



FlexPod for MEDITECH Directional Sizing Guide

FlexPod

NetApp
October 30, 2025

Table of Contents

FlexPod for MEDITECH Directional Sizing Guide	1
TR-4774: FlexPod for MEDITECH Directional Sizing	1
Purpose	1
Overall solution benefits	1
Scope	2
Audience	2
Related Documents	2
MEDITECH Workload Overview	3
MEDITECH and backup workloads	3
MEDITECH workload	3
Backup Workload Description	5
Cisco UCS Reference Architecture for MEDITECH	6
Technical specifications for small, medium and large architectures	7
Bill of material for small, medium, and large architectures	7
Additional Information	11
Acknowledgments	11

FlexPod for MEDITECH Directional Sizing Guide

TR-4774: FlexPod for MEDITECH Directional Sizing

Brandon Agee, John Duignan, NetApp
Mike Brennan, Jon Ebmeir, Cisco

In partnership with:



This report provides guidance for sizing FlexPod for a MEDITECH EHR application software environment.

Purpose

FlexPod systems can be deployed to host MEDITECH EXPANSE, 6.x, 5.x, and MAGIC services. FlexPod servers that host the MEDITECH application layer provide an integrated platform for a dependable, high-performance infrastructure. The FlexPod integrated platform is deployed rapidly by skilled FlexPod channel partners and is supported by Cisco and NetApp technical assistance centers.

Sizing is based on information in MEDITECH's hardware configuration proposal and the MEDITECH task document. The goal is to determine the optimal size for compute, network, and storage infrastructure components.

The [MEDITECH Workload Overview](#) section describes the types of compute and storage workloads that can be found in MEDITECH environments.

The [Technical Specifications for Small, Medium, and Large Architectures](#) section details a sample Bill of Materials for the different storage architectures described in the section. The configurations given are general guidelines only. Always size the systems using the sizers based on the workload and tune the configurations accordingly.

Overall solution benefits

Running a MEDITECH environment on the FlexPod architectural foundation can help healthcare organizations improve productivity and decrease capital and operating expenses. FlexPod provides a prevalidated, rigorously tested, converged infrastructure from the strategic partnership of Cisco and NetApp. It is engineered and designed specifically for delivering predictable low-latency system performance and high availability. This approach results in faster response time for users of the MEDITECH EHR system.

The FlexPod solution from Cisco and NetApp meets MEDITECH system requirements with a high performing, modular, prevalidated, converged, virtualized, efficient, scalable, and cost-effective platform. FlexPod Datacenter with MEDITECH delivers several benefits specific to the healthcare industry:

- **Modular architecture.** FlexPod addresses the various needs of the MEDITECH modular architecture with customized FlexPod systems for each specific workload. All components are connected through a clustered server and storage management fabric and use a cohesive management toolset.
- **Simplified operations and lowered costs.** You can eliminate the expense and complexity of legacy

platforms by replacing them with a more efficient and scalable shared resource that can support clinicians wherever they are. This solution delivers better resource usage for greater return on investment (ROI).

