Configuring a stretch MetroCluster system with SAS disk shelves and SAS optical cables in 7-mode

SAS optical cables can be used to cable SAS disk shelves in a stretch MetroCluster system to achieve greater distance connectivity. A stretch MetroCluster system can be a new system installation, an existing system cabled with SAS cables (not using FibreBridge 6500N bridges) for which you are replacing SAS cables, hot-adding a SAS disk shelf or hot-removing a disk shelf, or an existing system for which you are replacing SAS copper cables and FibreBridge 6500N bridges.

About this document

This document contains the following procedures:

- "Installing a new system with SAS disk shelves using SAS optical cables"

- "Replacing SAS cables in a multipath HA configuration"

- "Hot-adding a SAS disk shelf using SAS optical cables"

- "Hot-removing disk shelves in systems running Data ONTAP 8.2.1 or later"

- "Replacing FibreBridge 6500N bridges and SAS copper cables with SAS optical cables"

Before you begin any procedure in this document

The following overall requirements must be met before completing any procedures in this document:

- Your system platform, disk shelves, and version of Data ONTAP that your system is running must support SAS optical cables.
  The most current support information can be found in the *NetApp Hardware Universe*.

- SAS optical multimode QSFP-to-QSFP cables can be used for controller-to-shelf and shelf-to-shelf connections, and are available in lengths up to 50 meters.

- If you are using SAS optical multimode MPO cables with MPO QSFP modules, the following rules apply:

  ◦ You can use these cables for controller-to-shelf and shelf-to-shelf connections.

  ◦ The length of a single cable cannot exceed 150 meters for OM4 and 100 meters for OM3.

  ◦ The total end-to-end path (sum of point-to-point paths from the controller to the last shelf) cannot exceed 510 meters.
    The total path includes the set of breakout cables, patch panels, and inter-panel cables.

- If you are using SAS optical multimode breakout cables, the following rules apply:

  ◦ You can use these cables for controller-to-shelf and shelf-to-shelf connections.
    If you use multimode breakout cables for a shelf-to-shelf connection, you can only use it once within a stack of disk shelves. You must use SAS optical multimode QSFP-to-QSFP or MPO cables with MPO QSFP modules to connect the remaining shelf-to-shelf connections.

  ◦ The point-to-point (QSFP-to-QSFP) path of any multimode cable cannot exceed 150 meters for OM4 and 100 meters for OM3.
    The path includes the set of breakout cables, patch panels, and inter-panel cables.

  ◦ The total end-to-end path (sum of point-to-point paths from the controller to the last shelf) cannot exceed 510 meters.

The total path includes the set of breakout cables, patch panels, and inter-panel cables.

- ◦ Up to one pair of patch panels can be used in a path.

- ◦ You need to supply the patch panels and inter-panel cables.
  The inter-panel cables must be the same mode as the SAS optical breakout cable: multimode.

- ◦ You received a set of QSFP-to-MPO cable modules with each set of SAS optical breakout cables, which you must attach to the MPO end of each SAS optical breakout cable.
  The breakout cables have SC, LC, or MTRJ connectors on the opposite end, which connect to a patch panel.

- ◦ You must connect all eight (four pairs) of the SC, LC, or MTRJ breakout connectors to the patch panel.

- • If you are using SAS optical single-mode breakout cables, the following rules apply:

  Note: SAS optical single-mode breakout cables can only be used in MetroCluster systems.

- ◦ You can use these cables for controller-to-shelf connections.
  Shelf-to-shelf connections use multimode QSFP-to-QSFP cables or multimode MPO cables with MPO QSFP modules.

- ◦ The point-to-point (QSFP-to-QSFP) path of a single single-mode cable cannot exceed 500 meters.

- ◦ The total end-to-end path (sum of point-to-point paths from the controller to the last shelf) cannot exceed 510 meters.
  The total path includes the set of breakout cables, patch panels, and inter-panel cables.

- ◦ Up to one pair of patch panels can be used in a path.

- ◦ You need to supply the patch panels and inter-panel cables.
  The inter-panel cables must be the same mode as the SAS optical breakout cable: single-mode.

- ◦ You must connect all eight (four pairs) of the SC, LC, or MTRJ breakout connectors to the patch panel.

- • The SAS cables can be SAS copper, SAS optical, or a mix depending on whether or not your system meets the requirements for using the type of cable.
  If you are using a mix of SAS copper cables and SAS optical cables, the following rules apply:

- ◦ Shelf-to-shelf connections in a stack must be all SAS copper cables or all SAS optical cables.

- ◦ If the shelf-to-shelf connections are SAS optical cables, the shelf-to-controller connections to that stack must also be SAS optical cables.

- ◦ If the shelf-to-shelf connections are SAS copper cables, the shelf-to-controller connections to that stack can be SAS optical cables or SAS copper cables.

About these procedures

The following general information applies to all procedures in this document:

- • The use of SAS optical cables in a stack attached to FibreBridge 6500N bridges is not supported.

- • Disk shelves connected with SAS optical cables require a version of disk shelf firmware that supports SAS optical cables.
  Best practice is to update all disk shelves in the storage system with the latest version of disk shelf firmware.

  Note: Do not revert disk shelf firmware to a version that does not support SAS optical cables.

- • The cable QSFP connector end connects to a disk shelf or a SAS port on a controller.
  The QSFP connectors are keyed; when oriented correctly into a SAS port the QSFP connector clicks into place and the disk shelf SAS port link LED, labeled LNK (Link Activity), illuminates green. Do not force a connector into a port.

- • The terms *node* and *controller* are used interchangeably.

- • The following documents contain the most current compatibility and best practices information:

All documents can be found on the NetApp Support Site at *mysupport.netapp.com*.

- ◦ *Interoperability Matrix Tool (IMT)*
- ◦ *Best Practices for MetroCluster Design and Implementation* TR-3548

Installing a new system with SAS disk shelves using SAS optical cables

You can install a new system that has SAS disk shelves using all SAS optical cables for shelf-to-shelf, controller-to-shelf, controller-to-patch panel, and patch panel-to-shelf connections.

