Replacing a chassis

To replace the chassis, you must remove the power supplies, fan modules, and controller module or modules and I/O expansion module, when present, from the old chassis, remove the old chassis from the equipment rack or system cabinet, install the new chassis of the same model as the old chassis, and then reinstall the components into the new chassis.

Before you begin

All other components in the system must be functioning properly; if not, you must contact technical support.

About this task

- You can use this procedure with all versions of Data ONTAP supported by your system.

- This procedure is written with the assumption that you are moving the controller module or modules to the new chassis, and that the chassis is a new component from NetApp.

Steps

Shutting down the target node or nodes

You shut down or take over the target nodes using different procedures, depending on the storage system hardware configuration.

Shutting down the nodes

You must shut down the node or nodes in the chassis prior to moving them to the new chassis.

About this task

If you cannot gracefully shut down the nodes due to system constraints such as backplane failure, call technical support for additional guidance.

Steps

1. Change to the advance privilege level:

```
set -privilege advanced
```

a. How you proceed depends on where epsilon is assigned:

- If the target chassis does not contain the epsilon node in it, then do to the next step.

- If the target chassis does contain the epsilon node in it, remove and

2. Display the status of the nodes:

```
cluster show -epsilon *
```

You will see a screen similar to the following:

```
Node          Health  Eligibility   Epsilon
------------  ------- ------------  ------------
node1         true    true          true
node2         true    true          false
node3         true    true          false
node4         true    true          false

4 entries were displayed.
```

Note: Epsilon must not be on a node in the chassis to be replaced.

3. If epsilon is assigned to a node in the chassis you are replacing, reassign it to a different node in a chassis that is healthy:

a. Remove epsilon from the node:

```
cluster modify -node node1 -epsilon false
```

b. Assign epsilon to a node in a healthy chassis:

```
cluster modify -node node4 -epsilon true
```

c. Display epsilon for the cluster:

```
cluster show -epsilon *
```

Clustered Data ONTAP 8.3 System Administration Guide

4. If your system has two controller modules, you must disable the HA pair with one of the following commands:

If your system is configured as a...	Then issue this command...
System operating in 7-Mode	`cf disable`
System running clustered Data ONTAP with two nodes in the cluster	`node::> cluster ha modify -configured false` `node::> storage failover modify -node node0 -enabled false`
System running clustered Data ONTAP with more than two nodes in the cluster	`node::> storage failover modify -node node0 -enabled false`

5. Halt the node:

If your system is configured as a...	Then issue this command...
System operating in 7-Mode	`halt`
System running clustered Data ONTAP	`system node halt -node <node_name>`

6. Repeat the preceding step for the second controller module, if present.

7. If you are not already grounded, properly ground yourself.

8. Turn off the power supplies and unplug both power cords from the power source.

Shutting down a controller module in a two-node MetroCluster configuration

To shut down a controller module, you must determine the status of the impaired node and, if necessary, perform a MetroCluster switchover operation from the healthy controller module so that the healthy node continues to serve data from the impaired node's storage.

Steps

1. If the system is running clustered Data ONTAP, check the status of the nodes in the cluster:

   a. Enter the following command at the system console of either node:

   cluster show

   The command produces output similar to the following:

   ```
   Node          Health  Eligibility
   ------------ ------- -----------
   node1         true    true
   node2         true    true
   node3         true    true
   node4         true    true

   4 entries were displayed.
   ```

   b. Take one of the following actions, depending on the result of the command:

If...	Then...
All nodes show `true` for both health and eligibility.	Proceed to Step 3.
The impaired node shows `false` for health.	Proceed to the next step.
Any nodes show `false` for eligibility.	Resolve any cluster issues as needed before continuing with this procedure.
Any nodes other than the impaired node show `false` for health.	Correct the problems that cause the health issues on the nodes before continuing with this procedure.

2. Use the `metrocluster check run`, `metrocluster check show` and `metrocluster check config-replication show` commands to make sure no configuration updates are in progress or pending.