- **Quicker deployment of infrastructure.** The integrated design of FlexPod Datacenter with MEDITECH enables customers to have the new infrastructure up and running quickly and easily for both on-site and remote data centers.
- **Scale-out architecture.** You can scale SAN and NAS from terabytes to tens of petabytes without reconfiguring running applications.
- **Nondisruptive operations.** You can perform storage maintenance, hardware lifecycle operations, and software upgrades without interrupting the business.
- **Secure multitenancy.** This benefit supports the increased needs of virtualized server and shared storage infrastructure, enabling secure multitenancy of facility-specific information. This benefit is important if you are hosting multiple instances of databases and software.
- **Pooled resource optimization.** This benefit can help reduce physical server and storage controller counts, load balance workload demands, boost utilization, and simultaneously improve performance.
- **Quality of service (QoS).** FlexPod offers quality of service (QoS) on the entire stack. Industry-leading QoS storage policies enable differentiated service levels in a shared environment. These policies enable optimal performance for workloads and help in isolating and controlling runaway applications.
- **Storage efficiency.** You can reduce storage costs with NetApp 7:1 storage efficiency.
- **Agility.** The industry-leading workflow automation, orchestration, and management tools offered by FlexPod systems allow IT to be far more responsive to business requests. These business requests can range from MEDITECH backup and provisioning of more testing and training environments to analytics database replications for population health management initiatives.
- **Productivity.** You can quickly deploy and scale this solution for optimal clinician end-user experiences.
- **Data Fabric.** The NetApp Data Fabric architecture weaves data together across sites, beyond physical boundaries, and across applications. The NetApp Data Fabric is built for data-driven enterprises in a data-centric world. Data is created and used in multiple locations, and is often shared with applications and infrastructures. Data Fabric provides a way to manage data that is consistent and integrated. It also offers IT more control of the data and simplifies ever-increasing IT complexity.

Scope

This document covers environments that use Cisco UCS and NetApp ONTAP based storage. It provides sample reference architectures for hosting MEDITECH.

It does not cover:

- Detailed sizing guidance using NetApp System Performance Modeler (SPM) or other NetApp sizing tools.
- Sizing for nonproduction workloads.

Audience

This document is intended for NetApp and partner systems engineers and NetApp Professional Services personnel. NetApp assumes that the reader has a good understanding of compute and storage sizing concepts as well as technical familiarity with Cisco UCS and NetApp storage systems.

Related Documents

The following technical reports and other documents are relevant to this Technical Report, and make up a complete set of documents required for sizing, designing, and deploying MEDITECH on FlexPod infrastructure.

- [TR-4753: FlexPod Datacenter for MEDITECH Deployment Guide](#)
- [TR-4190: NetApp Sizing Guidelines for MEDITECH Environments](#)
- [TR-4319: NetApp Deployment Guidelines for MEDITECH Environments](#)



Login credentials for the NetApp Field Portal are required to access some of these reports.

MEDITECH Workload Overview

This section describes the types of compute and storage workloads that you might find in MEDITECH environments.

MEDITECH and backup workloads

When you size NetApp storage systems for MEDITECH environments, you must consider both the MEDITECH production workload and the backup workload.

MEDITECH Host

A MEDITECH host is a database server. This host is also referred to as a MEDITECH file server (for the EXPANSE, 6.x or C/S 5.x platform) or a MAGIC machine (for the MAGIC platform). This document uses the term MEDITECH host to refer to a MEDITECH file server and a MAGIC machine.

The following sections describe the I/O characteristics and performance requirements of these two workloads.

MEDITECH workload

In a MEDITECH environment, multiple servers that run MEDITECH software perform various tasks as an integrated system known as the MEDITECH system. For more information about the MEDITECH system, see the MEDITECH documentation:

- For production MEDITECH environments, consult the appropriate MEDITECH documentation to determine the number of MEDITECH hosts and the storage capacity that must be included as part of sizing the NetApp storage system.
- For new MEDITECH environments, consult the hardware configuration proposal document. For existing MEDITECH environments, consult the hardware evaluation task document. The hardware evaluation task is associated with a MEDITECH ticket. Customers can request either of these documents from MEDITECH.

You can scale the MEDITECH system to provide increased capacity and performance by adding hosts. Each host requires storage capacity for its database and application files. The storage available to each MEDITECH host must also support the I/O generated by the host. In a MEDITECH environment, a LUN is available for each host to support that host's database and application storage requirements. The type of MEDITECH category and the type of platform that you deploy determines the workload characteristics of each MEDITECH host and, therefore, of the system as a whole.