Before you begin

- You must have met the requirements given at the beginning of this document.

- You ordered and received the appropriate type, number, and length of SAS optical cables required for your configuration.

- Your system must already be installed in a rack if it was not shipped in a system cabinet.

- Your system must have the appropriate number of available SAS ports on each controller.
  If you are using SAS HBAs, your system must have the appropriate number of supported SAS HBAs installed.
  Each controller requires two SAS ports for each stack to which it is connected. For example, two stacks requires four SAS ports on each controller.

- You must have determined the controller port pairs you will be cabling.
  The *Universal SAS and ACP Cabling Guide* has completed SAS port pair worksheets for common configurations, and has a SAS port pairs worksheet template if a completed worksheet is not provided for your configuration.
  SAS Disk Shelves Universal SAS and ACP Cabling Guide

- You must have checked the *Interoperability Matrix Tool* (IMT) to verify that your system meets all configuration requirements for the SAS optical cable.

- You must have checked the *Data ONTAP High Availability and MetroCluster Configuration Guide for 7-Mode* to verify that your system meets all applicable stretch MetroCluster requirements as defined in the MetroCluster installation section.

- You must have downloaded the following documents from the NetApp Support Site at *mysupport.netapp.com* (for reference later in this procedure):

  - ◦ *Universal SAS and ACP Cabling Guide*

  - ◦ *Data ONTAP High Availability and MetroCluster Configuration Guide for 7-Mode*

Steps

1. Properly ground yourself.

2. Install the platforms.

   For instructions on installing the platforms, see the *Installation and Setup Instructions* that came with your platform.

3. Install the disk shelves, power them on, and set the shelf IDs.

   For instructions, see the *Installation and Setup Instructions* that came with your disk shelf.

4. Create a port list to assign disk drives to the pools appropriately, using the information in the *Data ONTAP High Availability and MetroCluster Configuration Guide for 7-Mode*.

5. Cable the shelf-to-shelf connection (daisy-chain the disk shelves) in each stack.

   For information about daisy-chaining disk shelves, see the *Installation and Setup Instructions* that came with your disk shelf.

6. Cable the controller port pairs to each stack using your completed SAS cabling worksheet.

   You verify the SAS connections later.

7. Connect and configure the controllers following the procedure for the stretch MetroCluster configuration in the *Data ONTAP High Availability and MetroCluster Configuration Guide for 7-Mode*.

   This includes cabling the HA interconnect link as appropriate for your configuration.

8. Boot the system to Maintenance mode by completing the following substeps:

   a. Boot the system by entering the `boot_ontap` command.

   b. Halt the boot process by pressing `Ctrl-C`.

   c. Select the Maintenance mode option from the display menu.

      The system boots to Maintenance mode.

9. Verify the SAS connections by entering the following command at the Maintenance mode prompt of either controller:

   sasadmin expander_map

10. The next step depends on the output.

If...	Then...
The output lists all IOMs	The IOMs have connectivity; go to the next step.
Any IOMs are not shown (either the output does not show an IOM because it is cabled incorrectly, or the output does not show all the IOMs downstream from the incorrectly cabled IOM.)	Repeat Steps 5 and 6 to correct cabling errors, then go to Step 9.

11. Assign the attached disk shelves to the appropriate pools by using the "Assigning disk pools in a stretch MetroCluster configuration" procedure in the *Data ONTAP High Availability and MetroCluster Configuration Guide for 7-Mode*.

    Be sure to also complete the "Verifying disk paths" procedure that you are pointed to afterwards.

12. Verify that the disk shelves in the storage system have the latest version of disk shelf firmware by completing the following substeps:

    a. Enter the following command at the storage system console (from normal mode):

       sasadmin expander_map

    b. Locate the disk shelf firmware information for the disk shelves in the output.

       Example

       0151 is the disk shelf firmware version for shelf number one (Slot A/IOM A) in the storage system:
       ```
       Expanders on channel 4a:
       Level 3: WWN 500a0980000840ff, ID 1, Serial Number ' SHU0954292G114C', Product 'DS424IOM6
       ', Rev '0151', Slot A
       ```

    c. Compare the firmware information in the command output with the disk shelf firmware information at *NetApp Downloads: Disk Shelf Firmware* to determine the most current disk shelf firmware version.

13. The next step depends on how current the disk shelf firmware is.

If the firmware version in the command output is...	Then...
The same as or later than the most current version on the NetApp Support Site	Complete Steps 14 and 15. No disk shelf firmware update is needed.
An earlier version than the most current version on the NetApp Support Site	Complete Steps 14 through 16. You need to update the disk shelf firmware.

14. Configure the system and enable licenses as needed using the information about configuring an HA pair in the *Data ONTAP High Availability and MetroCluster Configuration Guide for 7-Mode*.

15. Boot the storage system and begin setup.

16. If, in Step 13, you had an earlier version of disk shelf firmware than the most current version on the NetApp Support Site, download the disk shelf firmware file by using the procedure at *NetApp Downloads: Disk Shelf Firmware*.

    You can run the commands from either node.

Nondisruptively replacing SAS cables in a multipath HA configuration

You can nondisruptively replace SAS copper and SAS optical cables in a multipath HA configuration.

Before you begin

You must have met the requirements given at the beginning of this document.

About this task

- *Replacing* a SAS cable means that you are replacing a cable using the same exact ports for a controller-to-shelf or shelf-to-shelf connection.

  Situations where you might want to replace a SAS cable can include when a cable has failed, a longer cable is needed, or SAS optical cables are preferred instead of SAS copper cables.

  Attention: If your storage system is running Data ONTAP 8.2 or earlier and your storage system is up and serving data, you cannot *move* SAS cables (change the SAS ports to which a cable is connected) nondisruptively. If you need to correct system cabling, you can use a maintenance period to do so.

- You cannot change any disks, disk shelves, or components of a controller module as part of these procedures.

