



Deployment and configuration

FlexPod

NetApp
March 25, 2024

This PDF was generated from https://docs.netapp.com/us-en/flexpod/healthcare/ehr-meditech-deploy_deployment_and_configuration_overview.html on March 25, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Deployment and configuration 1
 - Overview 1
 - Base infrastructure configuration 2
 - Cisco UCS blade server and switch configuration 3
 - ESXi configuration best practices 8
 - NetApp configuration 9
 - Aggregate configuration 10
 - Storage virtual machine configuration 11
 - Volume configuration 12
 - LUN configuration 13
 - Initiator group configuration 14
 - LUN mappings 14

Deployment and configuration

Overview

The NetApp storage guidance for FlexPod deployment that is provided in this document covers:

- Environments that use ONTAP
- Environments that use Cisco UCS blade and rack-mount servers

This document does not cover:

- Detailed deployment of the FlexPod Datacenter environment

For more information, see [FlexPod Datacenter with FC Cisco Validated Design \(CVD\)](#).

- An overview of MEDITECH software environments, reference architectures, and integration best practices guidance.

For more information, see [TR-4300i: NetApp FAS and All-Flash Storage Systems for MEDITECH Environments Best Practices Guide](#) (NetApp login required).

- Quantitative performance requirements and sizing guidance.


For more information, see [TR-4190: NetApp Sizing Guidelines for MEDITECH Environments](#).

- Use of NetApp SnapMirror technologies to meet backup and disaster recovery requirements.
- Generic NetApp storage deployment guidance.

This section provides an example configuration with infrastructure deployment best practices and lists the various infrastructure hardware and software components and the versions that you can use.

Cabling diagram

The following figure illustrates the 32Gb FC/40GbE topology diagram for a MEDITECH deployment.



Always use the [Interoperability Matrix Tool \(IMT\)](#) to validate that all versions of software and firmware are supported. The table in section "[MEDITECH modules and components](#)" lists the infrastructure hardware and software components that were used in the solution testing.

Next: [Base infrastructure Configuration](#).

Base infrastructure configuration

Network connectivity

The following network connections must be in place before you configure the infrastructure:

- Link aggregation that uses port channels and virtual port channels (vPCs) is used throughout, enabling the design for higher bandwidth and high availability:
 - vPC is used between the Cisco FI and Cisco Nexus switches.
 - Each server has virtual network interface cards (vNICs) with redundant connectivity to the Unified Fabric. NIC failover is used between FIs for redundancy.
 - Each server has virtual host bus adapters (vHBAs) with redundant connectivity to the Unified Fabric.
- The Cisco UCS FI is configured in end- host mode as recommended, providing dynamic pinning of vNICs to uplink switches.

Storage connectivity

The following storage connections must be in place before you configure the infrastructure:

- Storage port interface groups (ifgroups, vPC)
- 10Gb link to switch N9K-A
- 10Gb link to switch N9K-B
- In- band management (active-passive bond):

- 1Gb link to management switch N9K-A
- 1Gb link to management switch N9K-B
- 32Gb FC end-to-end connectivity through Cisco MDS switches; single initiator zoning configured
- FC SAN boot to fully achieve stateless computing; servers are booted from LUNs in the boot volume that is hosted on the AFF storage cluster
- All MEDITECH workloads are hosted on FC LUNs, which are spread across the storage controller nodes

Host software

The following software must be installed:

- ESXi installed on the Cisco UCS blades
- VMware vCenter installed and configured (with all the hosts registered in vCenter)
- VSC installed and registered in VMware vCenter
- NetApp cluster configured

Next: [Cisco UCS Blade Server and Switch Configuration](#).

Cisco UCS blade server and switch configuration

The FlexPod for MEDITECH software is designed with fault tolerance at every level. There is no single point of failure in the system. For optimal performance, Cisco recommends the use of hot spare blade servers.

This document provides high-level guidance on the basic configuration of a FlexPod environment for MEDITECH software. In this section, we present high-level steps with some examples to prepare the Cisco UCS compute platform element of the FlexPod configuration. A prerequisite for this guidance is that the FlexPod configuration is racked, powered, and cabled per the instructions in the [FlexPod Datacenter with Fibre Channel Storage using VMware vSphere 6.5 Update 1, NetApp AFF A-series and Cisco UCS Manager 3.2 CVD](#).

