



Considerations for FC configurations

ONTAP 9

NetApp
December 04, 2021

Table of Contents

- Considerations for FC configurations 1
 - Considerations for FC configurations overview 1
 - Ways to configure FC and FC-NVMe SAN hosts with single nodes 1
 - Ways to configure FC & FC-NVMe SAN hosts with HA pairs 3
 - FC switch configuration best practices 5
 - Supported number of FC hop counts 5
 - Manage systems with FC adapters 5
 - FC target port supported speeds 14
 - FC Target port configuration recommendations 15

Considerations for FC configurations

Considerations for FC configurations overview

You should be aware of several things when setting up your FC configuration.

- You can set up your FC configuration with single nodes or HA pairs using a single fabric or multifabric.
- You should configure two FC data LIFs per node.

This creates redundancy and protects against loss of data access.

- You should configure one management LIF for every storage virtual machine (SVM) supporting SAN.
- Multiple hosts, using different operating systems, such as Windows, Linux, or UNIX, can access the storage solution at the same time.

Hosts require that a supported multipathing solution be installed and configured. Supported operating systems and multipathing solutions can be verified on the Interoperability Matrix.

- ONTAP supports single, dual, or multiple node solutions that are connected to multiple physically independent storage fabrics; a minimum of two are recommended for SAN solutions.

This provides redundancy at the fabric and storage system layers. Redundancy is particularly important because these layers typically support many hosts.

- The use of heterogeneous FC switch fabrics is not supported, except in the case of embedded blade switches.

Specific exceptions are listed on the Interoperability Matrix.

- Cascade, partial mesh, full mesh, core-edge, and director fabrics are all industry-standard methods of connecting FC switches to a fabric, and all are supported.

A fabric can consist of one or multiple switches, and the storage controllers can be connected to multiple switches.

Related information

[NetApp Interoperability Matrix Tool](#)

Ways to configure FC and FC-NVMe SAN hosts with single nodes


You can configure FC and FC-NVMe SAN hosts with single nodes through one or more fabrics. N-Port ID Virtualization (NPIV) is required and must be enabled on all FC switches in the fabric. You cannot directly attach FC or FC-NMVE SAN hosts to single nodes without using an FC switch.

You can configure FC or FC-NVMe SAN hosts with single nodes through a single fabric or multifabrics. The FC target ports (0a, 0c, 0b, 0d) in the illustrations are examples. The actual port numbers vary depending on the model of your storage node and whether you are using expansion adapters.

Single-fabric single-node configurations


In single-fabric single-node configurations, there is one switch connecting a single node to one or more hosts. Because there is a single switch, this configuration is not fully redundant. All hardware platforms that support FC and FC-NVMe support single-fabric single-node configurations. However, the FAS2240 platform requires the X1150A-R6 expansion adapter to support a single-fabric single-node configuration.

The following figure shows a FAS2240 single-fabric single-node configuration. It shows the storage controllers side by side, which is how they are mounted in the FAS2240-2. For the FAS2240-4, the controllers are mounted one above the other. There is no difference in the SAN configuration for the two models.



Multifabric single-node configurations

In multifabric single-node configurations, there are two or more switches connecting a single node to one or more hosts. For simplicity, the following figure shows a multifabric single-node configuration with only two fabrics, but you can have two or more fabrics in any multifabric configuration. In this figure, the storage controller is mounted in the top chassis and the bottom chassis can be empty or can have an IOMX module, as it does in this example.



Related information

[NetApp Technical Report 4684: Implementing and Configuring Modern SANs with NVMe/FC](#)

Ways to configure FC & FC-NVMe SAN hosts with HA pairs


You can configure FC and FC-NVMe SAN hosts to connect to HA pairs through one or more fabrics. You cannot directly attach FC or FC-NVMe SAN hosts to HA pairs without using a switch.

You can configure FC and FC-NVMe SAN hosts with single fabric HA pairs or with multifabric HA pairs. The FC target port numbers (0a, 0c, 0d, 1a, 1b) in the illustrations are examples. The actual port numbers vary depending on the model of your storage node and whether you are using expansion adapters.