Steps

1. Verify that your stretch MetroCluster system is `Multi-Path HA` by running the following command at the console of both controllers:

   sysconfig

   Note: It might take up to a minute for the system to complete discovery.

   The configuration is listed in the **System Storage** configuration field.

   Attention: If your system configuration is shown as something other than `Multi-Path HA`, you cannot continue with this procedure until you meet the procedure prerequisites.

2. If you are replacing SAS copper cables with SAS optical cables, verify that the disk shelves in the storage system have the latest version of disk shelf firmware by completing the following substeps; otherwise, go to Step 4.

   a. Enter the following command at the storage system console:

```
sasadmin expander_map
```

b. Locate the disk shelf firmware information for the disk shelves in the output.

   Example

   0151 is the disk shelf firmware version for shelf number one (Slot A/IOM A) in the storage system:
   ```
   Expanders on channel 4a:
   Level 3: WWN 500a0980000840ff, ID 1, Serial Number ' SHU0954292G114C', Product 'DS424IOM6
   ', Rev '0151', Slot A
   ```

c. Compare the firmware information in the command output with the disk shelf firmware information at *NetApp Downloads: Disk Shelf Firmware* to determine the most current disk shelf firmware version.

3. The next step depends on how current the disk shelf firmware is.

If the firmware version in the command output is...	Then...
The same as or later than the most current version on the NetApp Support Site	No disk shelf firmware update is needed.
An earlier version than the most current version on the NetApp Support Site	Download the disk shelf firmware file by using the procedure at *NetApp Downloads: Disk Shelf Firmware*. You can run the commands from either controller.

4. Replace SAS cables by completing the following substeps:

   Note: When replacing a SAS cable, wait a minimum of 10 seconds before plugging in the new cable so that the system can detect the cable change.

   You can ignore cabling messages that might appear on the console.

   a. Replace cables on Side A one cable at a time.

      The Side A cables are the cables connected to IOM A of each disk shelf.

   b. Verify that you have correctly replaced the SAS cables by entering the following command at the system console:

      ```
      sysconfig
      ```

      You can run the command from either controller.

      Note: It might take up to a minute for the system to complete discovery.

      The output should be the same as Step 1: the system should be Multi-Path HA, and the SAS port and attached disk shelf information should be the same.

      If the output is something other than Multi-Path HA, you must identify the cabling error, correct it, and run the **sysconfig** command again.

   c. Repeat substeps a and b for Side B.

      The Side B cables are the cables connected to IOM B of each disk shelf.

Configuring a stretch MetroCluster system with SAS disk shelves and SAS optical cables in 7-mode

Hot-adding a SAS disk shelf using SAS optical cables

You can use SAS optical cables to hot-add a SAS disk shelf to an existing stack of SAS disk shelves, or as a new stack to a SAS HBA or an onboard SAS port on the controller. Hot-adding a disk shelf involves installing, cabling, and verifying the disk drive and disk shelf firmware versions.

Before you begin

- You must have met the requirements given at the beginning of this document.

- You must have verified that your storage system meets the requirements for the disk shelf (and disk drives) that you are hot-adding.

- You ordered and received the appropriate type, number, and length of SAS optical cables required for your configuration.

- You have met the following requirements if you are hot-adding a single disk shelf or a stack of disk shelves directly to a system controller:

  ◦ Each controller in your storage system must have enough available PCI SAS HBA or onboard SAS ports.
    Each controller requires two SAS ports for each stack to which it is connected. For example, two stacks requires four SAS ports on each controller.

  ◦ You must have determined the controller port pairs you will be cabling.
    The *Universal SAS and ACP Cabling Guide* has completed SAS port pair worksheets for common configurations and has a SAS port pairs worksheet template if a completed worksheet is not provided for your configuration.
    SAS Disk Shelves Universal SAS and ACP Cabling Guide

- If you are hot-adding a disk shelf with SAS optical cables to a stack of disk shelves that is connected with SAS copper cables, you can temporarily have both cable types present in the stack.
  After hot-adding the disk shelf, you must replace the SAS copper cables for the rest of the shelf-to-shelf connections in the stack and the shelf-to-controller connections from the first and last disk shelf in the stack so that the stack meets the cabling rules for using SAS optical and SAS copper cables. This means that you must have ordered the appropriate number of SAS optical cables.

  > **Note:** You are instructed to replace the copper cables, which is a nondisruptive task, in the *Cabling the hot-added disk shelf* procedure.

About this task

- You can use longer cables if needed to connect the hot-added shelf.
  This is a cable replacement and can be done nondisruptively for Multipath HA systems (in Step 1 of the "Cabling the hot-added disk shelf" procedure).

- If you are hot-adding a disk shelf to an existing stack, this procedure is written for hot-adding the disk shelf to the logical last disk shelf of the stack.

- If you are hot-adding more than one disk shelf, you must hot-add one at a time.

- If you are installing the disk shelf in an equipment rack or NetApp cabinet, you must install the two-post telco tray kit or four-post rail kit that came with your disk shelf.

Installing a disk shelf for a hot-add

Installing the new SAS disk shelf involves securing the disk shelf in a rack using the applicable two-post telco tray kit or the four-post rail kit and setting the disk shelf ID.

About this task

Disk shelves do not need to be grounded; grounding is done through the power cords.

Steps

1. Properly ground yourself.

2. Install the two-post telco tray kit or the four-post rail kit for your disk shelf model using the installation flyer that came with the kit.

   Attention: If you are installing multiple disk shelves, you should install them from the bottom to the top of the rack for the best stability.

   Do not ear-mount the disk shelf into a telco-type rack; the disk shelf will collapse from the rack under its own weight.

   Attention: For two-post mid-mount installations, you must use the mid-mount brackets in addition to the two-post telco tray kit.

3. Install and secure the disk shelf onto the support brackets and rack.

   To make the disk shelf lighter and easier to maneuver, remove the power supplies and I/O modules (IOMs). Avoid removing the disk drives or carriers if possible because excessive handling can lead to internal damage.

