA interconnect, the following message appears on the system console and in the event log (accessible using the `event log show` command):

```
HA interconnect port <port> of this appliance seems to be connected to port <port> on  
the partner appliance.
```

Result

The nodes are connected to each other.

After you finish

Configure the system.

Cabling a mirrored HA pair

To cable a mirrored HA pair, you must identify the ports that you need to use on each node, cable those ports, and then cable the HA interconnect.

Before you begin

You must have cabled the SAS disk shelves. The disk shelf product library has instructions for cabling SAS disk shelves.

Related information

[NetApp Documentation: Disk Shelves](#)

Steps

1. [Cabling the HA interconnect \(all systems except 32xx or 80xx in separate chassis\)](#) on page 17
To cable the HA interconnect between the HA pair nodes, you must make sure that your interconnect adapter is in the correct slot. You must also connect the adapters on each node with the optical cable.
2. [Cabling the HA interconnect \(32xx systems in separate chassis\)](#) on page 18
To enable the HA interconnect between 32xx controller modules that reside in separate chassis, you must cable the onboard 10-GbE ports on one controller module to the onboard 10-GbE ports on the partner.
3. [Cabling the HA interconnect \(80xx systems in separate chassis\)](#) on page 18
To enable the HA interconnect between 80xx controller modules that reside in separate chassis, you must cable the QSFP InfiniBand ports on one I/O expansion module to the QSFP InfiniBand ports on the partner's I/O expansion module.

Cabling the HA interconnect (all systems except 32xx or 80xx in separate chassis)

To cable the HA interconnect between the HA pair nodes, you must make sure that your interconnect adapter is in the correct slot. You must also connect the adapters on each node with the optical cable.

About this task

This procedure applies to all dual-chassis HA pairs (HA pairs in which the two controller modules reside in separate chassis) except the 32xx or 80xx in separate chassis, regardless of disk shelf type.

Steps

1. Verify that your interconnect adapter is in the correct slot for your system in an HA pair.

hww.netapp.com

For systems that use an NVRAM adapter, the NVRAM adapter functions as the HA interconnect adapter.

2. Plug one end of the optical cable into one of the local node's HA adapter ports, then plug the other end into the partner node's corresponding adapter port.

You must not cross-cable the HA interconnect adapter. Cable the local node ports only to the identical ports on the partner node.

If the system detects a cross-cabled HA interconnect, the following message appears on the system console and in the event log (accessible using the `event log show` command):

```
HA interconnect port <port> of this appliance seems to be connected to port <port> on  
the partner appliance.
```

3. Repeat Step 2 for the two remaining ports on the HA adapters.

Result

The nodes are connected to each other.

After you finish

Configure the system.

Cabling the HA interconnect (32xx systems in separate chassis)

To enable the HA interconnect between 32xx controller modules that reside in separate chassis, you must cable the onboard 10-GbE ports on one controller module to the onboard 10-GbE ports on the partner.

About this task

This procedure applies to 32xx systems regardless of the type of attached disk shelves.

Steps

1. Plug one end of the 10 GbE cable to the c0a port on one controller module.
2. Plug the other end of the 10 GbE cable to the c0a port on the partner controller module.
3. Repeat the preceding steps to connect the c0b ports.

You must not cross-cable the HA interconnect adapter; you must cable the local node ports only to the identical ports on the partner node.

If the system detects a cross-cabled HA interconnect, the following message appears on the system console and in the event log (accessible using the `event log show` command):

```
HA interconnect port <port> of this appliance seems to be connected to port <port> on  
the partner appliance.
```

Result

The nodes are connected to each other.

After you finish

You should configure the system.

Cabling the HA interconnect (80xx systems in separate chassis)

To enable the HA interconnect between 80xx controller modules that reside in separate chassis, you must cable the QSFP InfiniBand ports on one I/O expansion module to the QSFP InfiniBand ports on the partner's I/O expansion module.

About this task

Because the 80xx storage controller modules do not include external HA interconnect ports, you must use the HA interconnect ports on the I/O expansion modules to deploy these controller models in separate chassis.

This procedure applies to 80xx systems, regardless of the type of attached disk shelves.

Steps

1. Plug one end of the QSFP InfiniBand cable to the ib0a port on one I/O expansion module.

2. Plug the other end of the QSFP InfiniBand cable to the ib0a port on the partner's I/O expansion module.
3. Repeat the preceding steps to connect the ib0b ports.

Do not cross-cable the HA interconnect ports; cable the local node ports only to the identical ports on the partner node.

If the system detects a cross-cabled HA interconnect, the following message appears on the system console and in the event log (accessible using the `event log show` command):

```
HA interconnect port <port> of this appliance seems to be connected to port <port> on  
the partner appliance.
```

Result

The nodes are connected to each other.

After you finish

Configure the system.

Required connections for using uninterruptible power supplies with standard or mirrored HA pairs

You can use a UPS (uninterruptible power supply) with your HA pair. The UPS enables the system to fail over gracefully if power fails for one of the nodes, or to shut down gracefully if power fails for both nodes. You must ensure that the correct equipment is connected to the UPS.

To gain the full benefit of the UPS, you must ensure that all the required equipment is connected to the UPS. The equipment that needs to be connected depends on whether your configuration is a standard or a mirrored HA pair.

For a standard HA pair, you must connect the controller, disks, and any FC switches in use.

For a mirrored HA pair, you must connect the controller and any FC switches to the UPS, as for a standard HA pair. However, if the two sets of disk shelves have separate power sources, you do not have to connect the disks to the UPS. If power is interrupted to the local controller and disks, the controller can access the remote disks until it shuts down gracefully or the power supply is restored. In this case, if power is interrupted to both sets of disks at the same time, the HA pair cannot shut down gracefully.

Configuring an HA pair

Bringing up and configuring a standard or mirrored HA pair for the first time can require enabling HA mode capability and failover, setting options, configuring network connections, and testing the configuration.

These tasks apply to all HA pairs regardless of disk shelf type.

Steps

1. [Verifying and setting the HA state on the controller modules and chassis](#) on page 20
2. [Setting the HA mode and enabling storage failover](#) on page 21
3. [Enabling cluster HA and switchless-cluster in a two-node cluster](#) on page 22
4. [Checking for common configuration errors using Config Advisor](#) on page 23
5. [Configuring hardware-assisted takeover](#) on page 23
6. [Configuring automatic takeover](#) on page 25
7. [Configuring automatic giveback](#) on page 25
8. [Testing takeover and giveback](#) on page 28

Verifying and setting the HA state on the controller modules and chassis

For systems that use the HA state value, the value must be consistent in all components in the HA pair. You can use the Maintenance mode `ha-config` command to verify and, if necessary, set the HA state.

About this task

The `ha-config` command only applies to the local controller module and, in the case of a dual-chassis HA pair, the local chassis. To ensure consistent HA state information throughout the system, you must also run these commands on the partner controller module and chassis, if necessary.

Note: When you boot a node for the first time, the HA state value for both controller and chassis is *default*.

The HA state is recorded in the hardware PROM in the chassis and in the controller module. It must be consistent across all components of the system, as shown in the following table:

If the system or systems are in a...	The HA state is recorded on these components...	The HA state on the components must be...
Stand-alone configuration (not in an HA pair)	<ul style="list-style-type: none">• The chassis• Controller module A	non-ha
A single-chassis HA pair	<ul style="list-style-type: none">• The chassis• Controller module A• Controller module B	ha
A dual-chassis HA pair	<ul style="list-style-type: none">• Chassis A• Controller module A• Chassis B• Controller module B	ha

If the system or systems are in a...	The HA state is recorded on these components...	The HA state on the components must be...
Each single-chassis HA pair in a MetroCluster configuration	<ul style="list-style-type: none"> • The chassis • Controller module A • Controller module B 	mcc
Each dual-chassis HA pair in a MetroCluster configuration	<ul style="list-style-type: none"> • Chassis A • Controller module A • Chassis B • Controller module B 	mcc

Use the following steps to verify the HA state is appropriate and, if not, to change it:

Steps

1. Reboot or halt the current controller module and use either of the following two options to boot into Maintenance mode:
 - a. If you rebooted the controller, press Ctrl-C when prompted to display the boot menu and then select the option for **Maintenance mode boot**.
 - b. If you halted the controller, enter the following command from the LOADER prompt:

```
boot_ontap maint
```

Note: This option boots directly into Maintenance mode; you do not need to press Ctrl-C.

2. After the system boots into Maintenance mode, enter the following command to display the HA state of the local controller module and chassis:

```
ha-config show
```

The HA state should be **ha** for all components if the system is in an HA pair.

3. If necessary, enter the following command to set the HA state of the controller:

```
ha-config modify controller ha-state
```

4. If necessary, enter the following command to set the HA state of the chassis:

```
ha-config modify chassis ha-state
```

5. Exit Maintenance mode by entering the following command:

```
halt
```

6. Boot the system by entering the following command at the boot loader prompt:

```
boot_ontap
```

7. If necessary, repeat the preceding steps on the partner controller module.

Related information

[Stretch MetroCluster installation and configuration](#)

[Fabric-attached MetroCluster installation and configuration](#)

[MetroCluster management and disaster recovery](#)

Setting the HA mode and enabling storage failover

You need to set the HA mode and enable storage failover functionality to get the benefits of an HA pair.

Commands for setting the HA mode

There are specific ONTAP commands for setting the HA mode. The system must be physically configured for HA before HA mode is selected. A reboot is required to implement the mode change.