Single-fabric HA pairs


In single-fabric HA pair configurations, there is one fabric connecting both controllers in the HA pair to one or more hosts. Because the hosts and controllers are connected through a single switch, single-fabric HA pairs are not fully redundant.

All platforms that support FC configurations support single-fabric HA pair configurations, except the FAS2240 platform. The FAS2240 platform only supports single-fabric single-node configurations.



Multifabric HA pairs

In multifabric HA pairs, there are two or more switches connecting HA pairs to one or more hosts. For simplicity, the following multifabric HA pair figure shows only two fabrics, but you can have two or more fabrics in any multifabric configuration:



FC switch configuration best practices

For best performance, you should consider certain best practices when configuring your FC switch.

A fixed link speed setting is the best practice for FC switch configurations, especially for large fabrics because it provides the best performance for fabric rebuilds and can significantly save time. Although autonegotiation provides the greatest flexibility, FC switch configuration does not always perform as expected, and it adds time to the overall fabric-build sequence.

All of the switches that are connected to the fabric must support N_Port ID virtualization (NPIV) and must have NPIV enabled. ONTAP uses NPIV to present FC targets to a fabric. For details about which environments are supported, see the Interoperability Matrix. For FC and iSCSI best practices, see TR-4080: *Best Practices for Scalable SAN - ONTAP 9*.

Related information

[NetApp Interoperability Matrix Tool](#)

[Best Practices for Scalable SAN - ONTAP 9](#)

Supported number of FC hop counts

The maximum supported FC hop count between a host and storage system depends on the switch supplier and storage system support for FC configurations.

The hop count is defined as the number of switches in the path between the initiator (host) and target (storage system). Cisco also refers to this value as the *diameter of the SAN fabric*.

Switch supplier	Supported hop count
Brocade	7 for FC5 for FCoE
Cisco	7 for FCUp to 3 of the switches can be FCoE switches.

Related information

[NetApp Downloads: Brocade Scalability Matrix Documents](#)

[NetApp Downloads: Cisco Scalability Matrix Documents](#)

Manage systems with FC adapters

Managing systems with FC adapters overview

Commands are available to manage onboard FC adapters and FC adapter cards. These commands can be used to configure the adapter mode, display adapter information, and change the speed.

Most storage systems have onboard FC adapters that can be configured as initiators or targets. You can also use FC adapter cards configured as initiators or targets. Initiators connect to back-end disk shelves, and possibly foreign storage arrays (FlexArray). Targets connect only to FC switches. Both the FC target HBA ports and the switch port speed should be set to the same value and should not be set to auto.

Commands for managing FC adapters

You can use FC commands to manage FC target adapters, FC initiator adapters, and onboard FC adapters for your storage controller. The same commands are used to manage FC adapters for the FC protocol and the FC-NVMe protocol.

FC initiator adapter commands work only at the node level. You must use the `run -node node_name` command before you can use the FC initiator adapter commands.

Commands for managing FC target adapters

If you want to...	Use this command...
Display FC adapter information on a node	<code>network fcp adapter show</code>
Modify FC target adapter parameters	<code>network fcp adapter modify</code>
Display FC protocol traffic information	<code>run -node node_name sysstat -f</code>
Display how long the FC protocol has been running	<code>run -node node_name uptime</code>
Display adapter configuration and status	<code>run -node node_name sysconfig -v adapter</code>
Verify which expansion cards are installed and whether there are any configuration errors	<code>run -node node_name sysconfig -ac</code>
View a man page for a command	<code>man command_name</code>

Commands for managing FC initiator adapters

If you want to...	Use this command...
Display information for all initiators and their adapters in a node	<code>run -node node_name storage show adapter</code>
Display adapter configuration and status	<code>run -node node_name sysconfig -v adapter</code>
Verify which expansion cards are installed and whether there are any configuration errors	<code>run -node node_name sysconfig -ac</code>

Commands for managing onboard FC adapters

If you want to...	Use this command...
Display the status of the onboard FC ports	<code>system node hardware unified-connect show</code>

Configure FC adapters for initiator mode

You can configure individual FC ports of onboard adapters and certain FC adapter cards for initiator mode. Initiator mode is used to connect the ports to tape drives, tape libraries, or third-party storage with FlexArray Virtualization or Foreign LUN Import (FLI).