Cisco Nexus switch configuration

A fault- tolerant pair of Cisco Nexus 9300 Series Ethernet switches is deployed for the solution. You should cable these switches as described in the [Cabling Diagram](#) section. The Cisco Nexus configuration helps ensure that Ethernet traffic flows are optimized for the MEDITECH application.

1. After you have completed the initial setup and licensing, run the following commands to set global configuration parameters on both switches:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
port-channel load-balance src-dst l4port
ntp server <global-ntp-server-ip> use-vrf management
ntp master 3
ip route 0.0.0.0/0 <ib-mgmt-vlan-gateway>
copy run start
```

2. Create the VLANs for the solution on each switch using the global configuration mode:

```
vlan <ib-mgmt-vlan-id>
name IB-MGMT-VLAN
vlan <native-vlan-id>
name Native-VLAN
vlan <vmotion-vlan-id>
name vMotion-VLAN
vlan <vm-traffic-vlan-id>
name VM-Traffic-VLAN
vlan <infra-nfs-vlan-id>
name Infra-NFS-VLAN
exit
copy run start
```

3. Create the Network Time Protocol (NTP) distribution interface, port channels, port channel parameters, and port descriptions for troubleshooting per [FlexPod Datacenter with Fibre Channel Storage using VMware vSphere 6.5 Update 1, NetApp AFF A-series and Cisco UCS Manager 3.2 CVD](#).

Cisco MDS 9132T configuration

The Cisco MDS 9100 Series FC switches provide redundant 32Gb FC connectivity between the NetApp AFF A200 or AFF A300 controllers and the Cisco UCS compute fabric. You should connect the cables as described in the [Cabling Diagram](#) section.

1. From the consoles on each MDS switch, run the following commands to enable the required features for the solution:

```
configure terminal
feature npiv
feature fport-channel-trunk
```

2. Configure individual ports, port channels, and descriptions as per the FlexPod Cisco MDS switch configuration section in [FlexPod Datacenter with FC Cisco Validated Design](#).
3. To create the necessary virtual SANs (VSANs) for the solution, complete the following steps while in global configuration mode:

- a. For the Fabric-A MDS switch, run the following commands:

```
vsan database
vsan <vsan-a-id>
vsan <vsan-a-id> name Fabric-A
exit
zone smart-zoning enable vsan <vsan-a-id>
vsan database
vsan <vsan-a-id> interface fc1/1
vsan <vsan-a-id> interface fc1/2
vsan <vsan-a-id> interface port-channel110
vsan <vsan-a-id> interface port-channel112
```

The port channel numbers in the last two lines of the command were created when the individual ports, port channels, and descriptions were provisioned by using the reference document.

- b. For the Fabric-B MDS switch, run the following commands:

```
vsan database
vsan <vsan-b-id>
vsan <vsan-b-id> name Fabric-B
exit
zone smart-zoning enable vsan <vsan-b-id>
vsan database
vsan <vsan-b-id> interface fc1/1
vsan <vsan-b-id> interface fc1/2
vsan <vsan-b-id> interface port-channel111
vsan <vsan-b-id> interface port-channel113
```

The port channel numbers in the last two lines of the command were created when the individual ports, port channels, and descriptions were provisioned by using the reference document.

4. For each FC switch, create device alias names that make the identification of each device intuitive for ongoing operations by using the details in the reference document.
5. Finally, create the FC zones by using the device alias names that were created in step 4 for each MDS switch as follows:
 - a. For the Fabric-A MDS switch, run the following commands:

```

configure terminal
zone name VM-Host-Infra-01-A vsan <vsan-a-id>
member device-alias VM-Host-Infra-01-A init
member device-alias Infra-SVM-fcp_lif01a target
member device-alias Infra-SVM-fcp_lif02a target
exit
zone name VM-Host-Infra-02-A vsan <vsan-a-id>
member device-alias VM-Host-Infra-02-A init
member device-alias Infra-SVM-fcp_lif01a target
member device-alias Infra-SVM-fcp_lif02a target
exit
zoneset name Fabric-A vsan <vsan-a-id>
member VM-Host-Infra-01-A
member VM-Host-Infra-02-A
exit
zoneset activate name Fabric-A vsan <vsan-a-id>
exit
show zoneset active vsan <vsan-a-id>

```

b. For the Fabric-B MDS switch, run the following commands:

```

configure terminal
zone name VM-Host-Infra-01-B vsan <vsan-b-id>
member device-alias VM-Host-Infra-01-B init
member device-alias Infra-SVM-fcp_lif01b target
member device-alias Infra-SVM-fcp_lif02b target
exit
zone name VM-Host-Infra-02-B vsan <vsan-b-id>
member device-alias VM-Host-Infra-02-B init
member device-alias Infra-SVM-fcp_lif01b target
member device-alias Infra-SVM-fcp_lif02b target
exit
zoneset name Fabric-B vsan <vsan-b-id>
member VM-Host-Infra-01-B
member VM-Host-Infra-02-B
exit
zoneset activate name Fabric-B vsan <vsan-b-id>
exit
show zoneset active vsan <vsan-b-id>

```

Cisco UCS configuration guidance


Cisco UCS enables you as a MEDITECH customer to leverage your subject-matter experts in network, storage, and compute to create policies and templates that tailor the environment to your specific needs. After

they are created, these policies and templates can be combined into service profiles that deliver consistent, repeatable, reliable, and fast deployments of Cisco blade and rack servers.

Cisco UCS provides three methods for managing a Cisco UCS system, called a domain:

- Cisco UCS Manager HTML5 GUI
- Cisco UCS CLI
- Cisco UCS Central for multidomain environments

The following figure shows a sample screenshot of the SAN node in Cisco UCS Manager.



In larger deployments, independent Cisco UCS domains can be built for more fault tolerance at the major MEDITECH functional component level.

In highly fault- tolerant designs with two or more data centers, Cisco UCS Central plays a key role in setting global policy and global service profiles for consistency between hosts throughout the enterprise.

To set up the Cisco UCS compute platform, complete the following procedures. Perform these procedures after the Cisco UCS B200 M5 Blade Servers are installed in the Cisco UCS 5108 AC blade chassis. Also, you must complete the cabling requirements as described in the [Cabling Diagram](#) section.

1. Upgrade the Cisco UCS Manager firmware to version 3.2(2f) or later.
2. Configure the reporting, Cisco call home features, and NTP settings for the domain.
3. Configure the server and uplink ports on each Fabric Interconnect.
4. Edit the chassis discovery policy.
5. Create the address pools for out- of- band management, universal unique identifiers (UUIDs), MAC address, servers, worldwide node name (WWNN), and worldwide port name (WWPN).
6. Create the Ethernet and FC uplink port channels and VSANs.
7. Create policies for SAN connectivity, network control, server pool qualification, power control, server BIOS,

and default maintenance.

8. Create vNIC and vHBA templates.
9. Create vMedia and FC boot policies.
10. Create service profile templates and service profiles for each MEDITECH platform element.
11. Associate the service profiles with the appropriate blade servers.

For the detailed steps to configure each key element of the Cisco UCS service profiles for FlexPod, see the [FlexPod Datacenter with Fibre Channel Storage using VMware vSphere 6.5 Update 1, NetApp AFF A-series and Cisco UCS Manager 3.2 CVD](#) document.

[Next: ESXi Configuration Best Practices.](#)

ESXi configuration best practices

For the ESXi host-side configuration, configure the VMware hosts as you would run any enterprise database workload:


- VSC for VMware vSphere checks and sets the ESXi host multipathing settings and HBA timeout settings that work best with NetApp storage systems. The values that VSC sets are based on rigorous internal testing by NetApp.
- For optimal storage performance, consider using storage hardware that supports VMware vStorage APIs - Array Integration (VAAI). The NetApp Plug-In for VAAI is a software library that integrates the VMware Virtual Disk Libraries that are installed on the ESXi host. The VMware VAAI package enables the offloading of certain tasks from the physical hosts to the storage array.

You can perform tasks such as thin provisioning and hardware acceleration at the array level to reduce the workload on the ESXi hosts. The copy offload feature and space reservation feature improve the performance of VSC operations. You can download the plug-in installation package and obtain the instructions for installing the plug-in from the NetApp Support site.

VSC sets ESXi host timeouts, multipath settings, and HBA timeout settings and other values for optimal performance and successful failover of the NetApp storage controllers. Follow these steps:

1. From the VMware vSphere Web Client home page, select vCenter > Hosts.
2. Right-click a host and then select Actions > NetApp VSC > Set Recommended Values.
3. In the NetApp Recommended Settings dialog box, select the values that work best with your system.

The standard recommended values are set by default.



4. Click OK.

[Next: NetApp Configuration.](#)

NetApp configuration

NetApp storage that is deployed for MEDITECH software environments uses storage controllers in a high-availability-pair configuration. Storage must be presented from both controllers to MEDITECH database servers over the FC Protocol. The configuration presents storage from both controllers to evenly balance the application load during normal operation.

ONTAP configuration

This section describes a sample deployment and provisioning procedures that use the relevant ONTAP commands. The emphasis is to show how storage is provisioned to implement the storage layout that NetApp recommends, which uses a high-availability controller pair. One of the major advantages with ONTAP is the ability to scale out without disturbing the existing high-availability pairs.

ONTAP licenses

After you have set up the storage controllers, apply licenses to enable the ONTAP features that NetApp recommends. The licenses for MEDITECH workloads are FC, CIFS, and NetApp Snapshot, SnapRestore, FlexClone, and SnapMirror technologies.

To configure licenses, open NetApp ONTAP System Manager, go to Configuration-Licenses, and then add the appropriate licenses.

Alternatively, run the following command to add licenses by using the CLI:

```
license add -license-code <code>
```

AutoSupport configuration

The NetApp AutoSupport tool sends summary support information to NetApp through HTTPS. To configure AutoSupport, run the following ONTAP commands:

```
autosupport modify -node * -state enable
autosupport modify -node * -mail-hosts <mailhost.customer.com>
autosupport modify -node prod1-01 -from prod1-01@customer.com
autosupport modify -node prod1-02 -from prod1-02@customer.com
autosupport modify -node * -to storageadmins@customer.com
autosupport modify -node * -support enable
autosupport modify -node * -transport https
autosupport modify -node * -hostnamesubj true
```

Hardware-assisted takeover configuration

On each node, enable hardware-assisted takeover to minimize the time that it takes to initiate a takeover in the unlikely event of a controller failure. To configure hardware-assisted takeover, complete the following steps:

1. Run the following ONTAP command to xxx.

Set the partner address option to the IP address of the management port for `prod1-01`.

```
MEDITECH::> storage failover modify -node prod1-01 -hwassist-partner-ip
<prod1-02-mgmt-ip>
```

2. Run the following ONTAP command to xxx:

Set the partner address option to the IP address of the management port for `cluster1-02`.

```
MEDITECH::> storage failover modify -node prod1-02 -hwassist-partner-ip
<prod1-01-mgmt-ip>
```

3. Run the following ONTAP command to enable hardware-assisted takeover on both the `prod1-01` and the `prod1-02` HA controller pair.