3. If the impaired node has not switched over, perform the switchover operation from the healthy node:

   metrocluster switchover

4. Monitor the completion of the switchover:

   metrocluster operation show

   Example

   ```
   mcc1A::*> metrocluster operation show
     Operation: Switchover
    Start time: 10/4/2012 19:04:13
         State: in-progress
      End time: -
        Errors:
   ```

```
mcc1A::*> metrocluster operation show
  Operation: Switchover
 Start time: 10/4/2012 19:04:13
      State: successful
   End time: 10/4/2012 19:04:22
     Errors: -
```

5. Shut down the impaired node.

   The method you use to shut down the node depends on whether remote management using a Service Processor (SP) is used:

Is the SP configured?	Then...
Yes	Log in to the impaired node's SP and issue the following command: **system power off**
No	At the impaired node's prompt, press **Ctrl-C** and respond **Y** to halt the node.

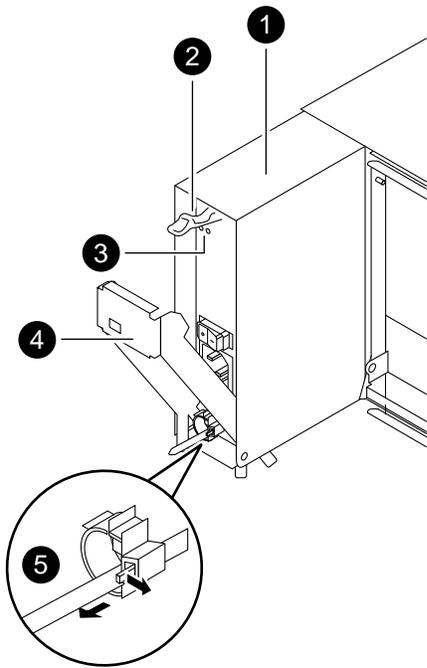6. If you are not already grounded, properly ground yourself.

Removing the power supplies

To remove the power supplies from the chassis, you must perform a specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.

2. If you have not already done so, turn off the power to the target power supply, open the power cord locking mechanism, and then remove the power cord from the power supply and the power source.

3. Press the release latch down on the power supply cam handle, lower the cam handle to the fully open position, then slide the power supply out of the chassis, and then set it aside.

   Make sure that you support the power supply with your free hand.

❶	Power supply
❷	Cam handle release latch
❸	Power and Attention LEDs
❹	Cam handle
❺	Power cord locking mechanism

4. Repeat the preceding steps for any remaining power supplies.

# Removing the controller module and/or I/O expansion module

To remove the controller module and if applicable, the I/O expansion module from the chassis, you must perform a specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.

2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, and keep track of where the cables were connected.

Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove the cable management device from the controller module and the I/O expansion module, and then set them both aside.

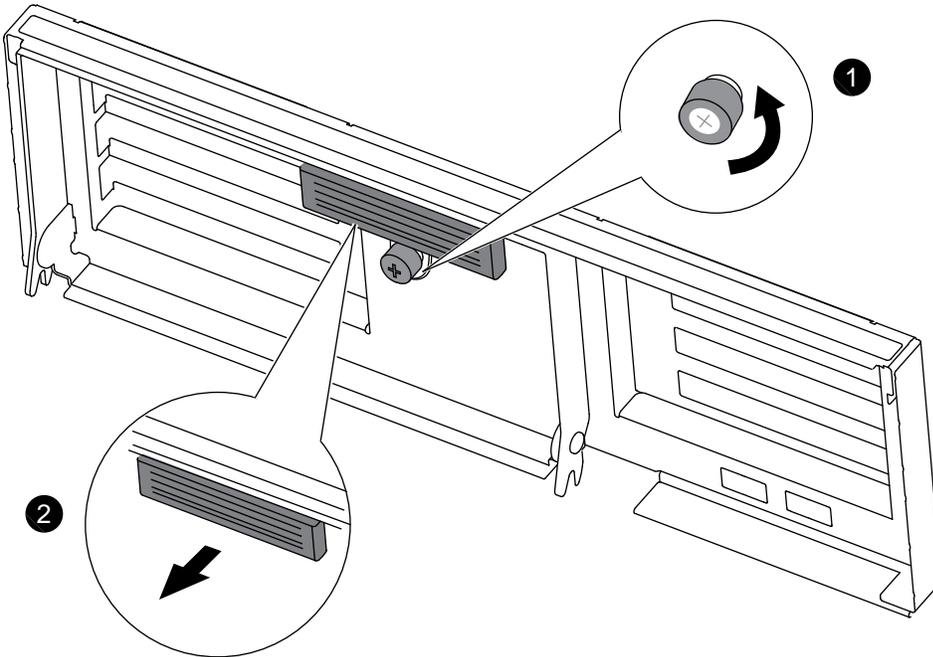**4.** Loosen the thumbscrew on the cam handle of the controller module and the I/O expansion module.
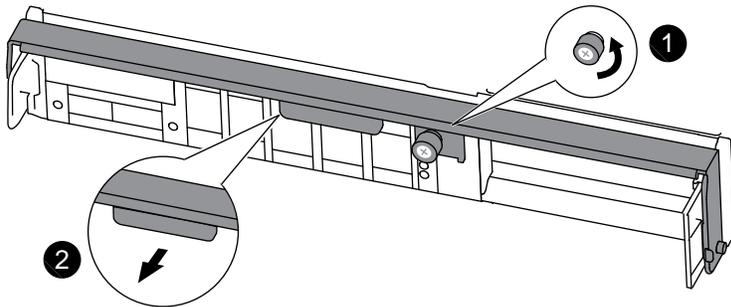
8020 controller module cam handle