MEDITECH Categories

MEDITECH associates the deployment size with a category number ranging from 1 to 6. Category 1 represents the smallest MEDITECH deployments; category 6 represents the largest. Examples of the MEDITECH application specification associated with each category include metrics such as:

- Number of hospital beds
- Inpatients per year
- Outpatients per year
- Emergency room visits per year
- Exams per year
- Inpatient prescriptions per day
- Outpatient prescriptions per day

For more information about MEDITECH categories, see the MEDITECH category reference sheet. You can obtain this sheet from MEDITECH through the customer or through the MEDITECH system installer.

MEDITECH Platforms

MEDITECH has four platforms:

- EXPANSE
- MEDITECH 6.x
- Client/Server 5.x (C/S 5.x)
- MAGIC

For the MEDITECH EXPANSE, 6.x and C/S 5.x platforms, the I/O characteristics of each host are defined as 100% random with a request size of 4,000. For the MEDITECH MAGIC platform, each host's I/O characteristics are defined as 100% random with a request size of either 8,000 or 16,000. According to MEDITECH, the request size for a typical MAGIC production deployment is either 8,000 or 16,000.

The ratio of reads and writes varies depending on the platform that is deployed. MEDITECH estimates the average mix of read and write and then expresses them as percentages. MEDITECH also estimates the average sustained IOPS value required for each MEDITECH host on a particular MEDITECH platform. The table below summarizes the platform-specific I/O characteristics that are provided by MEDITECH.

MEDITECH Category	MEDITECH Platform	Average Random Read %	Average Random Write %	Average Sustained IOPS per MEDITECH Host
1	EXPANSE, 6.x	20	80	750
2-6	EXPANSE	20	80	750
	6.x	20	80	750
	C/S 5.x	40	60	600
	MAGIC	90	10	400

In a MEDITECH system, the average IOPS level of each host must equal the IOPS values defined in the above table. To determine the correct storage sizing based on each platform, the IOPS values specified in the above table are used as part of the sizing methodology described in the [Technical Specifications for Small, Medium and Large Architectures](#) section.

MEDITECH requires the average random write latency to stay below 1ms for each host. However, temporary increases of write latency up to 2ms during backup and reallocation jobs are considered acceptable. MEDITECH also requires the average random read latency to stay below 7ms for category 1 hosts and below

5ms for category 2 hosts. These latency requirements apply to every host regardless of which MEDITECH platform is being used.

The table below summarizes the I/O characteristics that you must consider when you size NetApp storage for MEDITECH workloads.

Parameter	MEDITECH Category	EXPANSE	MEDITECH 6.x	C/S 5.x	MAGIC
Request size	1-6	4K	4K	4K	8K or 16K
Random/sequential		100% random	100% random	100% random	100% random
Average sustained IOPS	1	750	750	N/A	N/A
	2-6	750	750	600	400
Read/write ratio	1-6	20% read, 80% write	20% read, 80% write	40% read, 60% write	90% read, 10% write
Write latency		<1ms	<1ms	<1ms	<1ms
Temporary peak write latency	1-6	<2ms	<2ms	<2ms	<2ms
Read latency	1	<7ms	<7ms	N/A	N/A
	2-6	<5ms	<5ms	<5ms	<5ms



MEDITECH hosts in categories 3 through 6 have the same I/O characteristics as category 2. For MEDITECH categories 2 through 6, the number of hosts that are deployed in each category differs.

The NetApp storage system should be sized to satisfy the performance requirements described in previous sections. In addition to the MEDITECH production workload, the NetApp storage system must be able to maintain these MEDITECH performance targets during backup operations, as described in the following section.

Backup Workload Description

MEDITECH certified backup software backs up the LUN used by each MEDITECH host in a MEDITECH system. For the backups to be in an application-consistent state, the backup software quiesces the MEDITECH system and suspends I/O requests to disk. While the system is in a quiesced state, the backup software issues a command to the NetApp storage system to create a NetApp Snapshot copy of the volumes that contain the LUNs. The backup software later unquiesces the MEDITECH system, which enables production I/O requests to continue to the database. The software creates a NetApp FlexClone volume based on the Snapshot copy. This volume is used by the backup source while production I/O requests continue on the parent volumes that host the LUNs.