If you want to...	Use this command...
Set the mode to HA	<code>storage failover modify -mode ha -node nodename</code>
Set the mode to non-HA	<code>storage failover modify -mode non_ha -node nodename</code> Note: You must disable storage failover before disabling HA mode.

Related reference

[Description of node states displayed by storage failover show-type commands](#) on page 30
 You can use the `storage failover show`, `storage failover show-takeover`, and `storage failover show-giveback` commands to check the status of the HA pair and to troubleshoot issues.

Commands for enabling and disabling storage failover

There are specific ONTAP commands for enabling the storage failover functionality.

If you want to...	Use this command...
Enable takeover	<code>storage failover modify -enabled true -node nodename</code>
Disable takeover	<code>storage failover modify -enabled false -node nodename</code>

Enabling cluster HA and switchless-cluster in a two-node cluster

A cluster consisting of only two nodes requires special configuration settings. Cluster high availability (HA) differs from the HA provided by storage failover, and is required in a cluster if it contains only two nodes. If you have a switchless configuration, the `switchless-cluster` option must be enabled. Starting in ONTAP 9.2, detection of a switchless cluster is automatic.

About this task

Cluster HA ensures that the failure of one node does not disable the cluster. If your cluster contains only two nodes:

- Enabling cluster HA requires and automatically enables storage failover and auto-giveback.
- Cluster HA is enabled automatically when you enable storage failover.

Note: If the cluster contains or grows to more than two nodes, cluster HA is not required and is disabled automatically.

For ONTAP 9.0 and 9.1, if you have a two-node switchless configuration, the `switchless-cluster` network option must be enabled to ensure proper cluster communication between the nodes. In ONTAP 9.2, the `switchless-cluster` network option is automatically enabled. When the `detect-switchless-cluster` option is set to false, the `switchless-cluster` option will behave as it has in previous releases.

Steps

1. Enter the following command to enable cluster HA:

```
cluster ha modify -configured true
```

If storage failover is not already enabled, you are prompted to enable it.

2. ONTAP 9.0, 9.1: If you have a two-node switchless cluster, enter the following commands to verify that the switchless-cluster option is set:

a. Enter the following command to change to the advanced privilege level:

```
set -privilege advanced
```

Confirm when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

b. Enter the following command:

```
network options switchless-cluster show
```

If the output shows that the value is false, you must issue the following command:

```
network options switchless-cluster modify true
```

c. Enter the following command to return to the admin privilege level:

```
set -privilege admin
```

Related concepts

[If your cluster consists of a single HA pair](#) on page 8

Related reference

[Halting or rebooting a node without initiating takeover](#) on page 39

You can prevent an automatic storage failover takeover when you halt or reboot a node. This ability enables specific maintenance and reconfiguration operations.

Checking for common configuration errors using Config Advisor

You can use the Config Advisor tool to check for common configuration errors.

About this task

Config Advisor is a configuration validation and health check tool for NetApp systems. This tool can be deployed at both secure sites and nonsecure sites for data collection and system analysis.

Note: Support for Config Advisor is limited and is available only online.

Steps

1. Log in to the NetApp Support Site, and then navigate to **Downloads > Software > ToolChest**.
[NetApp Downloads: Config Advisor](#)
2. Click **Config Advisor**.
3. Download, install, and run Config Advisor by following the directions on the web page.
4. After running Config Advisor, review the tool's output, and follow the recommendations that are provided to address any issues that are discovered by the tool.

Configuring hardware-assisted takeover

You can configure hardware-assisted takeover to speed up takeover times. Hardware-assisted takeover uses the remote management device to quickly communicate local status changes to the partner node.

Commands for configuring hardware-assisted takeover

There are specific ONTAP commands for configuring the hardware-assisted takeover feature.

If you want to...	Use this command...
Disable or enable hardware-assisted takeover	<code>storage failover modify hwassist</code>

If you want to...	Use this command...
Set the partner address	<code>storage failover modify hwassist-partner-ip</code>
Set the partner port	<code>storage failover modify hwassist-partner-port</code>
Specify the interval between heartbeats	<code>storage failover modify hwassist-health-check-interval</code>
Specify the number of times the hardware-assisted takeover alerts are sent	<code>storage failover modify hwassist-retry-count</code>

Related concepts

[System events that trigger hardware-assisted takeover](#) on page 24

The remote management device (Service Processor) can detect many events and generate alerts. The partner node might initiate takeover, depending on the type of alert received.

Related information

[Command map for 7-Mode administrators](#)

System events that trigger hardware-assisted takeover

The remote management device (Service Processor) can detect many events and generate alerts. The partner node might initiate takeover, depending on the type of alert received.

Alert	Takeover initiated upon receipt?	Description
<code>power_loss</code>	Yes	A power loss has occurred on the node. The remote management device has a power supply that maintains power for a short period after a power loss, allowing it to report the power loss to the partner.
<code>l2_watchdog_reset</code>	Yes	The system watchdog hardware has detected an L2 reset. The remote management device detected a lack of response from the system CPU and reset the system.
<code>power_off_via_sp</code>	Yes	The remote management device was used to power off the system.
<code>power_cycle_via_sp</code>	Yes	The remote management device was used to cycle the system power off and on.
<code>reset_via_sp</code>	Yes	The remote management device was used to reset the system.
<code>abnormal_reboot</code>	No	An abnormal reboot of the node has occurred.
<code>loss_of_heartbeat</code>	No	The heartbeat message from the node was no longer received by the remote management device. Note: This alert does not refer to the heartbeat messages between the nodes in the HA pair; it refers to the heartbeat between the node and its local remote management device.
<code>periodic_message</code>	No	A periodic message has been sent during a normal hardware-assisted takeover operation.
<code>test</code>	No	A test message has been sent to verify a hardware-assisted takeover operation.

Related reference

[Commands for configuring hardware-assisted takeover](#) on page 23

There are specific ONTAP commands for configuring the hardware-assisted takeover feature.

Configuring automatic takeover

Automatic takeover is enabled by default. You can control when automatic takeovers occur by using specific commands.

Commands for controlling automatic takeover

There are specific ONTAP commands you can use to change the default behavior and control when automatic takeovers occur.

If you want takeover to occur automatically when the partner node...	Use this command...
Reboots	<code>storage failover modify -node nodename -onreboot true</code>
Panics	<code>storage failover modify -node nodename -onpanic true</code>

Related information

[Command map for 7-Mode administrators](#)

System events that result in an automatic takeover

Some events lead to an automatic takeover if storage failover is enabled.

The following system events cause an automatic and unavoidable takeover of the node.

- The node cannot send heartbeat messages to its partner due to events such as loss of power or watchdog reset.
- You halt the node without using the `-f` or `-inhibit-takeover` parameter.

If the `onpanic` parameter is set to `true`, a node panic also causes an automatic takeover. If `onpanic` is set to `false` a node panic does not cause an automatic takeover.

Configuring automatic giveback

You can configure automatic giveback so that when a node that has been taken over boots up to the Waiting for Giveback state, giveback automatically occurs.

Commands for configuring automatic giveback

There are specific ONTAP commands for enabling or disabling automatic giveback.

If you want to...	Use this command...
Enable automatic giveback so that giveback occurs as soon as the taken-over node boots, reaches the Waiting for Giveback state, and the Delay before Auto Giveback period has expired The default setting is <code>true</code> .	<code>storage failover modify -node nodename -auto-giveback true</code>

If you want to...	Use this command...
Disable automatic giveback The default setting is <code>true</code> . Note: Setting this parameter to <code>false</code> does not disable automatic giveback after takeover on panic and takeover on reboot; automatic giveback after takeover on panic must be disabled by setting the <code>-auto-giveback-after-panic</code> parameter to <code>false</code>	<pre>storage failover modify -node <i>nodename</i> -auto-giveback false</pre>
Disable automatic giveback after takeover on panic (this setting is enabled by default)	<pre>storage failover modify -node <i>nodename</i> -auto-giveback-after-panic false</pre>
Delay automatic giveback for a specified number of seconds (default is 600) This option determines the minimum time that a node will remain in takeover before performing an automatic giveback.	<pre>storage failover modify -node <i>nodename</i> -delay-seconds <i>seconds</i></pre>

Related information

[Command map for 7-Mode administrators](#)

How variations of the storage failover modify command affect automatic giveback

The operation of automatic giveback depends on how you configure the parameters of the `storage failover modify` command.

The effects of automatic giveback parameter combinations that apply to situations other than panic

The following table lists the `storage failover modify` command parameters that apply to takeover events not caused by a panic:

Parameter	Default setting
<code>-auto-giveback <i>true/false</i></code>	<code>true</code>
<code>-delay-seconds <i>integer</i> (seconds)</code>	600
<code>-onreboot <i>true/false</i></code>	<code>true</code>

The following table describes how combinations of the `-onreboot` and `-auto-giveback` parameters affect automatic giveback for takeover events not caused by a panic:

<code>storage failover modify</code> parameters used	Cause of takeover	Does automatic giveback occur?
<code>-onreboot <i>true</i></code> <code>-auto-giveback <i>true</i></code>	reboot command	Yes
	halt command, or power cycle operation issued from the Service Processor	Yes
<code>-onreboot <i>true</i></code> <code>-auto-giveback <i>false</i></code>	reboot command	Yes
	halt command, or power cycle operation issued from the Service Processor	No

storage failover modify parameters used	Cause of takeover	Does automatic giveback occur?
-onreboot <i>false</i> -auto-giveback <i>true</i>	reboot command	No
	halt command, or power cycle operation issued from the Service Processor	Yes
-onreboot <i>false</i> -auto-giveback <i>false</i>	reboot command	No
	halt command, or power cycle operation issued from the Service Processor	No

Note: If the `-onreboot` parameter is set to `true` and a takeover occurs due to a reboot, then automatic giveback is always performed, regardless of whether the `-auto-giveback` parameter is set to `true`.