❶	Thumbscrew
❷	Cam handle
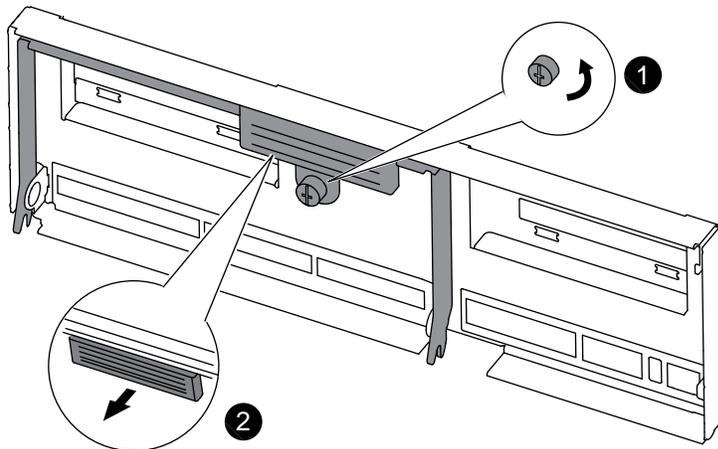
Any other 80xx controller module cam handle

❶	Thumbscrew
❷	Cam handle

I/O expansion module cam handle

❶	Thumbscrew
❷	Cam handle

5. Pull the cam handle downward and begin to slide the controller module or I/O expansion module out of the chassis.

If you have...	Then...
A 8020 system	Slide the controller module completely out of the system. **Note:** Make sure that you support the bottom of the controller module with your free hand and set it aside.
Any other 80xx model	Slide the controller module out of the system until it catches, press the release latch on the left side of the controller module, and then slide the controller module out of the system and set it aside. **Note:** Make sure that you support the bottom of the controller module with your free hand.
A system with an I/O expansion module	Slide the I/O expansion module completely out of the system and set it aside. **Note:** Make sure that you support the bottom of the module with your free hand.

6. Repeat the preceding steps if you need to remove the second controller module, or if you need to remove the blanking panel.
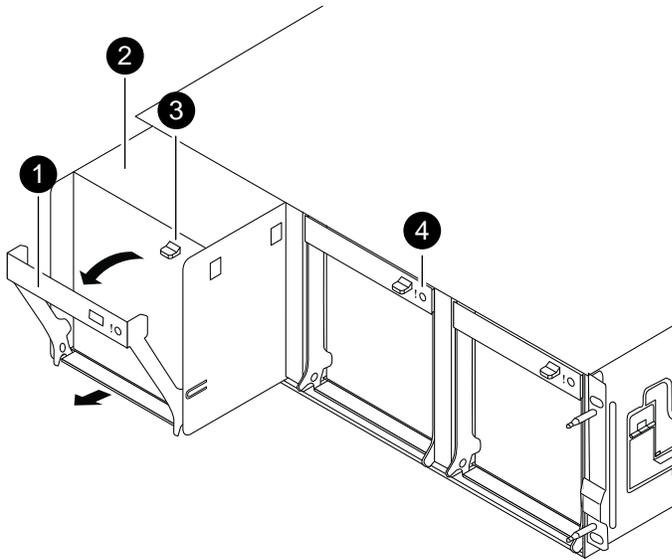
Removing the fans