   Caution: A fully populated DS4246 disk shelf can weigh approximately 110 lbs (49.9 kg).

   A fully populated DS2246 disk shelf can weigh approximately 49 lbs (22 kg).

4. Reinstall any power supplies and IOMs you removed to install your disk shelf into the rack.

5. Repeat Steps 3 and 4 for each disk shelf you are installing if you are adding multiple disk shelves.

6. Connect the power supplies for each disk shelf:

   a. Connect the power cords first to the disk shelves, securing them in place with the power cord retainer, and then to different power sources for resiliency.

      Note: If you have a disk shelf with four power supplies, connect power supplies in slots 1 and 3 to one power source and power supplies in slots 2 and 4 to a different power source.

   b. Turn on the power supplies for each disk shelf and wait for the disk drives to spin up.

      When the disk shelf has the maximum number of supported power supplies, all disk drives or carriers spin up at the same time. However, if one or two power supplies have faulted in a disk shelf with four power supplies, or if one power supply has faulted in a disk shelf with two power supplies, disk drives spin up in sets of six at 12-second intervals.

7. Change the shelf ID for each disk shelf you hot-added by completing the following substeps:

   You can verify IDs already in use by entering the `sasadmin shelf` command at the system console of either node.

   a. Change the shelf ID to a valid ID that is unique from the other SAS disk shelves in the storage system.

   b. Power-cycle the disk shelf to make the shelf ID take effect.

   See "Changing the disk shelf ID" in the *DS4243, DS2246, DS4486, and DS4246 Disk Shelf Installation and Service Guide* for more detailed instructions.

Cabling the hot-added disk shelf

Cabling the hot-added disk shelf involves cabling the SAS connections and, if applicable, assigning disk drive ownership.

Before you begin

- You must have met the requirements given at the beginning of this document.

- You must have met all the requirements in the "Hot-adding a SAS disk shelf using SAS optical cables" section and installed your disk shelf.

About this task

This procedure is written with the assumption that you originally cabled your system so that the controllers connect to the last disk shelf in the stack through the disk shelf's circle ports instead of the square ports.

Steps

1. Cable the SAS connections.

If you are cabling a disk shelf...	Then...
To an existing stack of disk shelves	**a.** Disconnect the SAS cable from the I/O Module (IOM) A circle port on the last shelf in the stack. You can leave the other end of the cable connected to the controller to minimize confusion, or replace the cable with a longer cable if needed. **b.** Connect (daisy-chain) the IOM A circle port of the last disk shelf in the stack to the IOM A square port of the new disk shelf, using the SAS cables that came with the new disk shelf. **c.** Connect the cable that you removed in Substep a to the IOM A circle port of the new disk shelf. **d.** Verify that all cables are securely fastened. **e.** Repeat substep a through substep d for IOM B. The storage system recognizes the new disk shelf after all the drives spin up.
To an existing SAS HBA or onboard SAS port	**a.** Use the completed SAS port pair worksheet. **b.** Verify that all cables are securely fastened.

2. Verify SAS connectivity by completing the applicable substeps:

    You can run these commands from the system console of either node.

    a. Enter the following command to find out what the adapter name is:

        sasadmin expander_map

    b. Enter the following command to verify that all disk drives can be seen by the system:

        sasadmin shell *adapter_name*

        The system displays a representation of your disk shelf populated with all the disk drives it sees.

    c. Enter the following command to verify that all IOMs (expanders) can be seen by the system (SAS channels and controller ports):

        sasadmin expander_map *adapter_name*

The following example of output from this command shows that a single expander, IOM B (slot B), in shelf 3 (ID 3) is attached to port 4a (channel 4a) on the controller:

```
Expanders on channel 4a:
Level 1: WWN 500a098000049c3f, ID 3, Serial Number 1006SZ00196, Product 'DS224IOM6 ', Rev
'0134', Slot B
```

3. Check whether your system has disk autoassignment enabled by entering the following command at the console of either node:

   `options disk.auto_assign`

   If disk autoassignment is enabled, the output shows `disk.auto_assign on`.

4. If your system does not have disk autoassignment enabled or disk drives in the same stack are owned by both nodes, assign disk drives to the appropriate pools by using the "Assigning disk pools in a stretch MetroCluster configuration" procedure in the *Data ONTAP High Availability and MetroCluster Configuration Guide for 7-Mode*

   Be sure to also complete the "Verifying disk paths" procedure that you are pointed to afterwards.

5. If you hot-added a disk shelf with SAS optical cables to a stack of disk shelves that are connected with SAS copper cables, replace the SAS copper cables for the rest of the shelf-to-shelf connections and the controller-to-shelf connections so that the stack meets the cabling rules stated at the beginning of this document; otherwise, go to the next step.

   You can use the cable replacement procedure in the "Replacing SAS Cables in a multipath HA configuration" section of this document.

6. Go to "Verifying the disk drive and disk shelf firmware versions".

Verifying the disk drive and disk shelf firmware versions

Because Data ONTAP does not always automatically update the disk drive and disk shelf firmware on hot-added SAS disk shelves, you must verify that the disk drive and disk shelf firmware are the most current versions. If they are not, you must manually update the firmware.

Steps

1. Check the console for a message containing `dbfu.selected:info` and text stating `selected for background disk firmware update` to determine whether or not you need to manually update the disk drive firmware.

   For example, actual output could look similar to `Fri Jul 19 13:05:23 PDT [svt-8040-02: bdfu.selected:info]: Disk svt-16g-sw4:4.126L64 [NETAPP X420_SFIRF300A10 NQ03] S/N [3SE0W95500009017R4SV] selected for background disk firmware.`

   After assigning disk drives on the hot-added disk shelf, the disk drive firmware updates should have begun automatically on each disk drive with downrev firmware. A repeated message similar to what is shown above appears on the console every three to five minutes—the time it takes to update downrev firmware on a disk drive—showing the firmware update progress.