The workload that is generated by the backup software comes from the sequential reading of the LUNs that reside in the FlexClone volumes. The workload is defined as a 100% sequential read workload with a request size of 64,000. For the MEDITECH production workload, the performance criterion is to maintain the required IOPS and the associated read/write latency levels. For the backup workload, however, the attention is shifted to the overall data throughput (MBps) that is generated during the backup operation. MEDITECH LUN backups are required to be completed in an eight-hour backup window, but NetApp recommends that the backup of all MEDITECH LUNs be completed in six hours or less. Aiming to complete the backup in less than six hours

mitigates for events such as an unplanned increase in the MEDITECH workload, NetApp ONTAP background operations, or data growth over time. Any of these events might incur extra backup time. Regardless of the amount of application data stored, the backup software performs a full block-level backup of the entire LUN for each MEDITECH host.

Calculate the sequential read throughput that is required to complete the backup within this window as a function of the other factors involved:

- The desired backup duration
- The number of LUNs
- The size of each LUN to be backed up

For example, in a 50-host MEDITECH environment in which each host's LUN size is 200GB, the total LUN capacity to backup is 10TB.

To back up 10TB of data in eight hours, the following throughput is required:

- $= (10 \times 10^6) \text{MB} (8 \times 3,600) \text{s}$
- $= 347.2 \text{MBps}$

However, to account for unplanned events, a conservative backup window of 5.5 hours is selected to provide headroom beyond the six hours that is recommended.

To back up 10TB of data in eight hours, the following throughput is required:

- $= (10 \times 10^6) \text{MB} (5.5 \times 3,600) \text{s}$
- $= 500 \text{MBps}$

At the throughput rate of 500MBps, the backup can complete within a 5.5-hour time frame, comfortably within the 8-hour backup requirement.

The table below summarizes the I/O characteristics of the backup workload to use when you size the storage system.

Parameter	All Platforms
Request size	64K
Random/sequential	100% sequential
Read/write ratio	100% read
Average throughput	Depends on the number of MEDITECH hosts and the size of each LUN: Backup must complete within 8 hours.
Required backup duration	8 hours

Cisco UCS Reference Architecture for MEDITECH

The architecture for MEDITECH on FlexPod is based on guidance from MEDITECH, Cisco, and NetApp and on partner experience in working with MEDITECH customers of all sizes. The architecture is adaptable and applies best practices for MEDITECH, depending on the customer's data center strategy: whether that is small or large, centralized, distributed, or multitenant.

When deploying MEDITECH, Cisco has designed Cisco UCS reference architectures that align directly with MEDITECH's best practices. Cisco UCS delivers a tightly integrated solution for high performance, high availability, reliability, and scalability to support physician practices and hospital systems with several thousand beds.

Technical specifications for small, medium and large architectures

This section discusses a sample Bill of Materials for different size storage architectures.

Bill of material for small, medium, and large architectures.

The FlexPod design is a flexible infrastructure that encompasses many different components and software versions. Use [TR-4036: FlexPod Technical Specifications](#) as a guide to assembling a valid FlexPod configuration. The configurations in the table below are the minimum requirements for FlexPod, and are just a sample. The configuration can be expanded for each product family as required for different environments and use cases.

For this sizing exercise small corresponds to a Category 3 MEDITECH environment, medium to a Category 5, and large to a Category 6.