```
MEDITECH::> storage failover modify -node prod1-01 -hwassist true
MEDITECH::> storage failover modify -node prod1-02 -hwassist true
```

[Next: Aggregate Configuration.](#)

Aggregate configuration

NetApp RAID DP

NetApp recommends NetApp RAID DP technology as the RAID type for all aggregates in a NetApp FAS or AFF system, including regular NetApp Flash Pool aggregates. MEDITECH documentation might specify the use of RAID 10, but MEDITECH has approved the use of RAID DP.

RAID group size and number of RAID groups

The default RAID group size is 16. This size might or might not be optimal for the aggregates for the MEDITECH hosts at your specific site. For the number of disks that NetApp recommends that you use in a RAID group, see [NetApp TR-3838: Storage Subsystem Configuration Guide](#).

The RAID group size is important for storage expansion because NetApp recommends that you add disks to an aggregate with one or more groups of disks equal to the RAID group size. The number of RAID groups depends on the number of data disks and the RAID group size. To determine the number of data disks that you need, use the NetApp System Performance Modeler (SPM) sizing tool. After you determine the number of data disks, adjust the RAID group size to minimize the number of parity disks to within the recommended range for RAID group size per disk type.

For details on how to use the SPM sizing tool for MEDITECH environments, see [NetApp TR-4190: NetApp Sizing Guidelines for MEDITECH Environments](#).

Storage expansion considerations

When you expand aggregates with more disks, add the disks in groups that are equal to the aggregate RAID group size. Following this approach helps provide performance consistency throughout the aggregate.

For example, to add storage to an aggregate that was created with a RAID group size of 20, the number of disks that NetApp recommends adding is one or more 20-disk groups. So, you should add 20, 40, 60, and so on, disks.

After you expand aggregates, you can improve performance by running reallocation tasks on the affected volumes or aggregate to spread existing data stripes over the new disks. This action is helpful particularly if the existing aggregate was nearly full.



You should plan reallocation of schedules during nonproduction hours because it is a high-CPU and disk-intensive task.

For more information about using reallocation after an aggregate expansion, see [NetApp TR-3929: Reallocate Best Practices Guide](#).

Aggregate-level Snapshot copies

Set the aggregate-level NetApp Snapshot copy reserve to zero and disable the default aggregate Snapshot schedule. Delete any preexisting aggregate-level Snapshot copies if possible.

Next: [Storage Virtual Machine Configuration](#).

Storage virtual machine configuration

This section pertains to deployment on ONTAP 8.3 and later versions.



A storage virtual machine (SVM) is also known as a Vserver in the ONTAP API and in the ONTAP CLI.

SVM for MEDITECH host LUNs

You should create one dedicated SVM per ONTAP storage cluster to own and to manage the aggregates that

contain the LUNs for the MEDITECH hosts.

SVM language encoding setting

NetApp recommends that you set the language encoding for all SVMs. If no language encoding setting is specified at the time that the SVM is created, the default language encoding setting is used. The default language encoding setting is C.UTF-8 for ONTAP. After the language encoding has been set, you cannot modify the language of an SVM with Infinite Volume later.

The volumes that are associated with the SVM inherit the SVM language encoding setting unless you explicitly specify another setting when the volumes are created. To enable certain operations to work, you should use the language encoding setting consistently in all volumes for your site. For example, SnapMirror requires the source and destination SVM to have the same language encoding setting.

[Next: Volume Configuration.](#)

Volume configuration

Volume provisioning

MEDITECH volumes that are dedicated for MEDITECH hosts can be either thick or thin provisioned.

Default volume-level Snapshot copies

Snapshot copies are created as part of the backup workflow. Each Snapshot copy can be used to access the data stored in the MEDITECH LUNs at different times. The MEDITECH- approved backup solution creates thin-provisioned FlexClone volumes based on these Snapshot copies to provide point-in-time copies of the MEDITECH LUNs. The MEDITECH environment is integrated with an approved backup software solution. Therefore, NetApp recommends that you disable the default Snapshot copy schedule on each of the NetApp FlexVol volumes that make up the MEDITECH production database LUNs.

Important: FlexClone volumes share parent data volume space, so it is vital for the volume to have enough space for the MEDITECH data LUNs and the FlexClone volumes that the backup servers create. FlexClone volumes do not occupy more space the way that data volumes do. However, if there are huge deletions on the MEDITECH LUNs in a short time, the clone volumes might grow.

Number of volumes per aggregate

For a NetApp FAS system that uses Flash Pool or NetApp Flash Cache caching, NetApp recommends provisioning three or more volumes per aggregate that are dedicated for storing the MEDITECH program, dictionary, and data files.