When the `-onreboot` parameter is set to `false`, a takeover does not occur in the case of a node reboot. Therefore, automatic giveback cannot occur, regardless of whether the `-auto-giveback` parameter is set to `true`. A client disruption occurs.

The effects of automatic giveback parameter combinations that apply to panic situations

The following table lists the `storage failover modify` command parameters that apply to panic situations:

Parameter	Default setting
-onpanic <i>true/false</i>	<i>true</i>
-auto-giveback-after-panic <i>true/false</i> (Privilege: Advanced)	<i>true</i>
-auto-giveback <i>true/false</i>	<i>true</i>

The following table describes how parameter combinations of the `storage failover modify` command affect automatic giveback in panic situations:

storage failover parameters used	Does automatic giveback occur after panic?
-onpanic <i>true</i> -auto-giveback-after-panic <i>true</i>	Yes
-onpanic <i>false</i> -auto-giveback-after-panic <i>true</i>	No
-onpanic <i>false</i> -auto-giveback-after-panic <i>false</i>	No
-onpanic <i>true</i> -auto-giveback <i>true</i> -auto-giveback-after-panic <i>true</i>	Yes
-onpanic <i>true</i> -auto-giveback <i>true</i> -auto-giveback-after-panic <i>false</i>	Yes

storage failover parameters used	Does automatic giveback occur after panic?
-onpanic <i>true</i> -auto-giveback <i>false</i> -auto-giveback-after-panic <i>true</i>	Yes
-onpanic <i>true</i> -auto-giveback <i>false</i> -auto-giveback-after-panic <i>false</i>	No
-onpanic <i>false</i> If -onpanic is set to <i>false</i> , takeover/giveback does not occur, regardless of the value set for -auto-giveback or -auto-giveback-after-panic	No

Note: If the -onpanic parameter is set to *true*, automatic giveback is always performed if a panic occurs, *unless* you have changed the default settings for the -auto-giveback and -auto-giveback-after-panic parameters. If both of these parameters are changed from their default (*true*) to *false*, then an automatic giveback will *not* occur after a panic, even if the -onpanic parameter is set to *true*.

If the -onpanic parameter is set to *false*, takeover does not occur. Therefore, automatic giveback cannot occur, even if the -auto-giveback-after-panic parameter is set to *true*. A client disruption occurs.

Testing takeover and giveback

After you configure all aspects of your HA pair, you need to verify that it is operating as expected in maintaining uninterrupted access to both nodes' storage during takeover and giveback operations. Throughout the takeover process, the local (or takeover) node should continue serving the data normally provided by the partner node. During giveback, control and delivery of the partner's storage should return to the partner node.

Steps

1. Check the cabling on the HA interconnect cables to make sure that they are secure.
2. Verify that you can create and retrieve files on both nodes for each licensed protocol.
3. Enter the following command:

```
storage failover takeover -ofnode partner_node
```

See the man page for command details.

4. Enter either of the following commands to confirm that takeover occurred:

```
storage failover show-takeover
```

```
storage failover show
```

If you have the storage failover command's -auto-giveback option enabled:

```
cluster::> storage failover show
                Takeover
Node  Partner Possible      State Description
-----
node1 node2      -      Waiting for giveback
node2 node1     false   In takeover, Auto giveback will be
                initiated in number of seconds seconds
```

If you have the storage failover command's -auto-giveback option disabled:

```
cluster::> storage failover show
                Takeover
Node   Partner Possible      State Description
-----
node1  node2    -          Waiting for giveback
node2  node1    false      In takeover.
```

5. Enter the following command to display all the disks that belong to the partner node (Node2) that the takeover node (Node1) can detect:

```
storage disk show -home node2 -ownership
```

The following command displays all disks belonging to Node2 that Node1 can detect:

```
cluster::> storage disk show -home node2 -ownership
Disk   Aggregate Home   Owner DR Home Home ID   Owner ID   DR Home ID Reserver   Pool
-----
1.0.2  -          node2 node2 -          4078312453 4078312453 -          4078312452 Pool10
1.0.3  -          node2 node2 -          4078312453 4078312453 -          4078312452 Pool10
...
```

6. Enter the following command to confirm that the takeover node (Node1) controls the partner node's (Node2) aggregates:

```
aggr show -fields home-id,home-name,is-home
```

```
cluster::> aggr show -fields home-id,home-name,is-home
aggregate home-id   home-name is-home
-----
aggr0_1   2014942045 node1     true
aggr0_2   4078312453 node2     false
aggr1_1   2014942045 node1     true
aggr1_2   4078312453 node2     false
4 entries were displayed.
```

During takeover, the *is-home* value of the partner node's aggregates is false.

7. Give back the partner node's data service after it displays the *Waiting for giveback* message by entering the following command:

```
storage failover giveback -ofnode partner_node
```

8. Enter either of the following commands to observe the progress of the giveback operation:

```
storage failover show-giveback
storage failover show
```

9. Proceed depending on whether you saw the message that giveback was completed successfully:

If takeover and giveback...	Then...
Is completed successfully	Repeat Step 2 through Step 8 on the partner node.
Fails	Correct the takeover or giveback failure and then repeat this procedure.

Related reference

Description of node states displayed by storage failover show-type commands on page 30
 You can use the `storage failover show`, `storage failover show-takeover`, and `storage failover show-giveback` commands to check the status of the HA pair and to troubleshoot issues.

Monitoring an HA pair

You can use a variety of commands to monitor the status of the HA pair. If a takeover occurs, you can also determine what caused the takeover.

Commands for monitoring an HA pair

There are specific ONTAP commands for monitoring the HA pair.

If you want to check...	Use this command...
Whether failover is enabled or has occurred, or reasons why failover is not currently possible	<code>storage failover show</code>
Displays the nodes on which the storage failover HA-mode setting is enabled You must set the value to <code>ha</code> for the node to participate in a storage failover (HA pair) configuration. The <code>non-ha</code> value is used only in a stand-alone, or single node cluster configuration.	<code>storage failover show -mode ha</code>
Whether hardware-assisted takeover is enabled	<code>storage failover hwassist show</code>
The history of hardware-assisted takeover events that have occurred	<code>storage failover hwassist stats show</code>
The progress of a takeover operation as the partner's aggregates are moved to the node doing the takeover	<code>storage failover show-takeover</code>
The progress of a giveback operation in returning aggregates to the partner node	<code>storage failover show-giveback</code>
Whether an aggregate is home during takeover or giveback operations	<code>aggregate show -fields home-id,owner-id,home-name,owner-name,is-home</code>
Whether cluster HA is enabled (applies only to two node clusters)	<code>cluster ha show</code>
The HA state of the components of an HA pair (on systems that use the HA state)	<code>ha-config show</code> Note: This is a Maintenance mode command.

Related tasks

[Enabling cluster HA and switchless-cluster in a two-node cluster](#) on page 22

A cluster consisting of only two nodes requires special configuration settings. Cluster high availability (HA) differs from the HA provided by storage failover, and is required in a cluster if it contains only two nodes. If you have a switchless configuration, the `switchless-cluster` option must be enabled. Starting in ONTAP 9.2, detection of a switchless cluster is automatic.

Description of node states displayed by storage failover show-type commands

You can use the `storage failover show`, `storage failover show-takeover`, and `storage failover show-giveback` commands to check the status of the HA pair and to troubleshoot issues.

The following table shows the node states that the `storage failover show` command displays:

State	Meaning
Connected to <i>partner_name</i> .	The HA interconnect is active and can transmit data to the partner node.
Connected to <i>partner_name</i> , Partial giveback.	The HA interconnect is active and can transmit data to the partner node. The previous giveback to the partner node was a partial giveback, or is incomplete.
Connected to <i>partner_name</i> , Takeover of <i>partner_name</i> is not possible due to reason(s): <i>reason1, reason2,...</i>	The HA interconnect is active and can transmit data to the partner node, but takeover of the partner node is not possible. A detailed list of reasons explaining why takeover is not possible is provided in the section following this table.
Connected to <i>partner_name</i> , Partial giveback, Takeover of <i>partner_name</i> is not possible due to reason(s): <i>reason1, reason2,...</i>	The HA interconnect is active and can transmit data to the partner node, but takeover of the partner node is not possible. The previous giveback to the partner was a partial giveback.
Connected to <i>partner_name</i> , Waiting for cluster applications to come online on the local node.	The HA interconnect is active and can transmit data to the partner node and is waiting for cluster applications to come online. This waiting period can last several minutes.
Waiting for <i>partner_name</i> , Takeover of <i>partner_name</i> is not possible due to reason(s): <i>reason1, reason2,...</i>	The local node cannot exchange information with the partner node over the HA interconnect. Reasons for takeover not being possible are displayed under <i>reason1, reason2,...</i>
Waiting for <i>partner_name</i> , Partial giveback, Takeover of <i>partner_name</i> is not possible due to reason(s): <i>reason1, reason2,...</i>	The local node cannot exchange information with the partner node over the HA interconnect. The previous giveback to the partner was a partial giveback. Reasons for takeover not being possible are displayed under <i>reason1, reason2,...</i>
Pending shutdown.	The local node is shutting down. Takeover and giveback operations are disabled.
In takeover.	The local node is in takeover state and automatic giveback is disabled.
In takeover, Auto giveback will be initiated in <i>number of seconds</i> seconds.	The local node is in takeover state and automatic giveback will begin in <i>number of seconds</i> seconds.
In takeover, Auto giveback deferred.	The local node is in takeover state and an automatic giveback attempt failed because the partner node was not in waiting for giveback state.
Giveback in progress, module <i>module name</i> .	The local node is in the process of giveback to the partner node. Module <i>module name</i> is being given back. <ul style="list-style-type: none"> • Run the <code>storage failover show-giveback</code> command for more information.
Normal giveback not possible: partner missing file system disks.	The partner node is missing some of its own file system disks.