If...	Then...
There is similar output	Go to the next step.
	Disk drives with downrev firmware have been detected and the firmware is being updated automatically.

If...	Then...

There is no similar output	Wait for an hourly message on the console and take the applicable action:
	
	• If there is no message on the console about disk drive firmware, this means that the disk drive firmware is current and no action is needed.
	
	• If there is a message containing `disk.fw.downrevWarning` and text stating `disks have downrev firmware that you need to update`, update the disk drive firmware by completing the following substeps:
	
	For example, actual output could look similar to `Sun May 5 04:00:01 PDT [svt-6040-01: disk.fw.downrevWarning:warning]: 1 disks have downrev firmware that you need to update.`
	
	  a. Download the disk drive firmware using the procedure at *NetApp Downloads: Disk Drive and Firmware*.
	
	  b. Enter the following command at the storage system console to update the disk drive firmware:
	
	    `disk_fw_update`
	
	  You must run this command on both controllers.
	
	    Attention: Running this command will delay I/O on the disk drives on which you are updating the firmware.

2. Verify that the disk shelf firmware is the most current version:

   a. Enter the following command at the storage system console:

      `sasadmin expander_map`

   b. Locate the disk shelf firmware information for the hot-added disk shelf in the output.

      Example

      0151 is the disk shelf firmware version for shelf number one (Slot A/IOM A) in the storage system:
      ```
      Expanders on channel 4a:
      Level 3: WWN 500a0980000840ff, ID 1, Serial Number ' SHU0954292G114C', Product 'DS424IOM6
      ', Rev '0151', Slot A
      ```

   c. Compare the firmware information in the command output with the disk shelf firmware information at *NetApp Downloads: Disk Shelf Firmware* to determine the most current disk shelf firmware version.

3. The next step depends on how current the disk shelf firmware is.

If the firmware version in the command output is...	Then...
The same as or later than the most current version on the NetApp Support Site	No disk shelf firmware update is needed.
An earlier version than the most current version on the NetApp Support Site	Download the disk shelf firmware file by using the procedure at *NetApp Downloads: Disk Shelf Firmware*. You can run the commands from either node.

Hot-removing disk shelves or stacks in systems running Data ONTAP 8.2.1 or later

If your system is running Data ONTAP 8.2.1 or later, you can hot-remove disk shelves—physically remove SAS disk shelves that have had the aggregates removed from the disk drives—in a stretch MetroCluster multipath HA configuration that is up and

serving data. You can hot-remove one or more disk shelves from anywhere within a stack of disk shelves or remove a stack of disk shelves.

Before you begin

- Your storage system must be running Data ONTAP 8.2.1 or later.

- Your storage system must be a multipath HA system.

- You must have already removed all aggregates from the disk drives in the disk shelves you are removing.

  Attention: Do not attempt this procedure unless you have already removed all aggregates from the disk shelves you are removing.

  For MetroCluster configurations, removing aggregates involves splitting the mirrored aggregates on the disk shelves you are removing and then re-creating the mirrored aggregates with another set of disk drives. For procedures, see "Splitting a mirrored aggregate" and "Rejoining split aggregates" in the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

- As a best practice, you should remove disk drive ownership after you remove the aggregates from the disk drives in the disk shelves you are removing.

  Note: This procedure follows the best practice of removing disk drive ownership; therefore, steps are written with the assumption that you have removed disk drive ownership.

  For information about removing disk drive ownership, see the "Removing ownership from a disk" procedure in the *Storage Management Guide for 7-Mode*. This document is available on the NetApp Support Site at *mysupport.netapp.com*.

  Note: The procedure for removing ownership from disk drives requires you to disable disk autoassignment. You reenable disk autoassignment when prompted at the end of this shelf hot-remove procedure.

- Multipath HA configurations cannot be in a takeover state.

- If you are removing one or more disk shelves from within a stack, you must have factored the distance to bypass the disk shelves you are removing; therefore, if the current cables are not long enough, you need to have longer cables available. The *Hardware Universe* at *hwu.netapp.com* contains information about supported SAS cables.

About this task

- This procedure follows the cabling rules in the *Universal SAS and ACP Cabling Guide*; therefore, references to IOMs and IOM SAS square and circle ports align with the cabling rules. If your multipath HA configuration is cabled differently from what is prescribed in the cabling guide, the IOM and/or IOM ports might be different.

- *Path A* refers to the A-side disk shelf I/O module (IOM A) located in the top or left of the disk shelf depending on your disk shelf model.

- *Path B* refers to the B-side disk shelf I/O module (IOM B) located in the bottom or right of the disk shelf depending on your disk shelf model.

- The *first* disk shelf in the stack is the disk shelf with the SAS square ports directly connected to the controllers.

- The *interim* disk shelf in the stack is the disk shelf directly connected to other disk shelves in the stack.

- The *last* disk shelf in the stack is the disk shelf with SAS circle ports directly connected to the controllers.

- The *next* disk shelf is the disk shelf downstream of the disk shelf being removed, in depth order.

- The *previous* disk shelf is the disk shelf upstream of the disk shelf being removed, in depth order.

- After completing this procedure for the local site, you are instructed to repeat this procedure for the remote site if you are also removing disk shelves for the remote site.

Steps

1. Verify that your system configuration is `Multi-Path HA` by entering the following command from the console of either controller:

   `sysconfig`

   It might take up to a minute for the system to complete discovery.

   The configuration is listed in the `System Storage Configuration` field.

   > **Attention:** If your system configuration is shown as something other than `Multi-Path HA`, you cannot continue with this procedure. Your system must meet the prerequisites stated in the "Before you begin" section of this procedure.

2. Verify that no aggregates are on the disk drives in the disk shelves you are removing, by completing the following substeps:

   a. Enter the following command at the console of either controller:

   `aggr status -r`

   If you have traditional volumes, you enter the **`vol status -r`** command.

   b. Check the output `SHELF` column to verify that the disk shelves from which you have removed aggregates do not appear.