To remove the fans from the chassis , you must perform a specific sequence of steps.
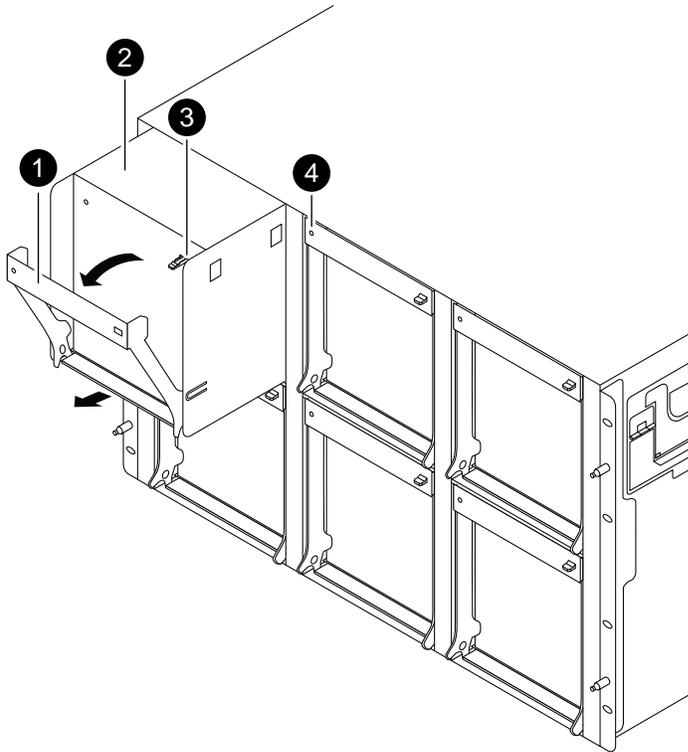
Steps

1. If you are not already grounded, properly ground yourself.

2. Remove the bezel from the front of the system.

3. Press down the release latch on the fan module cam handle and pull the cam handle down.

   The fan module moves a little bit away from the chassis.

8020 system fan modules

❶	Cam handle
❷	Fan module
❸	Cam handle release latch
❹	Fan module Attention LED

Any other 80xx system fan modules

1	Cam handle
2	Fan module
3	Cam handle release latch
4	Fan module Attention LED

4. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.

   Caution: The fan modules are short. Always support the bottom of the module with your free hand so that it does not suddenly drop free from the chassis and injure you.

5. Set the fan module aside.

6. Repeat the preceding steps for any remaining fan modules.

Removing the chassis from the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the new one.

Steps

1. Remove the screws from the chassis mount points.

   Note: If the system is in a system cabinet, you might need to remove the rear tie-down bracket.

2. With the help of two or three people, slide the empty chassis off the rack rails and set it aside.

Installing the new chassis into the equipment rack or system cabinet

To install the new chassis in the equipment rack or system cabinet, you must perform a specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.

2. Using two or three people, install the empty chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or *L* brackets in an equipment rack.

3. Slide the chassis all the way into the equipment rack or system cabinet.

4. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the chassis.

Installing the fan modules

To reinstall the fans into the chassis, you must perform a specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.

2. Insert the fan module into the chassis by aligning it with the opening and sliding it into the chassis.

3. Using the cam handle, firmly push the fan module into the chassis to ensure that it is seated all the way into the chassis.

4. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.

5. Repeat these steps for the remaining fan modules.

6. Align the bezel with the ball studs and gently push the bezel onto the ball studs.

Installing the power supplies

Installing the power supplies requires that you place each power supply in the chassis, secure it, and reconnect the power cord to it.

Steps

1. Make sure that the on/off switch of the power supply is in the "Off" position.

2. If you are not already grounded, properly ground yourself.

3. Align the edges of the power supply with the opening in the system chassis, and gently push the power supply into the chassis using the cam handle.