What you'll need

- LIFs on the adapter must be removed from any port sets of which they are members.
- All LIF's from every storage virtual machine (SVM) using the physical port to be modified must be migrated or destroyed before changing the personality of the physical port from target to initiator.

About this task

Each onboard FC port can be individually configured as an initiator or a target. Ports on certain FC adapters can also be individually configured as either a target port or an initiator port, just like the onboard FC ports. A list of adapters that can be configured for target mode is available in the *Hardware Universe*.



NVMe/FC does support initiator mode.

Steps

1. Remove all LIFs from the adapter:

```
network interface delete -vserver SVM_name -lif lif_name,lif_name
```

2. Take your adapter offline:

```
network fcp adapter modify -node node_name -adapter adapter_port -status-admin down
```

If the adapter does not go offline, you can also remove the cable from the appropriate adapter port on the system.

3. Change the adapter from target to initiator:

```
system hardware unified-connect modify -t initiator adapter_port
```

4. Reboot the node hosting the adapter you changed.

5. Verify that the FC ports are configured in the correct state for your configuration:

```
system hardware unified-connect show
```

6. Bring the adapter back online:

```
node run -node node_name storage enable adapter adapter_port
```

Related information

[NetApp Hardware Universe](#)

Configure FC adapters for target mode

You can configure individual FC ports of onboard adapters and certain FC adapter cards for target mode. Target mode is used to connect the ports to FC initiators.

About this task

Each onboard FC port can be individually configured as an initiator or a target. Ports on certain FC adapters can also be individually configured as either a target port or an initiator port, just like the onboard FC ports. A list of adapters that can be configured for target mode is available in the *Hardware Universe*.

The same steps are used when configuring FC adapters for the FC protocol and the FC-NVMe protocol. However, only certain FC adapters support FC-NVMe. See *Hardware Universe* for a list of adapters that support the FC-NVMe protocol.

Steps

1. Take the adapter offline:

```
node run -node node_name storage disable adapter -d adapter_port
```

If the adapter does not go offline, you can also remove the cable from the appropriate adapter port on the system.

2. Change the adapter from initiator to target:

```
system hardware unified-connect modify -t target adapter_port
```

3. Reboot the node hosting the adapter you changed.
4. Verify that the target port has the correct configuration:

```
network fcp adapter show -node node_name
```

5. Bring your adapter online:

```
network fcp adapter modify -node node_name -adapter adapter_port -state up
```

Related information

[NetApp Hardware Universe](#)

Display information about an FC target adapter

You can use the `network fcp adapter show` command to display system configuration and adapter information for any FC adapter in the system.

Step

1. Display information about the FC adapter by using the `network fcp adapter show` command.

The output displays system configuration information and adapter information for each slot that is used.

```
network fcp adapter show -instance -node node1 -adapter 0a
```

Change the FC adapter speed

You should set your adapter target port speed to match the speed of the device to which it connects, instead of using autonegotiation. A port that is set to autonegotiation can take longer time to reconnect after a takeover/giveback or other interruption.

What you'll need

All LIFs that use this adapter as their home port must be offline.

About this task

Because this task encompasses all storage virtual machines (SVMs) and all LIFs in a cluster, you must use the `-home-port` and `-home-lif` parameters to limit the scope of this operation. If you do not use these parameters, the operation applies to all LIFs in the cluster, which might not be desirable.

Steps

1. Take all of the LIFs on this adapter offline:

```
network interface modify -vserver * -lif * { -home-node node1 -home-port 0c }  
-status-admin down
```

2. Take the adapter offline:

```
network fcp adapter modify -node node1 -adapter 0c -state down
```

If the adapter does not go offline, you can also remove the cable from the appropriate adapter port on the system.