State	Meaning
Retrieving disk information. Wait a few minutes for the operation to complete, then try giveback.	The partner and takeover nodes have not yet exchanged disk inventory information. This state clears automatically.
Connected to <i>partner_name</i> , Takeover is not possible: Local node missing partner disks	<p>After a takeover or giveback operation (or in the case of MetroCluster, a disaster recovery operation including <i>switchover</i>, <i>healing</i>, or <i>switchback</i>), you might see disk inventory mismatch messages.</p> <p>If this is the case, you should wait at least five minutes for the condition to resolve before retrying the operation.</p> <p>If the condition persists, investigate possible disk or cabling issues.</p>
Connected to partner, Takeover is not possible: Storage failover mailbox disk state is invalid, Local node has encountered errors while reading the storage failover partner's mailbox disks. Local node missing partner disks	<p>After a takeover or giveback operation (or in the case of MetroCluster, a disaster recovery operation including <i>switchover</i>, <i>healing</i>, or <i>switchback</i>), you might see disk inventory mismatch messages.</p> <p>If this is the case, you should wait at least five minutes for the condition to resolve before retrying the operation.</p> <p>If the condition persists, investigate possible disk or cabling issues.</p>
Previous giveback failed in module <i>module name</i> .	<p>Giveback to the partner node by the local node failed due to an issue in <i>module name</i>.</p> <ul style="list-style-type: none"> Run the <code>storage failover show-giveback</code> command for more information.
Previous giveback failed. Auto giveback disabled due to exceeding retry counts.	Giveback to the partner node by the local node failed. Automatic giveback is disabled because of excessive retry attempts.
Takeover scheduled in <i>seconds</i> seconds.	Takeover of the partner node by the local node is scheduled due to the partner node shutting down or an operator-initiated takeover from the local node. The takeover will be initiated within the specified number of seconds.
Takeover in progress, module <i>module name</i> .	The local node is in the process of taking over the partner node. Module <i>module name</i> is being taken over.
Takeover in progress.	The local node is in the process of taking over the partner node.
<i>firmware-status</i> .	The node is not reachable and the system is trying to determine its status from firmware updates to its partner. A detailed list of possible firmware statuses is provided after this table.
Node unreachable.	The node is unreachable and its firmware status cannot be determined.
Takeover failed, reason: <i>reason</i> .	Takeover of the partner node by the local node failed due to reason <i>reason</i> .

State	Meaning
Previous giveback failed in module: <i>module name</i> . Auto giveback disabled due to exceeding retry counts.	Previously attempted giveback failed in module <i>module name</i> . Automatic giveback is disabled. <ul style="list-style-type: none"> Run the <code>storage failover show-giveback</code> command for more information.
Previous giveback failed in module: <i>module name</i> .	Previously attempted giveback failed in module <i>module name</i> . Automatic giveback is not enabled by the user. <ul style="list-style-type: none"> Run the <code>storage failover show-giveback</code> command for more information.
Connected to <i>partner_name</i> , Giveback of one or more SFO aggregates failed.	The HA interconnect is active and can transmit data to the partner node. Giveback of one or more SFO aggregates failed and the node is in partial giveback state.
Waiting for <i>partner_name</i> , Partial giveback, Giveback of one or more SFO aggregates failed.	The local node cannot exchange information with the partner node over the HA interconnect. Giveback of one or more SFO aggregates failed and the node is in partial giveback state.
Connected to <i>partner_name</i> , Giveback of SFO aggregates in progress.	The HA interconnect is active and can transmit data to the partner node. Giveback of SFO aggregates is in progress. <ul style="list-style-type: none"> Run the <code>storage failover show-giveback</code> command for more information.
Waiting for <i>partner_name</i> , Giveback of SFO aggregates in progress.	The local node cannot exchange information with the partner node over the HA interconnect. Giveback of SFO aggregates is in progress. <ul style="list-style-type: none"> Run the <code>storage failover show-giveback</code> command for more information.
Waiting for <i>partner_name</i> . Node owns aggregates belonging to another node in the cluster.	The local node cannot exchange information with the partner node over the HA interconnect, and owns aggregates that belong to the partner node.
Connected to <i>partner_name</i> , Giveback of partner spare disks pending.	The HA interconnect is active and can transmit data to the partner node. Giveback of SFO aggregates to the partner is done, but partner spare disks are still owned by the local node. <ul style="list-style-type: none"> Run the <code>storage failover show-giveback</code> command for more information.
Connected to <i>partner_name</i> , Automatic takeover disabled.	The HA interconnect is active and can transmit data to the partner node. Automatic takeover of the partner is disabled.

State	Meaning
Waiting for <i>partner_name</i> , Giveback of partner spare disks pending.	The local node cannot exchange information with the partner node over the HA interconnect. Giveback of SFO aggregates to the partner is done, but partner spare disks are still owned by the local node. <ul style="list-style-type: none"> Run the <code>storage failover show-giveback</code> command for more information.
Waiting for <i>partner_name</i> . Waiting for partner lock synchronization.	The local node cannot exchange information with the partner node over the HA interconnect, and is waiting for partner lock synchronization to occur.
Waiting for <i>partner_name</i> . Waiting for cluster applications to come online on the local node.	The local node cannot exchange information with the partner node over the HA interconnect, and is waiting for cluster applications to come online.
Takeover scheduled. <i>target node</i> relocating its SFO aggregates in preparation of takeover.	Takeover processing has started. The target node is relocating ownership of its SFO aggregates in preparation for takeover.
Takeover scheduled. <i>target node</i> has relocated its SFO aggregates in preparation of takeover.	Takeover processing has started. The target node has relocated ownership of its SFO aggregates in preparation for takeover.
Takeover scheduled. Waiting to disable background disk firmware updates on local node. A firmware update is in progress on the node.	Takeover processing has started. The system is waiting for background disk firmware update operations on the local node to complete.
Relocating SFO aggregates to <i>taking over node</i> in preparation of takeover.	The local node is relocating ownership of its SFO aggregates to the taking-over node in preparation for takeover.
Relocated SFO aggregates to <i>taking over node</i> . Waiting for <i>taking over node</i> to takeover.	Relocation of ownership of SFO aggregates from the local node to the taking-over node has completed. The system is waiting for takeover by the taking-over node.
Relocating SFO aggregates to <i>partner_name</i> . Waiting to disable background disk firmware updates on the local node. A firmware update is in progress on the node.	Relocation of ownership of SFO aggregates from the local node to the taking-over node is in progress. The system is waiting for background disk firmware update operations on the local node to complete.
Relocating SFO aggregates to <i>partner_name</i> . Waiting to disable background disk firmware updates on <i>partner_name</i> . A firmware update is in progress on the node.	Relocation of ownership of SFO aggregates from the local node to the taking-over node is in progress. The system is waiting for background disk firmware update operations on the partner node to complete.