   > **Attention:** If any disk drives in the disk shelves you are removing are anything other than spares or failed, you cannot continue with this procedure. Your system must meet the prerequisites stated in the "Before you begin" section of this procedure.

   Example

   The following `aggr status -r` output is for a system with disk shelf 3 being removed. Because aggregates have been removed from disk shelf 3, only disk shelves in the stack with aggregates (disk shelves 2 and 4) appear in the `SHELF` column:

   ```
   fas6200> aggr status -r

   RAID Disk   Device     HA  SHELF BAY  CHAN Pool Type  RPM  Used (MB/blks)     Phys (MB/
   blks)
   ---------   ------     --- ----- ---- ---- ---- ---- ----- --------------
   --------------
   ...
   dparity     4d.02.2    4d  2     2    SA:B  1   SAS 10000 418000/856064000
   420156/860480776
   parity      4d.02.1    4d  2     1    SA:B  1   SAS 10000 418000/856064000
   420156/860480776
   data        4d.04.3    4d  4     3    SA:B  1   SAS 10000 418000/856064000
   420156/860480768
   data        4d.04.5    4d  4     5    SA:B  1   SAS 10000 418000/856064000
   420156/860480768
   ```

3. Verify that ownership is removed from the disk drives on the disk shelves you are removing:

   a. Enter the following command at the console of either controller:

   `disk show -v`

   b. Check the output and take the appropriate action:

If the output shows...	Then...
`Not Owned` or `FAILED`	Go to the next step.
	The disk drives in the disk shelves you are removing are in the correct state.

If the output shows...	Then...
Any disk drives in the disk shelves you are removing have ownership	You can use the "Removing ownership from a disk" procedure referenced in the "Before you begin" section of this procedure.

Example

The following output for the `disk show -v` command shows that disk drives on the disk shelf being removed (disk shelf 3) are in a correct state for removal. Two disk drives have the ownership removed; therefore, `Not Owned` appears in the `OWNER` column. Two disk drives are failed; therefore, `FAILED` appears in the `POOL` column:

```
fas6200> disk show -v

DISK        OWNER                       POOL    SERIAL NUMBER         DR
HOME                        CHKSUM
---------   -------------               -----   -------------
-------------               -------
...
4b.05.5     Not Owned                   NONE    6XQ1FA1V0000N250MAQ5  de-
f3250-136-99(1574897553)    Block
4b.05.11    Not Owned                   NONE    6XQ1F9ZN0000B242DQPB  de-
f3250-136-99(1574897553)    Block
4a.01.15    de-f3250-136-97(1574937479) FAILED  S143NEACB00066        de-
f3250-136-97(1574937479)    Block
4a.01.13    de-f3250-136-97(1574937479) FAILED  S143NEACB00115        de-
f3250-136-97(1574937479)    Block
...
```

4. Turn on the LEDs for each disk shelf you are removing so that the disk shelves are physically identifiable, by entering the following command for each disk shelf you are removing:

storage shelf identify *adapter_name*.shelf*shelf_ID* on

You must be in advanced privilege level to enter this command. You can enter the command from the console of either controller.

All LEDs on the front and back of the specified disk shelves flash. LEDs continue to flash for up to five minutes.

Note: You need to be certain of which disk shelves you are removing so that you can correctly recable path A and path B later in this procedure. Even if you are removing an entire stack of disk shelves, it is recommended that you use this command.

Example

To physically identify disk shelf 3, you enter **storage shelf identify *0b.shelf3* on**.

5. If you are removing an entire stack of disk shelves, complete the following substeps; otherwise, go to the next step:

   a. Remove all SAS cables on path A and path B.

   This includes controller-to-shelf cables and shelf-to-shelf cables for all disk shelves in the stack you are removing.

   b. Go to Step 9.

6. If you are removing one or more disk shelves from a stack (but keeping the stack), recable the applicable path A stack connections to bypass the disk shelves you are removing by completing the applicable set of substeps:

If you are removing more than one disk shelf, complete the applicable set of substeps one disk shelf at a time.

If you are removing...	Then...
The first disk shelf in a stack	a. Remove the cable connecting the IOM A SAS circle port of the first disk shelf and the IOM A SAS square port of the second disk shelf in the stack and set it aside. b. Move the cable connecting the controller to the IOM A SAS square port of the first disk shelf to the IOM A SAS square port of the second disk shelf in the stack
An interim disk shelf in a stack	a. Remove the cable connecting the IOM A SAS circle port of the disk shelf being removed and the IOM A SAS square port of the next disk shelf in the stack and set it aside. b. Move the cable connecting the IOM A SAS square port of the disk shelf being removed to the IOM A SAS square port of the next disk shelf in the stack
The last disk shelf in a stack	a. Remove the cable connecting the IOM A SAS square port of the last disk shelf and the IOM A SAS circle port of the previous disk shelf in the stack and set it aside. b. Move the cable connecting the controller to the IOM A SAS circle port of the last disk shelf to the IOM A SAS circle port of the previous disk shelf in the stack

7. Verify that the cabling on path A has successfully bypassed the disk shelves you are removing and all disk drives on the disk shelves you are removing are still connected through path B, by entering the following command from the console of either controller:

storage show disk -p

It might take up to a minute for the system to complete discovery.

Example

In this example of how the output should appear for a stretch MetroCluster system, the disk shelf being removed is disk shelf 3. One line item appears for each disk drive connected through path B (now the primary path); therefore, the disk drives are listed in the PRIMARY column and B appears in the first PORT column. There is no connectivity through path A for any of the disk drives in the disk shelf being removed; therefore, no information is shown in the SECONDARY or second PORT columns:

```
fas6200> storage show disk -p

PRIMARY    PORT  SECONDARY PORT SHELF BAY
--------   ----  --------- ---- ---------
...
1a.03.0    B                    3     0
1a.03.1    B                    3     1
1a.03.2    B                    3     2
1a.03.3    B                    3     3
1a.03.4    B                    3     4
1a.03.5    B                    3     5
1a.03.6    B                    3     6
1a.03.7    B                    3     7
1a.03.8    B                    3     8
1a.03.9    B                    3     9
1a.03.10   B                    3     10
1a.03.11   B                    3     11
1a.03.12   B                    3     12
1a.03.13   B                    3     13
1a.03.14   B                    3     14
...
```

Attention: If the output shows anything other than all the disk drives connected only through path B, you must correct the cabling by repeating Step 6.