3. Determine the maximum speed for the port adapter:

```
fcp adapter show -instance
```

You cannot modify the adapter speed beyond the maximum speed.

4. Change the adapter speed:

```
network fcp adapter modify -node node1 -adapter 0c -speed 16
```

5. Bring the adapter online:

```
network fcp adapter modify -node node1 -adapter 0c -state up
```

6. Bring all of the LIFs on the adapter online:

```
network interface modify -vserver * -lif * { -home-node node1 -home-port 0c }  
-status-admin up
```

Supported FC ports

The number of onboard FC ports and CNA/UTA2 ports configured for FC varies based on the model of the controller. FC ports are also available through supported FC target expansion adapters or additional UTA2 cards configured with FC SFP+ adapters.

Onboard FC, UTA, and UTA2 ports

- Onboard ports can be individually configured as either target or initiator FC ports.
- The number of onboard FC ports differs depending on controller model.

The *Hardware Universe* contains a complete list of onboard FC ports on each controller model.

- FC ports are only available on FAS2240 systems through the X1150A-R6 expansion adapter.

FAS2220 and FAS2520 systems do not support FC.

Target expansion adapter FC ports

- Available target expansion adapters differ depending on controller model.

The *Hardware Universe* contains a complete list of target expansion adapters for each controller model.

- The ports on some FC expansion adapters are configured as initiators or targets at the factory and cannot be changed.

Others can be individually configured as either target or initiator FC ports, just like the onboard FC ports. A complete list is available in *Hardware Universe*.

Related information

[NetApp Hardware Universe](#)

Port configuration options for the X1143A-R6 adapter

By default the X1143A-R6 adapter is configured in FC target mode, but you can configure its ports as either 10 Gb Ethernet and FCoE (CNA) ports or as 16 Gb FC initiator or target ports, this will require different SFP+ adapters.

Port pairs connected to the same ASIC must be configured in the same mode.

In FC mode, the X1143A-R6 adapter behaves just like any existing FC device with speeds up to 16 Gbps. In CNA mode, you can use the X1143A-R6 adapter for concurrent NIC and FCoE traffic sharing the same 10 GbE port. CNA mode only supports FC target mode for the FCoE function.

Prevent loss of connectivity when using the X1133A-R6 adapter

You can prevent loss of connectivity during a port failure by configuring your system with redundant paths to separate X1133A-R6 HBAs.

The X1133A-R6 HBA is a 4-port, 16 Gb FC adapter consisting of two 2-port pairs. The X1133A-R6 adapter can

be configured as target mode or initiator mode. Each 2-port pair is supported by a single ASIC (for example, Port 1 and Port 2 on ASIC 1 and Port 3 and Port 4 on ASIC 2). Both ports on a single ASIC must be configured to operate in the same mode, either target mode or initiator mode. If an error occurs with the ASIC supporting a pair, both ports in the pair go offline.

To prevent this loss of connectivity, you configure your system with redundant paths to separate X1133A-R6 HBAs, or with redundant paths to ports supported by different ASICs on the HBA.

Supported port configurations for X1143A-R6 adapters

Supported port configurations for X1143A-R6 adapters overview

The FC target mode is the default configuration for X1143A-R6 adapter ports. However, ports on this adapter can be configured as either 10-Gb Ethernet and FCoE ports or as 16-Gb FC ports.

When configured for Ethernet and FCoE, X1143A-R6 adapters support concurrent NIC and FCoE target traffic on the same 10-GBE port. When configured for FC, each two-port pair that shares the same ASIC can be individually configured for FC target or FC initiator mode. This means that a single X1143A-R6 adapter can support FC target mode on one two-port pair and FC initiator mode on another two-port pair.

Related information

[NetApp Hardware Universe](#)

Configure the ports

To configure the unified target adapter (X1143A-R6), you must configure the two adjacent ports on the same chip in the same personality mode.