State	Meaning
<p>Connected to <i>partner_name</i>. Previous takeover attempt was aborted because <i>reason</i>. Local node owns some of partner's SFO aggregates.</p> <p>Reissue a takeover of the partner with the "-bypass-optimization" parameter set to true to takeover remaining aggregates, or issue a giveback of the partner to return the relocated aggregates.</p>	<p>The HA interconnect is active and can transmit data to the partner node. The previous takeover attempt was aborted because of the reason displayed under <i>reason</i>. The local node owns some of its partner's SFO aggregates.</p> <ul style="list-style-type: none"> • Either reissue a takeover of the partner node, setting the -bypass-optimization parameter to true to takeover the remaining SFO aggregates, or perform a giveback of the partner to return relocated aggregates.
<p>Connected to <i>partner_name</i>. Previous takeover attempt was aborted. Local node owns some of partner's SFO aggregates.</p> <p>Reissue a takeover of the partner with the "-bypass-optimization" parameter set to true to takeover remaining aggregates, or issue a giveback of the partner to return the relocated aggregates.</p>	<p>The HA interconnect is active and can transmit data to the partner node. The previous takeover attempt was aborted. The local node owns some of its partner's SFO aggregates.</p> <ul style="list-style-type: none"> • Either reissue a takeover of the partner node, setting the -bypass-optimization parameter to true to takeover the remaining SFO aggregates, or perform a giveback of the partner to return relocated aggregates.
<p>Waiting for <i>partner_name</i>. Previous takeover attempt was aborted because <i>reason</i>. Local node owns some of partner's SFO aggregates.</p> <p>Reissue a takeover of the partner with the "-bypass-optimization" parameter set to true to takeover remaining aggregates, or issue a giveback of the partner to return the relocated aggregates.</p>	<p>The local node cannot exchange information with the partner node over the HA interconnect. The previous takeover attempt was aborted because of the reason displayed under <i>reason</i>. The local node owns some of its partner's SFO aggregates.</p> <ul style="list-style-type: none"> • Either reissue a takeover of the partner node, setting the -bypass-optimization parameter to true to takeover the remaining SFO aggregates, or perform a giveback of the partner to return relocated aggregates.
<p>Waiting for <i>partner_name</i>. Previous takeover attempt was aborted. Local node owns some of partner's SFO aggregates.</p> <p>Reissue a takeover of the partner with the "-bypass-optimization" parameter set to true to takeover remaining aggregates, or issue a giveback of the partner to return the relocated aggregates.</p>	<p>The local node cannot exchange information with the partner node over the HA interconnect. The previous takeover attempt was aborted. The local node owns some of its partner's SFO aggregates.</p> <ul style="list-style-type: none"> • Either reissue a takeover of the partner node, setting the -bypass-optimization parameter to true to takeover the remaining SFO aggregates, or perform a giveback of the partner to return relocated aggregates.
<p>Connected to <i>partner_name</i>. Previous takeover attempt was aborted because failed to disable background disk firmware update (BDFU) on local node.</p>	<p>The HA interconnect is active and can transmit data to the partner node. The previous takeover attempt was aborted because the background disk firmware update on the local node was not disabled.</p>
<p>Connected to <i>partner_name</i>. Previous takeover attempt was aborted because <i>reason</i>.</p>	<p>The HA interconnect is active and can transmit data to the partner node. The previous takeover attempt was aborted because of the reason displayed under <i>reason</i>.</p>
<p>Waiting for <i>partner_name</i>. Previous takeover attempt was aborted because <i>reason</i>.</p>	<p>The local node cannot exchange information with the partner node over the HA interconnect. The previous takeover attempt was aborted because of the reason displayed under <i>reason</i>.</p>

State	Meaning
Connected to <i>partner_name</i> . Previous takeover attempt by <i>partner_name</i> was aborted because <i>reason</i> .	The HA interconnect is active and can transmit data to the partner node. The previous takeover attempt by the partner node was aborted because of the reason displayed under <i>reason</i> .
Connected to <i>partner_name</i> . Previous takeover attempt by <i>partner_name</i> was aborted.	The HA interconnect is active and can transmit data to the partner node. The previous takeover attempt by the partner node was aborted.
Waiting for <i>partner_name</i> . Previous takeover attempt by <i>partner_name</i> was aborted because <i>reason</i> .	The local node cannot exchange information with the partner node over the HA interconnect. The previous takeover attempt by the partner node was aborted because of the reason displayed under <i>reason</i> .
Previous giveback failed in module: <i>module name</i> . Auto giveback will be initiated in <i>number of seconds</i> seconds.	<p>The previous giveback attempt failed in module <i>module_name</i>. Auto giveback will be initiated in <i>number of seconds</i> seconds.</p> <ul style="list-style-type: none"> Run the <code>storage failover show-giveback</code> command for more information.
Node owns partner's aggregates as part of the non-disruptive controller upgrade procedure.	The node owns its partner's aggregates due to the non-disruptive controller upgrade procedure currently in progress.
Connected to <i>partner_name</i> . Node owns aggregates belonging to another node in the cluster.	The HA interconnect is active and can transmit data to the partner node. The node owns aggregates belonging to another node in the cluster.
Connected to <i>partner_name</i> . Waiting for partner lock synchronization.	The HA interconnect is active and can transmit data to the partner node. The system is waiting for partner lock synchronization to complete.
Connected to <i>partner_name</i> . Waiting for cluster applications to come online on the local node.	The HA interconnect is active and can transmit data to the partner node. The system is waiting for cluster applications to come online on the local node.
Non-HA mode, reboot to use full NVRAM.	<p>Storage failover is not possible. The HA mode option is configured as <i>non_ha</i>.</p> <ul style="list-style-type: none"> You must reboot the node to use all of its NVRAM.
Non-HA mode, remove HA interconnect card from HA slot to use full NVRAM.	<p>Storage failover is not possible. The HA mode option is configured as <i>non_ha</i>.</p> <ul style="list-style-type: none"> You must move the HA interconnect card from the HA slot to use all of the node's NVRAM.
Non-HA mode, remove partner system to use full NVRAM.	<p>Storage failover is not possible. The HA mode option is configured as <i>non_ha</i>.</p> <ul style="list-style-type: none"> You must remove the partner controller from the chassis to use all of the node's NVRAM.
Non-HA mode. Reboot node to activate HA.	<p>Storage failover is not possible.</p> <ul style="list-style-type: none"> The node must be rebooted to enable HA capability.

State	Meaning
Non-HA mode. See documentation for procedure to activate HA.	<p>Storage failover is not possible. The HA mode option is configured as <i>non_ha</i>.</p> <ul style="list-style-type: none"> You must run the <code>storage failover modify -mode ha -node nodename</code> command on both nodes in the HA pair and then reboot the nodes to enable HA capability.

Possible reasons automatic takeover is not possible

If automatic takeover is not possible, the reasons are displayed in the `storage failover show` command output. The output has the following form:

Takeover of *partner_name* is not possible due to reason(s): *reason1, reason2, ...*

Possible values for *reason* are as follows:

- Automatic takeover is disabled
- Disk shelf is too hot
- Disk inventory not exchanged
- Failover partner node is booting
- Failover partner node is performing software revert
- Local node about to halt
- Local node has encountered errors while reading the storage failover partner's mailbox disks
- Local node is already in takeover state
- Local node is performing software revert
- Local node missing partner disks
- Low memory condition
- NVRAM log not synchronized
- Storage failover interconnect error
- Storage failover is disabled
- Storage failover is disabled on the partner node
- Storage failover is not initialized
- Storage failover mailbox disk state is invalid
- Storage failover mailbox disk state is uninitialized
- Storage failover mailbox version mismatch
- Takeover disabled by operator
- The size of NVRAM on each node of the HA pair is different
- The version of software running on each node of the HA pair is incompatible
- Partner node attempting to take over this node
- Partner node halted after disabling takeover
- Takeover disallowed due to unknown reason
- Waiting for partner node to recover

Possible firmware states

- Boot failed
- Booting
- Dumping core
- Dumping sparecore and ready to be taken-over
- Halted
- In power-on self test

- In takeover
- Initializing
- Operator completed
- Rebooting
- Takeover disabled
- Unknown
- Up
- Waiting
- Waiting for cluster applications to come online on the local node
- Waiting for giveback
- Waiting for operator input

Related reference

[Commands for setting the HA mode](#) on page 22

There are specific ONTAP commands for setting the HA mode. The system must be physically configured for HA before HA mode is selected. A reboot is required to implement the mode change.

Halting or rebooting a node without initiating takeover

You can prevent an automatic storage failover takeover when you halt or reboot a node. This ability enables specific maintenance and reconfiguration operations.

Commands for halting or rebooting a node without initiating takeover

Inhibiting automatic storage failover takeover when halting or rebooting a node requires specific commands. If you have a two-node cluster, you must perform additional steps to ensure continuity of service.

To prevent the partner from taking over when you...	Use this command...
Halt the node	<code>system node halt -node node -inhibit-takeover</code> If you have a two-node cluster, this command causes all data LIFs in the cluster to go offline unless you first disable cluster HA and then assign epsilon to the node that you intend to keep online.
Reboot the node Including the <code>-inhibit-takeover</code> parameter overrides the <code>-takeover-on-reboot</code> setting of the partner node to prevent it from initiating takeover.	<code>system node reboot -node node -inhibit-takeover</code> If you have a two-node cluster, this command causes all data LIFs in the cluster to go offline unless you first disable cluster HA and then assign epsilon to the node that you intend to keep online.
Reboot the node By default, a node automatically takes over for its partner if the partner reboots. You can change the <code>-onreboot</code> parameter of the <code>storage failover</code> command to change this behavior.	<code>storage failover modify -node node -onreboot false</code> Takeover can still occur if the partner exceeds the user-configurable expected time to reboot, even when the <code>-onreboot</code> parameter is set to <code>false</code> .

Related tasks

[Halting or rebooting a node without initiating takeover in a two-node cluster](#) on page 39

In a two-node cluster, cluster HA ensures that the failure of one node does not disable the cluster. If you halt or reboot a node in a two-node cluster without takeover by using the `-inhibit-takeover true` parameter, both nodes will stop serving data unless you change specific configuration settings.

Halting or rebooting a node without initiating takeover in a two-node cluster

In a two-node cluster, cluster HA ensures that the failure of one node does not disable the cluster. If you halt or reboot a node in a two-node cluster without takeover by using the `-inhibit-takeover true` parameter, both nodes will stop serving data unless you change specific configuration settings.

About this task

Before a node in a cluster configured for cluster HA is rebooted or halted using the `-inhibit-takeover true` parameter, you must first disable cluster HA and then assign epsilon to the node that you want to remain online.

Steps

1. Enter the following command to disable cluster HA:

```
cluster ha modify -configured false
```

Note that this operation does not disable storage failover.

