



I/O module

Install and maintain

NetApp

February 13, 2026

This PDF was generated from <https://docs.netapp.com/us-en/ontap-systems/a70-90/io-module-overview.html> on February 13, 2026. Always check docs.netapp.com for the latest.

Table of Contents

I/O module	1
Overview of add and replace an I/O module - AFF A70 and AFF A90	1
Add an I/O module - AFF A70 and AFF A90	1
Step 1: Shut down the impaired controller module	2
Step 2: Add the new I/O module	5
Hot swap an I/O module - AFF A70 and AFF A90	8
Step 1: Ensure the storage system meets the procedure requirements	8
Step 2: Prepare the storage system and I/O module slot	9
Step 3: Replace the failed I/O module	10
Step 4: Bring the replacement I/O module online	11
Step 5: Restore the storage system to normal operation	14
Step 6: Return the failed part to NetApp	15

I/O module

Overview of add and replace an I/O module - AFF A70 and AFF A90

The AFF A70 and AFF A90 systems offer flexibility in expanding or replacing I/O modules to enhance network connectivity and performance. Adding or replacing an I/O module is essential when upgrading network capabilities or addressing a failed module.

You can replace a failed I/O module in your AFF A70 or AFF A90 storage system with the same type of I/O module, or with a different kind of I/O module. You can also add an I/O module into a system with empty slots.

- [Add an I/O module](#)

Adding additional modules can improve redundancy, helping to ensure that the system remains operational even if one module fails.

- [Hot-swap an I/O module](#)

You can hot-swap certain I/O modules for an equivalent I/O module to restore the storage system to its optimal operating state. Hot-swapping is done without having to perform a manual takeover.

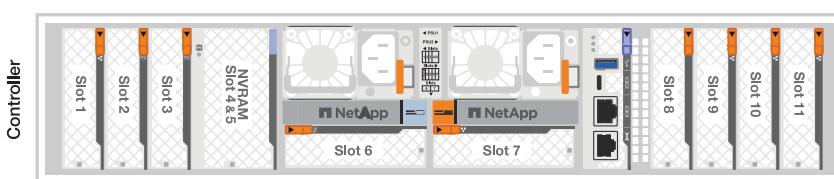
To use this procedure, your storage system must be running ONTAP 9.18.1 or later.

- [Replace an I/O module](#)

Replacing a failing I/O module can restore the system to its optimal operating state.

I/O slot numbering

The I/O slots on AFF A70 and AFF A90 controllers are numbered 1 through 11, as shown in the following illustration.



Add an I/O module - AFF A70 and AFF A90

Add an I/O module to your AFF A70 or AFF A90 system to enhance network connectivity and expand your system's ability to handle data traffic.

You can add an I/O module to your system when there are empty slots available or when all slots are fully populated.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

A storage system has two location LEDs, one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller module

Shut down or take over the impaired controller module using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, **quorum status** of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Option 2: MetroCluster configuration



Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport command:

```
system node autosupport invoke -node * -type all -message
MAINT=number_of_hours_downh
```

The following AutoSupport command suppresses automatic case creation for two hours:

```
cluster1:*> system node autosupport invoke -node * -type all -message
MAINT=2h
```

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next Step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.

If the impaired controller is displaying...	Then...
System prompt or password prompt (enter system password)	<p>Halt or take over the impaired controller from the healthy controller: <code>storage failover takeover -ofnode <i>impaired_node_name</i></code></p> <p>When the impaired controller shows Waiting for giveback..., press Ctrl-C, and then respond <code>y</code>.</p>

Step 2: Add the new I/O module

If the storage system has available slots, install the new I/O module into one of the available slots. If all slots are occupied, remove an existing I/O module to make space and then install the new one.

Before you begin

- Check the [NetApp Hardware Universe](#) to make sure that the new I/O module is compatible with your storage system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in [NetApp Hardware Universe](#) and use the best one available for your I/O module.
- Make sure that all other components are functioning properly.
- Make sure you have the replacement component you received from NetApp.

Add I/O module to an available slot

You can add a new I/O module into a storage system with available slots.

Steps

1. If you are not already grounded, properly ground yourself.
2. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
3. Remove the target slot blanking module from the carrier:
 - a. Depress the cam latch on the blanking module in the target slot.
 - b. Rotate the cam latch away from the module as far as it will go.
 - c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.
4. Install the I/O module:
 - a. Align the I/O module with the edges of the enclosure slot opening.
 - b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
5. Cable the I/O module to the designated device.