	Small	Medium	Large
Platform	One NetApp AFF A220 all-flash storage system HA pair	One NetApp AFF A220 HA pair	One NetApp AFF A300 all-flash storage system HA pair
Disk shelves	9TB x 3.8TB	13TB x 3.8TB	19TB x 3.8TB
MEDITECH database size	3TB-12TB	17TB	>30TB
MEDITECH IOPS	<22,000 IOPs	>25,000 IOPs	>32,000 IOPs
Total IOPS	22000	27000	35000
Raw	34.2TB	49.4TB	68.4TB
Usable capacity	18.53TiB	27.96TiB	33.82TiB
Effective capacity (2:1 storage efficiency)	55.6TiB	83.89TiB	101.47TiB

 Some customer environments might have multiple MEDITECH production workloads running simultaneously or might have higher IOPS requirements. In such cases, work with the NetApp account team to size the storage systems according to the required IOPS and capacity. You should be able to determine the right platform to serve the workloads. For example, there are customers successfully running multiple MEDITECH environments on a NetApp AFF A700 all-flash storage system HA pair.

The following table shows the standard software required for MEDITECH configurations.

Software	Product family	Version or release	Details
Storage	ONTAP	ONTAP 9.4 general availability (GA)	

Software	Product family	Version or release	Details
Network	Cisco UCS fabric interconnects	Cisco UCSM 4.x	Current recommended release
	Cisco Nexus Ethernet switches	7.0(3)I7(6)	Current recommended release
	Cisco FC: Cisco MDS 9132T	8.3(2)	Current recommended release
Hypervisor	Hypervisor	VMware vSphere ESXi 6.7	
	Virtual machines (VMs)	Windows 2016	
Management	Hypervisor management system	VMware vCenter Server 6.7 U1 (VCSA)	
	NetApp Virtual Storage Console (VSC)	VSC 7.0P1	
	NetApp SnapCenter	SnapCenter 4.0	
	Cisco UCS Manager	4.x	

The following table shows an small (category 3) example configuration – infrastructure components.

Layer	Product family	Quantity and model	Details
Compute	Cisco UCS 5108 Chassis	1	Supports up to eight half-width or four full-width blades. Add chassis as server requirement grows.
	Cisco Chassis I/O Modules	2 x 2208	8GB x 10GB uplink ports
	Cisco UCS blade servers	4 x B200 M5	Each with 2 x 14 cores, 2.6GHz or higher clock speed, and 384GB BIOS 3.2(3#)
	Cisco UCS Virtual Interface Cards	4 x UCS 1440	VMware ESXi fNIC FC driver: 1.6.0.47 VMware ESXi eNIC Ethernet driver: 1.0.27.0 (See interoperability matrix: https://ucshcltool.cloudapps.cisco.com/public/)
	2 x Cisco UCS Fabric Interconnects (FI)	2 x UCS 6454 FI	4th-generation fabric interconnects supporting 10/25/100GB Ethernet and 32GB FC
Network	Cisco Ethernet switches	2 x Nexus 9336c-FX2	1GB, 10GB, 25GB, 40GB, 100GB

Layer	Product family	Quantity and model	Details
Storage network	IP Network Nexus 9k for BLOB storage		FI and UCS chassis
	FC: Cisco MDS 9132T		Two Cisco 9132T switches
Storage	NetApp AFF A300 all-flash storage system	1 HA Pair	2-node cluster for all MEDITECH workloads (File Server, Image Server, SQL Server, VMware, and so on)
	DS224C disk shelf	1 DS224C disk shelf	
	Solid-state drive (SSD)	9 x 3.8TB	

The following table shows medium (category 5) example configuration – Infrastructure components

Layer	Product family	Quantity and model	Details
Compute	Cisco UCS 5108 chassis	1	Supports up to eight half-width or four full-width blades. Add chassis as server requirement grows.
	Cisco chassis I/O modules	2 x 2208	8GB x 10GB uplink ports
	Cisco UCS blade servers	6 x B200 M5	Each with 2 x 16 cores, 2.5GHz/or higher clock speed, and 384GB or more memory BIOS 3.2(3#)
	Cisco UCS virtual interface card (VIC)	6 x UCS 1440 VICs	VMware ESXi fNIC FC driver: 1.6.0.47 VMware ESXi eNIC Ethernet driver: 1.0.27.0 (See interoperability matrix:)
	2 x Cisco UCS Fabric Interconnects (FI)	2 x UCS 6454 FI	4th-generation fabric interconnects supporting 10GB/25GB/100GB Ethernet and 32GB FC
Network	Cisco Ethernet switches	2 x Nexus 9336c-FX2	1GB, 10GB, 25GB, 40GB, 100GB
Storage network	IP Network Nexus 9k for BLOB storage		
	FC: Cisco MDS 9132T		Two Cisco 9132T switches