8. Complete the following substeps:

a. Repeat Step 6 and Step 7 for path B.

b. Repeat Step 1 to confirm that your system configuration is the same as before you began this procedure.

c. Go to the next step.

9. If, when you removed ownership from the disk drives as part of the preparation for this procedure, you disabled disk autoassignment, then reenable disk autoassignment by entering the following command; otherwise, go to the next step:

   `options disk.auto_assign on`

   Enter the applicable command from the console of each controller.

10. Power off the disk shelves you disconnected and unplug the power cords from the disk shelves.

11. Remove the disk shelves from the rack or cabinet.

    To make disk shelves lighter and easier to maneuver, remove the power supplies and IOMs. Avoid removing the disk drives if possible, because excessive handling can lead to internal damage.

    It is recommended that you use a mechanical hoist or lift if you are moving a fully populated 4U disk shelf. For example, a fully populated DS4246 disk shelf can weigh approximately 110 lbs (49.9 kg).

12. If you are removing disk shelves at the remote site, repeat this procedure.

Replacing FibreBridge 6500N bridges and SAS copper cables with SAS optical cables

You must halt both controllers of the stretch MetroCluster to replace all FibreBridge 6500N bridges, shelf-to-shelf SAS copper cables, and bridge-to-shelf SAS copper cables with SAS optical cables.

Before you begin

- You must have met the requirements given at the beginning of this document.

- You have ordered and received the appropriate type, number, and length of SAS optical cables required for your configuration.

- You have ordered and received the appropriate number and type of SAS HBAs for each controller.
  Each controller requires two SAS ports for each stack to which it is connected. For example, one stack at Site 1 and one stack as Site 2 requires four ports on each controller.
  You install the SAS HBAs in Step 8 of this procedure, after you halt your system.

- You must have downloaded the *Universal SAS and ACP Cabling Guide* from the NetApp Support Site at *mysupport.netapp.com*.

About this task

- Commands are entered at the console of either controller unless otherwise noted.

- This procedure assumes that you are replacing bridges and SAS copper cables on all the stacks on your system.
  You replace bridges and cables one stack at a time (by repeating step 10 through 12).

- You cannot change any disks, disk shelves, or components of a controller module as part of this procedure.

- You can use the following illustration of a stretch MetroCluster system with FibreBridge 6500N bridges and SAS copper cables for reference when replacing the bridges and SAS copper cables:
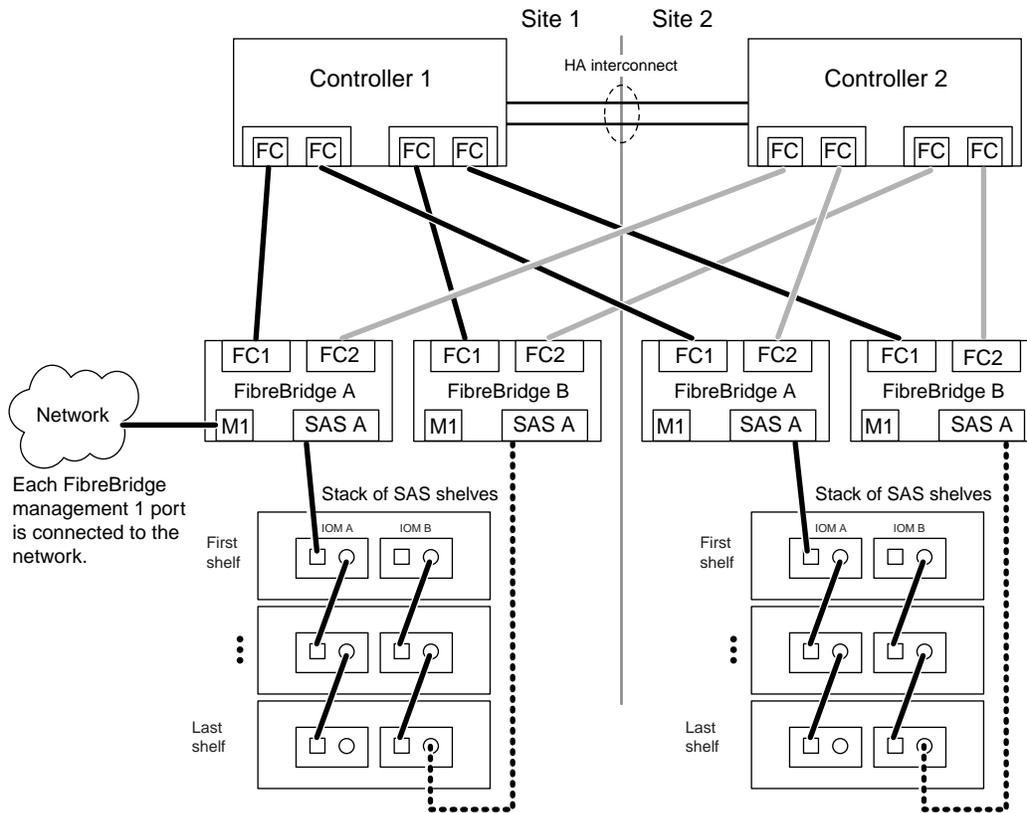
The following illustration is an example of what a 62xx looks like after cabling the system with SAS optical cables (having replaced the FibreBridge 6500N bridges and SAS copper cables):
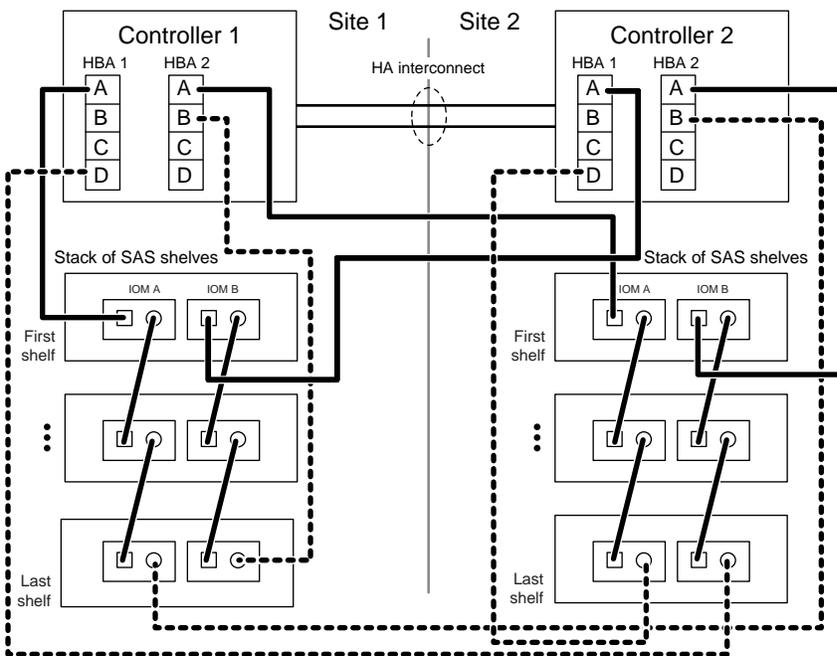
Steps