Make sure that any unused I/O slots have blanks installed to prevent possible thermal issues.

6. Rotate the cable management tray up to the closed position.
7. From the LOADER prompt, reboot the node:

bye



This reinitializes the I/O module and other components and reboots the node.

8. Give back the controller from the partner controller:

```
storage failover giveback -ofnode target_node_name
```

9. Repeat these steps for controller B.
10. From the healthy node, restore automatic giveback if you disabled it:

```
storage failover modify -node local -auto-giveback true
```

11. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Add I/O module to a fully-populated system

You can add an I/O module to a fully-populated system by removing an existing I/O module and installing a new one in its place.

About this task

Make sure you understand the following scenarios for adding a new I/O module to a fully-populated system:

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See Migrating a LIF for more information.
NIC to storage I/O module	Use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF .

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.
3. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
4. Remove the target I/O module from the chassis:
 - a. Depress the cam latch button.
 - b. Rotate the cam latch away from the module as far as it will go.
 - c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.

Make sure that you keep track of which slot the I/O module was in.

5. Install the I/O module into the target slot in the enclosure:
 - a. Align the module with the edges of the enclosure slot opening.
 - b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
6. Cable the I/O module to the designated device.
7. Repeat the remove and install steps to replace additional modules for the controller.
8. Rotate the cable management tray up to the closed position.
9. Reboot the controller from the LOADER prompt: `_bye_`

This reinitializes the PCIe cards and other components and reboots the node.



If you encounter an issue during reboot, see [BURT 1494308 - Environment shutdown might be triggered during I/O module replacement](#)

10. Give back the controller from the partner controller:

```
storage failover giveback -ofnode target_node_name
```

11. Enable automatic giveback if it was disabled:

```
storage failover modify -node local -auto-giveback true
```

12. Do one of the following:

- If you removed a NIC I/O module and installed a new NIC I/O module, use the following network command for each port:

```
storage port modify -node *<node name> -port *<port name> -mode network
```

- If you removed a NIC I/O module and installed a storage I/O module, install and cable your NS224 shelves, as described in [Hot-add workflow](#).

13. Repeat these steps for controller B.

Hot swap an I/O module - AFF A70 and AFF A90

You can hot swap an Ethernet I/O module in your AFF A70 or AFF A90 storage system if a module fails and your storage system meets all ONTAP version requirements.

To hot swap an I/O module, make sure your storage system is running ONTAP 9.18.1 GA or later, prepare your storage system and I/O module, hot-swap the failed module, bring the replacement module online, restore the storage system to normal operation, and return the failed module to NetApp.

About this task

- You do not need to perform a manual takeover before replacing the failed I/O module.
- Apply commands to the correct controller and I/O slot during the hot-swap:
 - The *impaired controller* is the controller where you are replacing the I/O module.
 - The *healthy controller* is the HA partner of the impaired controller.
- You can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

The storage system includes three location LEDs: one on the operator display panel and one on each controller. The LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Ensure the storage system meets the procedure requirements

To use this procedure, your storage system must be running ONTAP 9.18.1 GA or later, and your storage system must meet all requirements.



If your storage system is not running ONTAP 9.18.1 GA or later, you cannot use this procedure, you must use the [replace an I/O module procedure](#).

- You are hot swapping an Ethernet I/O module in any slot having any combination of ports used for cluster, HA, and client with an equivalent I/O module. You cannot change the I/O module type.

Ethernet I/O modules with ports used for storage or MetroCluster are not hot-swappable.

- Your storage system (switchless or switched cluster configuration) can have any number of nodes supported for your storage system.
- All nodes in the cluster must be running the same ONTAP version (ONTAP 9.18.1GA or later) or running different patch levels of the same ONTAP version.

If nodes in your cluster are running different ONTAP versions, this is considered a mixed-version cluster and hot-swapping an I/O module is not supported.

- The controllers in your storage system can be in either of the following states:
 - Both controllers can be up and running I/O (serving data).
 - Either controller can be in a takeover state if the takeover was caused by the failed I/O module and the nodes are otherwise functioning properly.

In certain situations, ONTAP can automatically perform a takeover of either controller due to the failed I/O module. For example, if the failed I/O module contained all of the cluster ports (all of the cluster links on that controller go down) ONTAP automatically performs a takeover.