4. Push the power supply cam handle to the closed position, making sure that the cam handle release latch clicks into its locked position.

5. Reconnect the power cord and secure it to the power supply using the power cord locking mechanism.

   Only connect the power cord to the power supply; do not connect the power cord to a power source at this time.

Reinstalling the controller module and I/O expansion module and booting the system

To install and cable controller modules, and, if applicable, the I/O expansion module in the new chassis and then boot the system, you must perform a specific sequence of steps.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module all the way into the chassis using the cam handle.

2. Push the cam handle into the closed position, and then tighten the thumbscrew on the cam handle.

3. Repeat the previous steps if you need to install the I/O expansion module in the chassis

4. Reinstall the cable management device on the controller module and I/O expansion module, if present.

5. Recable the console port only.

6. The next step depends on your system configuration.

If your system is in...	Then...
A stand-alone configuration, single controller module only	Reinstall the blanking panel and then go to the next step.
An HA pair with a second controller module	Repeat the preceding steps for the second controller module.
An HA pair with an I/O expansion module or stand-alone system with an I/O expansion module instead of a second controller module	a. Install the I/O expansion module using the preceding steps for installing the first controller module. b. Recable the adapters in the I/O expansion module.

7. Connect the power supplies to different power sources, and then turn them on.

8. Boot each node to Maintenance mode.

Verifying and setting the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration (HA pair, MetroCluster, or stand-alone).

Steps

1. In Maintenance mode, enter the following command from either controller module to display the HA state of the local controller module and chassis:

   ```
   ha-config show
   ```

The HA state should be the same for all components.

2. Proceed based on the results of Step 1.

If the displayed system state for the chassis...	Then...
Matches your system configuration	Go to the next step.

Does not match your system configuration	a. Set the HA state for the chassis:
	
	`ha-config modify chassis HA-state`
	
	The value for `HA-state` can be one of the following:
	
	• `ha`
	
	• `non-ha`
	
	• `mcc`
	
	b. Confirm that the setting has changed:
	
	`ha-config show`

3. If you have not already done so, recable the rest of your system.

4. The next step depends on your system configuration.

If your system is in...	Then...
A stand-alone configuration	a. Exit Maintenance mode: `halt` b. Go to *"Completing the replacement process* on page 15.
An HA pair with a second controller module	Exit Maintenance mode: `halt` The LOADER prompt appears.
An HA pair with an I/O expansion module instead of a second controller module	Exit Maintenance mode by entering the following command: `halt` The LOADER prompt appears.

Running diagnostics on the chassis

After installing a new chassis, you should run interconnect diagnostics.

Steps

1. Enter the following command at the LOADER prompt:

`boot_diags`

> **Note:** You must enter this command from the LOADER prompt for system-level diagnostics to function properly. The `boot_diags` command starts special drivers designed specifically for system-level diagnostics.

You are prompted to continue the boot. Enter **y**, and then the Maintenance mode prompt (*>) appears.

2. Repeat the previous step on the second controller if you are in an HA configuration.

Note: Both controllers must be in Maintenance mode to run the interconnect test.

3. Clear the status logs by entering the following command:

   `sldiag device clearstatus`

4. Verify that the log is cleared by entering the following command:

   `sldiag device status`

   The following default response is displayed:
   ```
   SLDIAG: No log messages are present.
   ```

5. Enable the interconnect diagnostics tests by entering the following command from the Maintenance mode prompt:

   `sldiag device modify -dev interconnect -sel enable`

   Note: The interconnect tests are disabled by default and must be enabled to run separately.

6. Run the interconnect diagnostics test by entering the following command from the Maintenance mode prompt:

   `sldiag device run -dev interconnect`

   Note: You only need to run the interconnect test from one controller.

7. View the status of the test by entering the following command:

   `sldiag device status`

   Your storage system provides the following output while the tests are still running:
   ```
   There are still test(s) being processed.
   ```
   After all the tests are complete, the following response appears by default:
   ```
   *> <SLDIAG:_ALL_TESTS_COMPLETED>
   ```

8. Verify that no hardware problems resulted from the replacement of the chassis by entering the following command from the Maintenance mode prompt:

   `sldiag device status -dev interconnect -long -state failed`

   System-level diagnostics returns you to the prompt if there are no test failures, or lists the full status of failures resulting from testing the component.