For AFF systems, NetApp recommends dedicating four or more volumes per aggregate for storing the MEDITECH program, dictionary, and data files.

Volume-level reallocate schedule

The data layout of storage becomes less optimal over time, especially when it is used by write-intensive workloads such as the MEDITECH Expanse, 6.x, and C/S 5.x platforms. Over time, this situation might increase sequential read latency, resulting in a longer time to complete the backup. Bad data layout or fragmentation can also affect the write latency. You can use volume-level reallocation to optimize the layout of data on disk to improve write latencies and sequential read access. The improved storage layout helps to complete the backup within the allocated time window of 8 hours.

Best practice

At a minimum, NetApp recommends that you implement a weekly volume reallocation schedule to run reallocation operations during the allocated maintenance downtime or during off-peak hours on a production site.



NetApp highly recommends that you run the reallocation task on one volume at a time per controller.

For more information about determining an appropriate volume reallocation schedule for your production database storage, see section 3.12 in [NetApp TR-3929: Reallocate Best Practices Guide](#). That section also guides you on how to create a weekly reallocation schedule for a busy site.

Next: [LUN Configuration](#).

LUN configuration

The number of MEDITECH hosts in your environment determines the number of LUNs that are created within the NetApp FAS or AFF system. The Hardware Configuration Proposal specifies the size of each LUN.

LUN provisioning

MEDITECH LUNs that are dedicated for MEDITECH hosts can be either thick or thin provisioned.

LUN operating system type

To properly align the LUNs that are created, you must correctly set the operating system type for the LUNs. Misaligned LUNs incur unnecessary write operation overhead, and it is costly to correct a misaligned LUN.

The MEDITECH host server typically runs in the virtualized Windows Server environment by using the VMware vSphere hypervisor. The host server can also run in the Windows Server environment on a bare-metal server. To determine the correct operating system type value to set, refer to the “LUN Create” section of [Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#).

LUN size

To determine the LUN size for each MEDITECH host, see the Hardware Configuration Proposal (new deployment) or the Hardware Evaluation Task (existing deployment) document from MEDITECH.

LUN presentation

MEDITECH requires that storage for program, dictionary, and data files be presented to MEDITECH hosts as LUNs by using the FC Protocol. In the VMware virtual environment, the LUNs are presented to the VMware ESXi servers that host the MEDITECH hosts. Then each LUN that is presented to the VMware ESXi server is mapped to each MEDITECH host VM by using RDM in the physical compatibility mode.

You should present the LUNs to the MEDITECH hosts by using the proper LUN naming conventions. For example, for easy administration, you must present the LUN `MTFS01E` to the MEDITECH host `mt-host-01`.

Refer to the MEDITECH Hardware Configuration Proposal when you consult with the MEDITECH and backup

system installer to devise a consistent naming convention for the LUNs that the MEDITECH hosts use.

An example of a MEDITECH LUN name is `MTFS05E`, in which:

- `MTFS` denotes the MEDITECH file server (for the MEDITECH host).
- `05` denotes host number 5.
- `E` denotes the Windows E drive.

[Next: Initiator Group Configuration.](#)

Initiator group configuration

When you use FC as the data network protocol, create two initiator groups (igroups) on each storage controller. The first igroup contains the WWPNs of the FC host interface cards on the VMware ESXi servers that host the MEDITECH host VMs (igroup for MEDITECH).

You must set the MEDITECH igroup operating system type according to the environment setup. For example:

- Use the igroup operating system type `Windows` for applications that are installed on bare-metal-server hardware in a Windows Server environment.
- Use the igroup operating system type `VMware` for applications that are virtualized by using the VMware vSphere hypervisor.



The operating system type for an igroup might be different from the operating system type for a LUN. As an example, for virtualized MEDITECH hosts, you should set the igroup operating system type to `VMware`. For the LUNs that are used by the virtualized MEDITECH hosts, you should set the operating system type to `Windows 2008 or later`. Use this setting because the MEDITECH host operating system is the Windows Server 2008 R2 64-bit Enterprise Edition.

To determine the correct value for the operating system type, see the sections “LUN Igroup Create” and “LUN Create” in the [Clustered Data ONTAP 8.2 Commands: Manual Page Reference](#).

[Next: LUN Mappings.](#)

LUN mappings

LUN mappings for the MEDITECH hosts are established when the LUNs are created.

Copyright information

Copyright © 2024 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.