Steps

1. Configure the ports as needed for Fibre Channel (FC) or Converged Network Adapter (CNA) using the `system node hardware unified-connect modify` command.
2. Attach the appropriate cables for FC or 10 Gb Ethernet.
3. Verify that you have the correct SFP+ installed:

```
network fcp adapter show -instance -node -adapter
```

For CNA, you should use a 10Gb Ethernet SFP. For FC, you should either use an 8 Gb SFP or a 16 Gb SFP, based on the FC fabric being connected to.

Change the UTA2 port from CNA mode to FC mode

You should change the UTA2 port from Converged Network Adapter (CNA) mode to Fibre Channel (FC) mode to support the FC initiator and FC target mode. You should change the personality from CNA mode to FC mode when you need to change the physical medium that connects the port to its network.

Steps

1. Take the adapter offline:

```
network fcp adapter modify -node node_name -adapter adapter_name -status-admin
down
```

2. Change the port mode:

```
ucadmin modify -node node_name -adapter adapter_name -mode fcp
```

3. Reboot the node, and then bring the adapter online:

```
network fcp adapter modify -node node_name -adapter adapter_name -status-admin
up
```

4. Notify your admin or VIF manager to delete or remove the port, as applicable:

- If the port is used as a home port of a LIF, is a member of an interface group (ifgrp), or hosts VLANs, then an admin should do the following:
 - i. Move the LIFs, remove the port from the ifgrp, or delete the VLANs, respectively.
 - ii. Manually delete the port by running the `network port delete` command.

If the `network port delete` command fails, the admin should address the errors, and then run the command again.

- If the port is not used as the home port of a LIF, is not a member of an ifgrp, and does not host VLANs, then the VIF manager should remove the port from its records at the time of reboot.

If the VIF manager does not remove the port, then the admin must remove it manually after the reboot by using the `network port delete` command.

```
net-f8040-34::> network port show
```

```
Node: net-f8040-34-01
```

Port	IPspace	Broadcast	Domain	Link	MTU	Speed (Mbps) Admin/Oper	Health Status
...							
e0i	Default	Default		down	1500	auto/10	-
e0f	Default	Default		down	1500	auto/10	-
...							

```
net-f8040-34::> ucadmin show
```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
net-f8040-34-01	0e	cna	target	-	-	offline

```

net-f8040-34-01
          0f          cna          target          -          -
offline
...

net-f8040-34::> network interface create -vs net-f8040-34 -lif m
-role
node-mgmt-home-node net-f8040-34-01 -home-port e0e -address 10.1.1.1
-netmask 255.255.255.0

net-f8040-34::> network interface show -fields home-port, curr-port

vserver lif                               home-port curr-port
-----
Cluster net-f8040-34-01_clus1 e0a         e0a
Cluster net-f8040-34-01_clus2 e0b         e0b
Cluster net-f8040-34-01_clus3 e0c         e0c
Cluster net-f8040-34-01_clus4 e0d         e0d
net-f8040-34
      cluster_mgmt          e0M         e0M
net-f8040-34
      m                      e0e         e0i
net-f8040-34
      net-f8040-34-01_mgmt1 e0M         e0M
7 entries were displayed.

net-f8040-34::> ucadmin modify local 0e fc

Warning: Mode on adapter 0e and also adapter 0f will be changed to
fc.
Do you want to continue? {y|n}: y
Any changes will take effect after rebooting the system. Use the
"system node reboot" command to reboot.

net-f8040-34::> reboot local
(system node reboot)

Warning: Are you sure you want to reboot node "net-f8040-34-01"?
{y|n}: y

```

5. Verify that you have the correct SFP+ installed:

```
network fcp adapter show -instance -node -adapter
```

For CNA, you should use a 10Gb Ethernet SFP. For FC, you should either use an 8 Gb SFP or a 16 Gb SFP, before changing the configuration on the node.

Change the CNA/UTA2 target adapter optical modules

You should change the optical modules on the unified target adapter (CNA/UTA2) to support the personality mode you have selected for the adapter.