2. Because disabling cluster HA automatically assigns epsilon to one of the two nodes, you must determine which node holds it, and if necessary, reassign it to the node that you wish to remain online.

- a. Enter the following command to change to the advanced privilege level:

```
set -privilege advanced
```

Confirm when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

- b. Determine which node holds epsilon by using the following command:

```
cluster show
```

In the following example, *Node1* holds epsilon:

```
cluster::*> cluster show
Node      Health  Eligibility  Epsilon
-----
Node1     true    true         true
Node2     true    true         false
```

If the node you wish to halt or reboot does not hold epsilon, proceed to step 3.

- c. If the node you wish to halt or reboot holds epsilon, you must remove it from the node by using the following command:

```
cluster modify -node Node1 -epsilon false
```

At this point, neither node holds epsilon.

- d. Assign epsilon to the node that you wish to remain online (in this example, *Node2*) by using the following command:

```
cluster modify -node Node2 -epsilon true
```

3. Halt or reboot and inhibit takeover of the node that does not hold epsilon (in this example, *Node2*) by using either of the following commands as appropriate:

```
system node halt -node Node2 -inhibit-takeover true
```

```
system node reboot -node Node2 -inhibit-takeover true
```

4. After the halted or rebooted node is back online, you must enable cluster HA by using the following command:

```
cluster ha modify -configured true
```

Enabling cluster HA automatically removes epsilon from both nodes.

5. Enter the following command to return to the admin privilege level:

```
set -privilege admin
```

Related tasks

[Moving epsilon for certain manually initiated takeovers](#) on page 42

You should move epsilon if you expect that any manually initiated takeovers could result in your storage system being one unexpected node failure away from a cluster-wide loss of quorum.

Performing a manual takeover

You can perform a takeover manually when maintenance is required on the partner, and in other similar situations. Depending on the state of the partner, the command you use to perform the takeover varies.

Commands for performing and monitoring manual takeovers

You can manually initiate the takeover of a node in an HA pair to perform maintenance on that node while it is still serving the data on its disks, array LUNs, or both to users.

The following table lists and describes the commands you can use when initiating a takeover:

If you want to...	Use this command...
Take over the partner node	<code>storage failover takeover</code>
Monitor the progress of the takeover as the partner's aggregates are moved to the node doing the takeover	<code>storage failover show-takeover</code>
Display the storage failover status for all nodes in the cluster	<code>storage failover show</code>
Take over the partner node without migrating LIFs	<code>storage failover takeover -skip-lif-migration-before-takeover true</code>
Take over the partner node even if there is a disk mismatch	<code>storage failover takeover -allow-disk-inventory-mismatch</code>
Take over the partner node even if there is an ONTAP version mismatch Note: This option is only used during the nondisruptive ONTAP upgrade process.	<code>storage failover takeover -option allow-version-mismatch</code>
Take over the partner node without performing aggregate relocation	<code>storage failover takeover -bypass-optimization true</code>
Take over the partner node before the partner has time to close its storage resources gracefully	<code>storage failover takeover -option immediate</code>

Note: Before you issue the storage failover command with the `immediate` option, you must migrate the data LIFs to another node by using the following command:

network interface migrate-all -node *node*

- If you specify the `storage failover takeover -option immediate` command without first migrating the data LIFs, data LIF migration from the node is significantly delayed even if the `skip-lif-migration-before-takeover` option is not specified.
- Similarly, if you specify the `immediate` option, negotiated takeover optimization is bypassed even if the `bypass-optimization` option is set to `false`.

Related information

[Disk and aggregate management](#)

Moving epsilon for certain manually initiated takeovers

You should move epsilon if you expect that any manually initiated takeovers could result in your storage system being one unexpected node failure away from a cluster-wide loss of quorum.

About this task

To perform planned maintenance, you must take over one of the nodes in an HA pair. Cluster-wide quorum must be maintained to prevent unplanned client data disruptions for the remaining nodes. In some instances, performing the takeover can result in a cluster that is one unexpected node failure away from cluster-wide loss of quorum.

This can occur if the node being taken over holds epsilon or if the node with epsilon is not healthy. To maintain a more resilient cluster, you can transfer epsilon to a healthy node that is not being taken over. Typically, this would be the HA partner.

Only healthy and eligible nodes participate in quorum voting. To maintain cluster-wide quorum, more than $N/2$ votes are required (where N represents the sum of healthy, eligible, online nodes). In clusters with an even number of online nodes, epsilon adds additional voting weight toward maintaining quorum for the node to which it is assigned.

Note: Although cluster formation voting can be modified by using the `cluster modify -eligibility false` command, you should avoid this except for situations such as restoring the node configuration or prolonged node maintenance. If you set a node as ineligible, it stops serving SAN data until the node is reset to eligible and rebooted. NAS data access to the node might also be affected when the node is ineligible.

For further information about cluster administration, quorum and epsilon, see the document library on the NetApp Support Site.

[NetApp Documentation: Product Library A-Z](#)

[System administration](#)

Steps

1. Verify the cluster state and confirm that epsilon is held by a healthy node that is not being taken over:
 - a. Change to the advanced privilege level, confirming that you want to continue when the advanced mode prompt appears (*>):

```
set -privilege advanced
```

- b. Determine which node holds epsilon:

```
cluster show
```

In the following example, Node1 holds epsilon:

```
cluster::*> cluster show
Node           Health  Eligibility  Epsilon
-----
Node1          true    true         true
Node2          true    true         false
```

If the node you want to take over does not hold epsilon, proceed to Step 4.

2. Remove epsilon from the node that you want to take over:

```
cluster modify -node Node1 -epsilon false
```
3. Assign epsilon to the partner node (in this example, Node2) by using the following command:

```
cluster modify -node Node2 -epsilon true
```
4. Perform the takeover operation:

```
storage failover takeover -ofnode node_name
```

5. Return to the admin privilege level:

```
set -privilege admin
```

Related tasks

[Halting or rebooting a node without initiating takeover in a two-node cluster](#) on page 39

In a two-node cluster, cluster HA ensures that the failure of one node does not disable the cluster.

If you halt or reboot a node in a two-node cluster without takeover by using the

`-inhibit-takeover true` parameter, both nodes will stop serving data unless you change specific configuration settings.

Related reference

[Halting or rebooting a node without initiating takeover](#) on page 39

You can prevent an automatic storage failover takeover when you halt or reboot a node. This ability enables specific maintenance and reconfiguration operations.

Performing a manual giveback

You can perform a normal giveback, a giveback in which you terminate processes on the partner node, or a forced giveback.

Note: Prior to performing a giveback, you must remove the failed drives in the taken-over system as described in the *Disks and Aggregates Power Guide*.

Related information

[Disk and aggregate management](#)

If giveback is interrupted

If the takeover node experiences a failure or a power outage during the giveback process, that process stops and the takeover node returns to takeover mode until the failure is repaired or the power is restored.

However, this depends upon the stage of giveback in which the failure occurred. If the node encountered failure or a power outage during partial giveback state (after it has given back the root aggregate), it will not return to takeover mode. Instead, the node returns to partial-giveback mode. If this occurs, complete the process by repeating the giveback operation.

If giveback is vetoed

If giveback is vetoed, you must check the EMS messages to determine the cause. Depending on the reason or reasons, you can decide whether you can safely override the vetoes.

The `storage failover show-giveback` command displays the giveback progress and shows which subsystem vetoed the giveback, if any. Soft vetoes can be overridden, while hard vetoes cannot be, even if forced. The following tables summarize the soft vetoes that should not be overridden, along with recommended workarounds.

You can review the EMS details for any giveback vetoes by using the following command:

```
event log show -node * -event gb*
```

Giveback of the root aggregate

These vetoes do not apply to aggregate relocation operations:

Vetoing subsystem module	Workaround
vfiler_low_level	Terminate the CIFS sessions causing the veto, or shutdown the CIFS application that established the open sessions. Overriding this veto might cause the application using CIFS to disconnect abruptly and lose data.
Disk Check	All failed or bypassed disks should be removed before attempting giveback. If disks are sanitizing, you should wait until the operation completes. Overriding this veto might cause an outage caused by aggregates or volumes going offline due to reservation conflicts or inaccessible disks.

Giveback of SFO aggregates

Vetoing subsystem module	Workaround
Lock Manager	Gracefully shutdown the CIFS applications that have open files, or move those volumes to a different aggregate. Overriding this veto results in loss of CIFS lock state, causing disruption and data loss.
Lock Manager NDO	Wait until the locks are mirrored. Overriding this veto causes disruption to Microsoft Hyper-V virtual machines.
RAID	Check the EMS messages to determine the cause of the veto: <ul style="list-style-type: none"> • If the veto is due to nvfile, bring the offline volumes and aggregates online. • If disk add or disk ownership reassignment operations are in progress, wait until they complete. • If the veto is due to an aggregate name or UUID conflict, troubleshoot and resolve the issue. • If the veto is due to mirror resync, mirror verify, or offline disks, the veto can be overridden and the operation restarts after giveback.
Disk Inventory	Troubleshoot to identify and resolve the cause of the problem. The destination node might be unable to see disks belonging to an aggregate being migrated. Inaccessible disks can result in inaccessible aggregates or volumes.
Volume Move Operation	Troubleshoot to identify and resolve the cause of the problem. This veto prevents the volume move operation from aborting during the important cutover phase. If the job is aborted during cutover, the volume might become inaccessible.