1. Determine the SAS port pairs to cable your system.

The *Universal SAS and ACP Cabling Guide* has completed SAS port pair worksheets for common configurations and has a SAS port pairs worksheet template if a completed worksheet is not provided for your configuration.

SAS Disk Shelves Universal SAS and ACP Cabling Guide

2. Verify that your system is `Multi-Path HA` by running the following command at the console of both nodes:

   sysconfig

   Note: It might take up to a minute for the system to complete discovery.

   The configuration is listed in the **System Storage** configuration field. It should be the fourth line of output.

   Caution: If your system configuration is shown as something other than `Multi-Path HA`, identify the cabling issue and correct it before continuing with this procedure.

3. Verify that the disk shelves in the storage system have the latest version of disk shelf firmware by completing the following substeps:

   a. Enter the following command at the storage system console:

      sysconfig -v

   b. Locate the disk shelf firmware information for the disk shelves in the output.

      Example

      0151 is the disk shelf firmware version for shelf number one (for each IOM6) in the storage system:
      `Shelf 1: IOM6 Firmware rev. IOM6 A: 0151 IOM6 B: 0151`

   c. Compare the firmware information in the command output with the disk shelf firmware information at *NetApp Downloads: Disk Shelf Firmware* to determine the most current disk shelf firmware version.

4. The next step depends on how current the disk shelf firmware is.

If the firmware version in the command output is...	Then...
The same as or later than the most current version on the NetApp Support Site	No disk shelf firmware update is needed.
An earlier version than the most current version on the NetApp Support Site	Download the disk shelf firmware file by using the procedure at *NetApp Downloads: Disk Shelf Firmware*. You can run the commands from either node.

5. Check the status of both nodes by entering the following command at the system console of either node:

   cf status

6. Take one of the following actions, depending on the result of the `cf status` command:

If...	Then...
Neither node is in takeover mode	Go to Step 7 in this procedure.
One of the nodes is in takeover mode	a. Correct the problem that caused the takeover. b. Enter the **cf giveback** command from the target node console. c. Go back to the beginning of this procedure.

Configuring a stretch MetroCluster system with SAS disk shelves and SAS optical cables in 7-mode

7. Disable controller failover by entering the following command from either node:

   cf disable

8. If you ordered additional HBAs, install them at this time; otherwise, go to the next step.

9. Enter the following command from the system console to perform a clean shutdown:

   halt

10. Select one of the disk shelf stacks and remove the cabling by completing the following substeps:

    You can begin on any stack.

    a. Remove the controller-to-bridge FC cables.

       You are removing a total of four FC cables, two from each controller.

          Note: The best practice is to remove the cables from the controller ports first.

    b. Remove the M1 port-to-bridge cables.

       You are removing a total of two cables, one from each bridge to the network.

    c. Remove the bridge-to-stack SAS copper cables.

       You are removing a total of two SAS copper cables, one from the first shelf in the stack and one from the last shelf in the stack.

    d. If needed, you can remove the bridges from the rack.

11. Cable the controller-to-stack connections with SAS optical cables by completing the following substeps:

    Use the cabling worksheet you completed in Step 1 so you know which port pairs to use.

    a. Connect the first shelf in the stack to each controller.

    b. Connect the last shelf in the stack to each controller.

    The stack of disk shelves is now connected to both controllers with SAS optical cables.

12. Replace the stack shelf-to-shelf SAS copper cables with SAS optical cables.

13. Repeat Step 10 through 12 for each remaining stack; and then proceed to the next step.

14. Boot the nodes by entering the following command on either node:

    boot_ontap

15. Verify that you correctly replaced the SAS cables by entering the following command on the console of either node:

    sysconfig

    The output should be the same as Step 2: The system should be Multi-Path HA. However, the SAS ports and attached disk shelf information will have changed because disk shelves were moved from FC ports (connected through the bridges) to SAS ports on the controllers.

16. Enable controller failover by entering the following command on either node:

    cf enable

17. Verify that controller failover is enabled by entering the following command on either node:

    cf status

Copyright information

Copyright © 1994–2015 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

Trademark information

NetApp, the NetApp logo, Go Further, Faster, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. A current list of NetApp trademarks is available on the web at *http://www.netapp.com/us/legal/netapptmlist.aspx*.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

How to send comments about documentation and receive update notification

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email to *doccomments@netapp.com*. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.

- Telephone: +1 (408) 822-6000

- Fax: +1 (408) 822-4501

- Support telephone: +1 (888) 463-8277