- All other components in the storage system must be functioning properly; if not, contact [NetApp Support](#) before continuing with this procedure.

Step 2: Prepare the storage system and I/O module slot

Prepare the storage system and I/O module slot so that it is safe to remove the failed I/O module:

Steps

1. Properly ground yourself.
2. Label the cables to identify where they came from, and then unplug all cables from the target I/O module.

The I/O module should be failed (ports should be in the link down state); however, if the links are still up and they contain the last functioning cluster port, unplugging the cables triggers an automatic takeover.



Wait five minutes after unplugging the cables to ensure any takeovers or LIF failovers complete before continuing with this procedure.

3. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<number of
hours down>h
```

For example, the following AutoSupport message suppresses automatic case creation for two hours:

```
node2::> system node autosupport invoke -node * -type all -message MAINT=2h
```

4. Disable automatic giveback if the partner node has been taken over:

If...	Then...
If either controller took over its partner automatically	<p>Disable automatic giveback:</p> <ol style="list-style-type: none"> 1. Enter the following command from the console of the controller that took over its partner: <pre>storage failover modify -node local -auto -giveback false</pre> <ol style="list-style-type: none"> 2. Enter <i>y</i> when you see the prompt <i>Do you want to disable auto-giveback?</i>
Both controllers are up and running I/O (serving data)	Go to the next step.

5. Prepare the failed I/O module for removal by removing it from service and powering it off:

a. Enter the following command:

```
system controller slot module remove -node impaired_node_name -slot
slot_number
```

b. Enter *y* when you see the prompt *Do you want to continue?*

For example, the following command prepares the failed module in slot 7 on node 2 (the impaired controller) for removal, and displays a message that it is safe to remove:

```
node2::> system controller slot module remove -node node2 -slot 7

Warning: IO_2X_100GBE_NVDA_NIC module in slot 7 of node node2 will be
powered off for removal.

Do you want to continue? {y|n}: y

The module has been successfully removed from service and powered
off. It can now be safely removed.
```

6. Verify the failed I/O module is powered off:

```
system controller slot module show
```

The output should show *powered-off* in the *status* column for the failed module and its slot number.

Step 3: Replace the failed I/O module

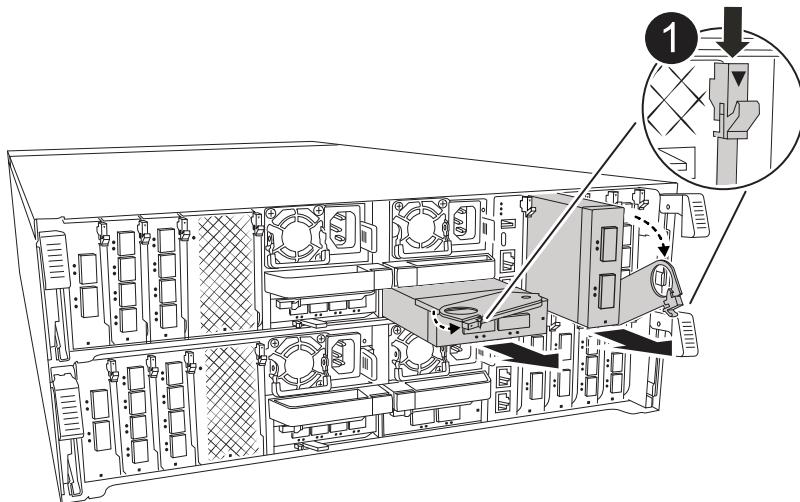
Replace the failed I/O module with an equivalent I/O module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
3. Remove the I/O module from the controller module:



The following illustration shows removing a horizontal and vertical I/O module. Typically, you will only remove one I/O module.



1	Cam locking button
----------	--------------------

- a. Depress the cam latch button.
- b. Rotate the cam latch away from the module as far as it will go.
- c. Remove the module from the controller module by hooking your finger into the cam lever opening and pulling the module out of the controller module.

Keep track of which slot the I/O module was in.

4. Set the I/O module aside.
5. Install the replacement I/O module into the target slot:
 - a. Align the I/O module with the edges of the slot.
 - b. Gently slide the module into the slot all the way into the controller module, and then rotate the cam latch all the way up to lock the module in place.
6. Cable the I/O module.
7. Rotate the cable management tray into the locked position.

Step 4: Bring the replacement I/O module online

Bring the replacement I/O module online, verify the I/O module ports initialized successfully, verify the slot is powered on, and then verify the I/O module is online and recognized.

About this task

After the I/O module is replaced and the ports are returned to a healthy state, LIFs are reverted to the replaced

I/O module.

Steps

1. Bring the replacement I/O module online:

- a. Enter the following command:

```
system controller slot module insert -node impaired_node_name -slot slot_number
```

- b. Enter *y* when you see the prompt, *Do you want to continue?*

The output should confirm the I/O module was successfully brought online (powered on, initialized, and placed into service).

For example, the following command brings slot 7 on node 2 (the impaired controller) online, and displays a message that the process was successful:

```
node2::> system controller slot module insert -node node2 -slot 7

Warning: IO_2X_100GBE_NVDA_NIC module in slot 7 of node node2 will be
powered on and initialized.

Do you want to continue? {y|n}: `y'

The module has been successfully powered on, initialized and placed
into service.
```

2. Verify that each port on the I/O module successfully initialized:

- a. Enter the following command from the console of the impaired controller:

```
event log show -event *hotplug.init*
```



It might take several minutes for any required firmware updates and port initialization.

The output should show one or more *hotplug.init.success* EMS events and *hotplug.init.success*: in the *Event* column, indicating each port on the I/O module initialized successfully.

For example, the following output shows initialization succeeded for I/O ports e7b and e7a:

```
node2::> event log show -event *hotplug.init*
Time           Node           Severity      Event
-----
-----
7/11/2025 16:04:06  node2      NOTICE        hotplug.init.success:
Initialization of ports "e7b" in slot 7 succeeded

7/11/2025 16:04:06  node2      NOTICE        hotplug.init.success:
Initialization of ports "e7a" in slot 7 succeeded

2 entries were displayed.
```

b. If the port initialization fails, review the EMS log for the next steps to take.

3. Verify the I/O module slot is powered on and ready for operation:

```
system controller slot module show
```

The output should show the slot status as *powered-on* and therefore ready for operation of the I/O module.

4. Verify that the I/O module is online and recognized.

Enter the command from the console of the impaired controller:

```
system controller config show -node local -slot slot_number
```

If the I/O module was successfully brought online and is recognized, the output shows I/O module information, including port information for the slot.

For example, you should see output similar to the following for an I/O module in slot 7:

```

node2::> system controller config show -node local -slot 7

Node: node2
Sub- Device/
Slot slot Information
-----
7 - Dual 40G/100G Ethernet Controller CX6-DX
    e7a MAC Address: d0:39:ea:59:69:74 (auto-100g_cr4-fd-
up)
        QSFP Vendor: CISCO-BIZLINK
        QSFP Part Number: L45593-D218-D10
        QSFP Serial Number: LCC2807GJFM-B
    e7b MAC Address: d0:39:ea:59:69:75 (auto-100g_cr4-fd-
up)
        QSFP Vendor: CISCO-BIZLINK
        QSFP Part Number: L45593-D218-D10
        QSFP Serial Number: LCC2809G26F-A
        Device Type: CX6-DX PSID(NAP0000000027)
        Firmware Version: 22.44.1700
        Part Number: 111-05341
        Hardware Revision: 20
        Serial Number: 032403001370

```

Step 5: Restore the storage system to normal operation

Restore your storage system to normal operation by giving back storage to the controller that was taken over (as needed), restoring automatic giveback (as needed), verifying LIFs are on their home ports, and reenabling AutoSupport automatic case creation.

Steps

1. As needed for the version of ONTAP your storage system is running and the state of the controllers, give back storage and restore automatic giveback on the controller that was taken over:

If...	Then...
If either controller took over its partner automatically	<ol style="list-style-type: none"> 1. Return the controller that was taken over to normal operation by giving back its storage: <code>storage failover giveback -ofnode controller that was taken over_name</code> 2. Restore automatic giveback from the console of the controller that was taken over: <code>storage failover modify -node local -auto -giveback true</code>

If...	Then...
Both controllers are up and running I/O (serving data)	Go to the next step.

2. Verify that the logical interfaces are reporting to their home node and ports: `network interface show -is-home false`

If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

3. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

Copyright information

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

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

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

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

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

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

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

Trademark information

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