Related reference

[Description of node states displayed by storage failover show-type commands](#) on page 30
 You can use the `storage failover show`, `storage failover show-takeover`, and `storage failover show-giveback` commands to check the status of the HA pair and to troubleshoot issues.

Commands for performing a manual giveback

You can manually initiate a giveback on a node in an HA pair to return storage to the original owner after completing maintenance or resolving any issues that caused the takeover.

If you want to...	Use this command...
Give back storage to a partner node	<code>storage failover giveback -ofnode nodename</code>
Give back storage even if the partner is not in the waiting for giveback mode	<code>storage failover giveback -ofnode nodename -require-partner-waiting false</code> Do not use this option unless a longer client outage is acceptable.
Give back storage even if processes are vetoing the giveback operation (force the giveback)	<code>storage failover giveback -ofnode nodename -override-vetoes true</code> Use of this option can potentially lead to longer client outage, or aggregates and volumes not coming online after the giveback.

If you want to...	Use this command...
Give back only the CFO aggregates (the root aggregate)	<code>storage failover giveback -ofnode <i>nodename</i> -only-cfo-aggregates true</code>
Monitor the progress of giveback after you issue the giveback command	<code>storage failover show-giveback</code>

Related information

[Command map for 7-Mode administrators](#)

Nondisruptive operations with HA pairs

By taking advantage of an HA pair's takeover and giveback operations, you can change hardware components and perform software upgrades in your configuration without disrupting access to the system's storage.

You can perform nondisruptive operations on a system by having its partner take over the system's storage, performing maintenance, and then giving back the storage. Aggregate relocation extends the range of nondisruptive capabilities by enabling storage controller upgrade and replacement operations.

Where to find procedures for nondisruptive operations with HA pairs

An HA pair enables you to perform nondisruptive system maintenance and upgrade operations. You can refer to the specific documents for the required procedures.

The following table lists where you can find information about nondisruptive operations:

If you want to perform this task nondisruptively...	See the...
Upgrade ONTAP	Upgrade, revert, or downgrade
Replace a hardware FRU component	FRU procedures for your platform NetApp Documentation: Product Library A-Z You can find a list of all FRUs for your platform in the <i>Hardware Universe</i> . NetApp Hardware Universe

Relocating aggregate ownership within an HA pair

You can change the ownership of aggregates among the nodes in an HA pair without interrupting service from the aggregates.

Both nodes in an HA pair are physically connected to each other's disks or array LUNs. Each disk or array LUN is owned by one of the nodes. Although ownership of disks temporarily changes when a takeover occurs, the aggregate relocation operations either permanently (for example, if done for load balancing) or temporarily (for example, if done as part of takeover) change the ownership of all disks or array LUNs within an aggregate from one node to the other. The ownership changes without any data-copy processes or physical movement of the disks or array LUNs.

Relocating aggregate ownership

You can change the ownership of an aggregate only between the nodes within an HA pair.

About this task

- Because volume count limits are validated programmatically during aggregate relocation operations, it is not necessary to check for this manually.
If the volume count exceeds the supported limit, the aggregate relocation operation fails with a relevant error message.
- You should not initiate aggregate relocation when system-level operations are in progress on either the source or the destination node; likewise, you should not start these operations during the aggregate relocation.

These operations can include the following:

- Takeover
- Giveback
- Shutdown
- Another aggregate relocation operation
- Disk ownership changes
- Aggregate or volume configuration operations
- Storage controller replacement
- ONTAP upgrade
- ONTAP revert
- If you have a MetroCluster configuration, you should not initiate aggregate relocation while disaster recovery operations (*switchover*, *healing*, or *switchback*) are in progress.
- If you have a MetroCluster configuration and initiate aggregate relocation on a switched-over aggregate, the operation might fail because it exceeds the DR partner's volume limit count.
- You should not initiate aggregate relocation on aggregates that are corrupt or undergoing maintenance.
- Before initiating the aggregate relocation, you should save any core dumps on the source and destination nodes.

Steps

1. View the aggregates on the node to confirm which aggregates to move and ensure they are online and in good condition:

```
storage aggregate show -node source-node
```

The following command shows six aggregates on the four nodes in the cluster. All aggregates are online. Node1 and Node3 form an HA pair and Node2 and Node4 form an HA pair.

```
cluster::> storage aggregate show
Aggregate      Size Available Used% State #Vols  Nodes  RAID Status
-----
aggr_0        239.0GB  11.13GB  95% online      1 node1  raid_dp,
normal
aggr_1        239.0GB  11.13GB  95% online      1 node1  raid_dp,
normal
aggr_2        239.0GB  11.13GB  95% online      1 node2  raid_dp,
normal
aggr_3        239.0GB  11.13GB  95% online      1 node2  raid_dp,
normal
aggr_4        239.0GB  238.9GB   0% online      5 node3  raid_dp,
normal
aggr_5        239.0GB  239.0GB   0% online      4 node4  raid_dp,
normal
6 entries were displayed.
```

2. Issue the command to start the aggregate relocation:

```
storage aggregate relocation start -aggregate-list aggregate-1, aggregate-2... -node
source-node -destination destination-node
```

The following command moves the aggregates aggr_1 and aggr_2 from Node1 to Node3. Node3 is Node1's HA partner. The aggregates can be moved only within the HA pair.

```
cluster::> storage aggregate relocation start -aggregate-list aggr_1, aggr_2 -node
node1 -destination node3
Run the storage aggregate relocation show command to check relocation status.
node1::storage aggregate>
```

3. Monitor the progress of the aggregate relocation with the storage aggregate relocation show command:

```
storage aggregate relocation show -node source-node
```

The following command shows the progress of the aggregates that are being moved to Node3:

```
cluster::> storage aggregate relocation show -node node1
Source Aggregate Destination Relocation Status
-----
node1
    aggr_1      node3      In progress, module: waf1
    aggr_2      node3      Not attempted yet
2 entries were displayed.
node1::storage aggregate>
```

When the relocation is complete, the output of this command shows each aggregate with a relocation status of Done.

Related reference

[Commands for aggregate relocation](#) on page 49

There are specific ONTAP commands for relocating aggregate ownership within an HA pair.

Commands for aggregate relocation

There are specific ONTAP commands for relocating aggregate ownership within an HA pair.

If you want to...	Use this command...
Start the aggregate relocation process	storage aggregate relocation start
Monitor the aggregate relocation process	storage aggregate relocation show

Related tasks

[Relocating aggregate ownership](#) on page 48

You can change the ownership of an aggregate only between the nodes within an HA pair.

Related information

[*ONTAP 9 commands*](#)

Appendix: Understanding takeover and giveback

Takeover and giveback are the operations that let you take advantage of the HA configuration to perform nondisruptive operations and avoid service interruptions. Takeover is the process in which a node takes over the storage of its partner. Giveback is the process in which the storage is returned to the partner.

When takeovers occur

You can initiate takeovers manually or they can occur automatically when a failover event happens, depending on how you configure the HA pair. In some cases, takeovers occur automatically, regardless of configuration.

Takeovers can occur under the following conditions:

- When you manually initiate takeover with the `storage failover takeover` command
- When a node in an HA pair with the default configuration for immediate takeover on panic undergoes a software or system failure that leads to a panic

By default, the node automatically performs a giveback, returning the partner to normal operation after the partner has recovered from the panic and booted up.

- When a node in an HA pair undergoes a system failure (for example, a loss of power) and cannot reboot

Note: If the storage for a node also loses power at the same time, a standard takeover is not possible.

- When a node does not receive heartbeat messages from its partner
This could happen if the partner experienced a hardware or software failure that did not result in a panic but still prevented it from functioning correctly.
- When you halt one of the nodes without using the `-f` or `-inhibit-takeover true` parameter

Note: In a two-node cluster with cluster HA enabled, halting or rebooting a node using the `-inhibit-takeover true` parameter causes both nodes to stop serving data unless you first disable cluster HA and then assign epsilon to the node that you want to remain online.

- When you reboot one of the nodes without using the `-inhibit-takeover true` parameter
The `-onreboot` parameter of the `storage failover` command is enabled by default.
- When hardware-assisted takeover is enabled and it triggers a takeover when the remote management device (Service Processor) detects failure of the partner node

How hardware-assisted takeover speeds up takeover

Hardware-assisted takeover speeds up the takeover process by using a node's remote management device (Service Processor) to detect failures and quickly initiate the takeover rather than waiting for ONTAP to recognize that the partner's heartbeat has stopped.

Without hardware-assisted takeover, if a failure occurs, the partner waits until it notices that the node is no longer giving a heartbeat, confirms the loss of heartbeat, and then initiates the takeover.

The hardware-assisted takeover feature uses the following process to take advantage of the remote management device and avoid that wait:

1. The remote management device monitors the local system for certain types of failures.
2. If a failure is detected, the remote management device immediately sends an alert to the partner node.
3. Upon receiving the alert, the partner initiates takeover.

Hardware-assisted takeover is enabled by default.

How automatic takeover and giveback work

The automatic takeover and giveback operations can work together to reduce and avoid client outages.

With the default settings, if one node in the HA pair panics or reboots, the partner node automatically takes over and then automatically gives back storage when the affected node reboots. This returns the HA pair to a normal operating state.