9. Proceed based on the result of the preceding step.

If the system-level diagnostics tests...	Then...

Were completed without any failures	a. Clear the status logs by entering the following command:
	
	`sldiag device clearstatus`
	
	b. Verify that the log is cleared by entering the following command:
	
	`sldiag device status`
	The following default response is displayed:
	
	`SLDIAG: No log messages are present.`
	
	c. Exit Maintenance mode on both controllers by entering following command:
	
	`halt`
	
	   Attention: You must exit Maintenance mode on both controllers before proceeding any further.
	
	d. Enter the following command on both controllers at the LOADER prompt:
	
	`boot_ontap`
	
	e. If your system is a:
	
	  • 7-Mode system, enter
	
	    `cf enable`
	
	  • System running clustered Data ONTAP with two nodes in the cluster, enter the following commands:
	
	    `node::> cluster ha modify -configured true`
	
	    `node::> storage failover modify -node node0 -enabled true`
	
	  • System running clustered Data ONTAP with more than two nodes in the cluster, enter the following command:
	
	    `node::> storage failover modify -node node0 -enabled true`
	
	You have completed system-level diagnostics.

Resulted in some test failures	Determine the cause of the problem. a. Exit Maintenance mode by entering the following command: **`halt`** b. Perform a clean shutdown and disconnect the power supplies. c. Verify that you have observed all the considerations identified for running system-level diagnostics, that cables are securely connected, and that hardware components are properly installed in the storage system. d. Reconnect the power supplies and power on the storage system. e. Rerun the system-level diagnostics tests.

Related information

System-Level Diagnostics Guide

Completing the replacement process

After you replace the part, you can return the failed part to NetApp, as described in the RMA instructions shipped with the kit. Contact technical support at *NetApp Support*, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.
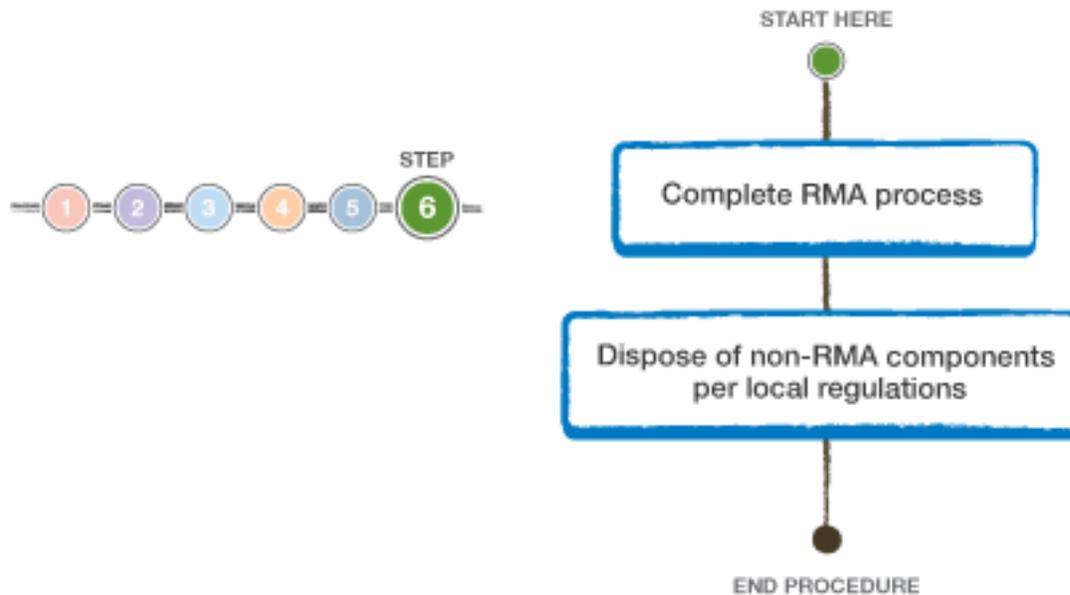
Copyright information

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

Trademark information

How to send comments about documentation and receive update notifications

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