Layer	Product family	Quantity and model	Details
Storage	NetApp AFF A220 all-flash storage system	2 HA Pair	2-node cluster for all MEDITECH workloads (File Server, Image Server, SQL Server, VMware, and so on)
	DS224C disk shelf	1 x DS224C disk shelf	
	SSD	13 x 3.8TB	

The following table shows a large (category 6) example configuration – infrastructure components.

Layer	Product family	Quantity and model	Details
Compute	Cisco UCS 5108 chassis	1	
	Cisco chassis I/O modules	2 x 2208	8 x 10GB uplink ports
	Cisco UCS blade servers	8 x B200 M5	Each with 2 x 24 cores, 2.7GHz and 768GB BIOS 3.2(3#)
	Cisco UCS virtual interface card (VIC)	8 x UCS 1440 VICs	VMware ESXi fNIC FC driver: 1.6.0.47 VMware ESXi eNIC Ethernet driver: 1.0.27.0 (review interoperability matrix: https://ucshcltool.cloudapps.cisco.com/public/)
	2 x Cisco UCS fabric interconnects (FI)	2 x UCS 6454 FI	4th-generation fabric interconnects supporting 10GB/25GB/100GB Ethernet and 32GB FC
Network	Cisco Ethernet switches	2 x Nexus 9336c-FX2	2 x Cisco Nexus 9332PQ1, 10GB, 25GB, 40GB, 100GB
Storage network	IP Network N9k for BLOB storage		
	FC: Cisco MDS 9132T		Two Cisco 9132T switches
Storage	AFF A300	1 HA Pair	2-node cluster for all MEDITECH workloads (File Server, Image Server, SQL Server, VMware, and so on)
	DS224C disk shelf	1 x DS224C disk shelves	
	SSD	19 x 3.8TB	



These configurations provide a starting point for sizing guidance. Some customer environments might have multiple MEDITECH production and non-MEDITECH workloads running simultaneously, or they might have higher IOP requirements. You should work with the NetApp account team to size the storage systems based on the required IOPS, workloads, and capacity to determine the right platform to serve the workloads.

Additional Information

To learn more about the information that is described in this document, see the following documents or websites:

- FlexPod Datacenter with FC Cisco Validated Design.

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_esxi65u1_n9fc.html

- NetApp Deployment Guidelines for MEDITECH Environments.

<https://fieldportal.netapp.com/content/248456> (NetApp login required)

- NetApp Sizing Guidelines for MEDITECH Environments.

www.netapp.com/us/media/tr-4190.pdf

- FlexPod Datacenter for Epic EHR Deployment

www.netapp.com/us/media/tr-4693.pdf

- FlexPod Design Zone

<https://www.cisco.com/c/en/us/solutions/design-zone/data-center-design-guides/flexpod-design-guides.html>

- FlexPod DC with FC Storage (MDS Switches) Using NetApp AFF, vSphere 6.5U1, and Cisco UCS Manager

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_esxi65u1_n9fc.html

- Cisco Healthcare

<https://www.cisco.com/c/en/us/solutions/industries/healthcare.html?dtid=osscdc000283>

Acknowledgments

The following people contributed to the writing and creation of this guide.

- Brandon Agee, Technical Marketing Engineer, NetApp
- John Duignan, Solutions Architect - Healthcare, NetApp
- Ketan Mota, Product Manager, NetApp
- Jon Ebmeier, Technical Solutions Architect, Cisco Systems, Inc
- Mike Brennan, Product Manager, Cisco Systems, Inc

Copyright information

Copyright © 2025 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—with 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.