The automatic giveback occurs by default. You can also configure the system to not perform automatic giveback if you want to control giveback impact on clients using the `storage failover modify-auto-giveback false-node < node` command. Before performing the automatic giveback (regardless of what triggered it), the partner node waits for a fixed amount of time as controlled by the `-delay-seconds` parameter of the `storage failover modify` command. The default delay is 600 seconds. By delaying the giveback, the process results in two brief outages:

1. One outage during the takeover operation
2. One outage during the giveback operation

This process avoids a single, prolonged outage that includes:

1. The time for the takeover operation
2. The time it takes for the taken-over node to boot up to the point at which it is ready for the giveback
3. The time for the giveback operation

If the automatic giveback fails for any of the non-root aggregates, the system automatically makes two additional attempts to complete the giveback.

Note: During the takeover process, the auto giveback process starts before the partner node is ready for the giveback. When the time limit of the auto giveback process expires and the partner node is still not ready, the timer restarts. As a result, the time between the partner node being ready and the actual giveback being performed might be shorter than the auto giveback time.

What happens during takeover

When a node takes over its partner, it continues to serve and update data in the partner's aggregates and volumes. To do this, the node takes ownership of the partner's aggregates, and the partner's LIFs migrate according to network interface failover rules. Except for specific SMB 3.0 connections, existing SMB (CIFS) sessions are disconnected when the takeover occurs.

The following steps occur when a node takes over its partner:

1. If the negotiated takeover is user-initiated, aggregate relocation is performed to move data aggregates one at a time from the partner node to the node that is performing the takeover. The current owner of each aggregate (except for the root aggregate) is changed from the target node to the node that is performing the takeover. There is a brief outage for each aggregate as ownership changes. This outage is briefer than an outage that occurs during a takeover without aggregate relocation.
 - You can monitor the progress using the `storage failover show-takeover` command.
 - The aggregate relocation can be avoided during this takeover instance by using the `-bypass-optimization` parameter with the `storage failover takeover` command. To bypass aggregate relocation during all future planned takeovers, set the `-bypass-takeover-optimization` parameter of the `storage failover modify` command to `true`.

Note: Aggregates are relocated serially during planned takeover operations to reduce client outage. If aggregate relocation is bypassed, longer client outage occurs during planned takeover events. Setting the `-bypass-takeover-optimization` parameter of the `storage failover modify` command to `true` is not recommended in environments that have stringent outage requirements.

2. If the user-initiated takeover is a negotiated takeover, the target node gracefully shuts down, followed by takeover of the target node's root aggregate and any aggregates that were not relocated in Step 1.
3. Before the storage takeover begins, data LIFs migrate from the target node to the node performing the takeover or to any other node in the cluster based on LIF failover rules. The LIF migration can be avoided by using the `-skip-lif-migration` parameter with the `storage failover takeover` command.

[SMB/CIFS management](#)

[NFS management](#)

[Network and LIF management](#)

4. Existing SMB (CIFS) sessions are disconnected when takeover occurs.



Attention: Due to the nature of the SMB protocol, all SMB sessions except for SMB 3.0 sessions connected to shares with the *Continuous Availability* property set, will be disruptive. SMB 1.0 and SMB 2.x sessions cannot reconnect after a takeover event. Therefore, takeover is disruptive and some data loss could occur.

5. SMB 3.0 sessions established to shares with the *Continuous Availability* property set can reconnect to the disconnected shares after a takeover event.

If your site uses SMB 3.0 connections to Microsoft Hyper-V and the *Continuous Availability* property is set on the associated shares, takeover will be nondisruptive for those sessions.

[SMB/CIFS management](#)

If the node doing the takeover panics

If the node that is performing the takeover panics within 60 seconds of initiating takeover, the following events occur:

- The node that panicked reboots.
- After it reboots, the node performs self-recovery operations and is no longer in takeover mode.
- Failover is disabled.
- If the node still owns some of the partner's aggregates, after enabling storage failover, return these aggregates to the partner using the `storage failover giveback` command.

What happens during giveback

The local node returns ownership of the aggregates and volumes to the partner node after you resolve any issues on the partner node or complete maintenance operations. In addition, the local node returns ownership when the partner node has booted up and giveback is initiated either manually or automatically.

The following process takes place in a normal giveback. In this discussion, *Node A* has taken over *Node B*. Any issues on Node B have been resolved and it is ready to resume serving data.

1. Any issues on Node B have been resolved and it displays the following message:
Waiting for giveback

2. The giveback is initiated by the `storage failover giveback` command or by automatic giveback if the system is configured for it.

This initiates the process of returning ownership of Node B's aggregates and volumes from Node A back to Node B.

3. Node A returns control of the root aggregate first.
4. Node B completes the process of booting up to its normal operating state.
5. As soon as Node B reaches the point in the boot process where it can accept the non-root aggregates, Node A returns ownership of the other aggregates, one at a time, until giveback is complete.

You can monitor the progress of the giveback with the `storage failover show-giveback` command.

Note: The `storage failover show-giveback` command does not (nor is it intended to) display information about all operations occurring during the storage failover giveback operation.

You can use the `storage failover show` command to display additional details about the current failover status of the node, such as whether the node is fully functional, whether takeover is possible, and whether giveback is complete.

I/O resumes for each aggregate once giveback is complete for that aggregate; this reduces the overall outage window for each aggregate.

HA policy and how it affects takeover and giveback operations

ONTAP automatically assigns an HA policy of CFO or SFO to an aggregate that determines how storage failover operations (takeover and giveback) occur for the aggregate and its volumes.

HA policy is assigned to and required by each aggregate on the system. The two options, CFO (controller failover), and SFO (storage failover), determine the aggregate control sequence ONTAP uses during storage failover and giveback operations.

Although the terms CFO and SFO are sometimes used informally to refer to storage failover (takeover and giveback) operations, they actually represent the HA policy assigned to the aggregates. For example, the terms *SFO aggregate* or *CFO aggregate* simply refer to the aggregate's HA policy assignment.

- Aggregates created on ONTAP systems (*except* for the root aggregate containing the root volume) have an HA policy of SFO.
Manually initiated takeover is optimized for performance by relocating SFO (non-root) aggregates serially to the partner prior to takeover. During the giveback process, aggregates are given back serially after the taken-over system boots and the management applications come online, enabling the node to receive its aggregates.
- Because aggregate relocation operations entail reassigning aggregate disk ownership and shifting control from a node to its partner, only aggregates with an HA policy of SFO are eligible for aggregate relocation.
- The root aggregate always has an HA policy of CFO and is given back at the start of the giveback operation since this is necessary to allow the taken-over system to boot.
All other aggregates are given back serially after the taken-over system completes the boot process and the management applications come online, enabling the node to receive its aggregates.

Note: Changing the HA policy of an aggregate from SFO to CFO is a Maintenance mode operation. Do not modify this setting unless directed to do so by a customer support representative.

Background disk firmware update and takeover, giveback, and aggregate relocation

Background disk firmware updates affect HA pair takeover, giveback, and aggregate relocation operations differently, depending on how those operations are initiated.

The following list describes how background disk firmware update affects takeover, giveback, and aggregate relocation:

- If a background disk firmware update is occurring on a disk on either node, manually initiated takeover operations are delayed until the disk firmware update finishes on that disk.
If the background disk firmware update takes longer than 120 seconds, takeover operations are aborted and must be restarted manually after the disk firmware update finishes. If the takeover was initiated with the `-bypass-optimization` parameter of the `storage failover takeover` command set to `true`, the background disk firmware update occurring on the destination node does not affect the takeover.
- If a background disk firmware update is occurring on a disk on the source (or takeover) node and the takeover was initiated manually with the `-options` parameter of the `storage failover takeover` command set to `immediate`, takeover operations start immediately.
- If a background disk firmware update is occurring on a disk on a node and it panics, takeover of the panicked node begins immediately.
- If a background disk firmware update is occurring on a disk on either node, giveback of data aggregates is delayed until the disk firmware update finishes on that disk.
If the background disk firmware update takes longer than 120 seconds, giveback operations are aborted and must be restarted manually after the disk firmware update completes.
- If a background disk firmware update is occurring on a disk on either node, aggregate relocation operations are delayed until the disk firmware update finishes on that disk.
If the background disk firmware update takes longer than 120 seconds, aggregate relocation operations are aborted and must be restarted manually after the disk firmware update finishes. If aggregate relocation was initiated with the `-override-destination-checks` of the `storage aggregate relocation` command set to `true`, background disk firmware update occurring on the destination node does not affect aggregate relocation.

Types of disk ownership

The HA or Disaster Recovery (DR) state of the system that owns a disk can affect which system has access to the disk. This means that there are several types of ownership for disks.

Disk ownership information is set either by ONTAP or by the administrator, and recorded on the disk, in the form of the controller module's unique system ID (obtained from a node's NVRAM card or NVMEM board).

Disk ownership information displayed by ONTAP can take one or more of the following forms. Note that the names used vary slightly depending on the context.

- Owner (or current owner)
This is the system that can currently access the disk.
- Original owner (or home owner)
If the system is in HA takeover, then owner is changed to the system that took over the node, and original owner or home owner reflects the system that owned the disk before the takeover.
- DR home owner
If the system is in a MetroCluster switchover, DR home owner reflects the value of the home owner field before the switchover occurred.

Copyright, trademark, and machine translation

Copyright

Copyright © 2021 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, 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).

Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<https://www.netapp.com/company/legal/trademarks/>

Machine translation

See important information about localized content at [netapp.com](https://www.netapp.com).

<https://www.netapp.com/company/legal/machine-translation/>