Steps

1. Verify the current SFP+ used in the card. Then, replace the current SFP+ with the appropriate SFP+ for the preferred personality (FC or CNA).
2. Remove the current optical modules from the X1143A-R6 adapter.
3. Insert the correct modules for your preferred personality mode (FC or CNA) optics.
4. Verify that you have the correct SFP+ installed:

```
network fcp adapter show -instance -node -adapter
```

Supported SFP+ modules and Cisco-branded Copper (Twinax) cables are listed in the *Hardware Universe*.

Related information

[NetApp Hardware Universe](#)

View adapter settings

To view the settings for your unified target adapter (X1143A-R6), you must run the `system hardware unified-connect show` command to display all modules on your controller.

Steps

1. Boot your controller without the cables attached.
2. Run the `system hardware unified-connect show` command to see the port configuration and modules.
3. View the port information before configuring the CNA and ports.

FC target port supported speeds

FC target ports can be configured to run at different speeds. You should set the target port speed to match the speed of the device to which it connects. All target ports used by a given host should be set to the same speed. FC target ports can be used for FC-NVMe configurations in the exact same way they are used for FC configurations.

You should set the target port speed to match the speed of the device to which it connects instead of using autonegotiation. A port that is set to autonegotiation can take longer to reconnect after a takeover/giveback or other interruption.

You can configure onboard ports and expansion adapters to run at the following speeds. Each controller and expansion adapter port can be configured individually for different speeds as needed.

4 Gb ports	8 Gb ports	16 Gb ports	32 Gb ports
<ul style="list-style-type: none"> • 4 Gb • 2 Gb • 1 Gb 	<ul style="list-style-type: none"> • 8 Gb • 4 Gb • 2 Gb 	<ul style="list-style-type: none"> • 16 Gb • 8 Gb • 4 Gb 	<ul style="list-style-type: none"> • 32 Gb • 16 Gb • 8 Gb



UTA2 ports can use an 8 Gb SFP+ adapter to support 8, 4, and 2 Gb speeds, if required.

FC Target port configuration recommendations

For best performance and highest availability, you should use the recommended FC target port configuration.

The following table shows the preferred port usage order for onboard FC and FC-NVMe target ports. For expansion adapters, the FC ports should be spread so that they do not use the same ASIC for connectivity. The preferred slot order is listed in the *Hardware Universe* for the version of ONTAP software used by your controller.

FC-NVMe is supported on the following models:

- AFF A300



The AFF A300 onboard ports do not support FC-NVMe.

- AFF A700
- AFF A700s
- AFF A800



The FAS22xx and FAS2520 systems do not have onboard FC ports and do not support add-on adapters.

Controller	Port pairs with shared ASIC	Number of target ports: Preferred ports
FAS9000, AFF A700, AFF A700s and AFF A800	None	All data ports are on expansion adapters. See the <i>Hardware Universe</i> for more information.
8080, 8060 and 8040	0e+0f 0g+0h	1: 0e 2: 0e, 0g 3: 0e, 0g, 0h 4: 0e, 0g, 0f, 0h

Controller	Port pairs with shared ASIC	Number of target ports: Preferred ports
FAS8200 and AFF A300	0g+0h	1: 0g 2: 0g, 0h
8020	0c+0d	1: 0c 2: 0c, 0d
62xx	0a+0b 0c+0d	1: 0a 2: 0a, 0c 3: 0a, 0c, 0b 4: 0a, 0c, 0b, 0d
32xx	0c+0d	1: 0c 2: 0c, 0d
FAS2554, FAS2552, FAS2600 series, FAS2720, FAS2750, AFF A200 and AFF A220	0c+0d 0e+0f	1: 0c 2: 0c, 0e 3: 0c, 0e, 0d 4: 0c, 0e, 0d, 0f

Related information

[NetApp Hardware Universe](#)

Copyright Information

Copyright © 2021 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means-graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system- without prior written permission of the copyright owner.

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

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

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

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

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

Trademark Information

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