



Technical Report

FlexPod Datacenter for MEDITECH Deployment Guide

Healthcare Industry Solution

Brandon Agee and John Duignan, NetApp
Mike Brennan and Jon Ebmeir, Cisco
February 2019 | TR-4753

Reviewed by

MEDITECH

Abstract

This technical report is for customers who plan to deploy MEDITECH on FlexPod® systems. It provides a brief overview of the FlexPod architecture for MEDITECH and covers the FlexPod for MEDITECH setup and installation procedures for the healthcare industry.

FlexPod systems that are deployed to host MEDITECH Expanse, MEDITECH 6.x, Client/Server 5.x (C/S 5.x), MAGIC, and services servers hosting the MEDITECH application layer provide an integrated platform for a dependable, high-performance infrastructure that can be deployed rapidly. The FlexPod integrated platform is deployed by skilled FlexPod channel partners and is supported by Cisco and NetApp® technical assistance centers.

TABLE OF CONTENTS

1 Overview	4
1.1 Overall Solution Benefits	4
1.2 FlexPod	4
1.3 MEDITECH Overview	10
1.4 Comprehensive Management Tools and Automation Capabilities	12
2 Design	14
2.1 Storage Layout	14
2.2 Storage Placement	15
2.3 Storage Controller Configuration	15
3 Deployment and Configuration	16
3.1 Cabling Diagram	17
3.2 Base Infrastructure Configuration	17
3.3 Cisco UCS Blade Server and Switch Configuration	18
3.4 ESXi Configuration Best Practices	22
3.5 NetApp Configuration	22
3.6 Aggregate Configuration	23
3.7 Storage Virtual Machine Configuration	24
3.8 Volume Configuration	24
3.9 LUN Configuration	25
3.10 Initiator Group Configuration	26
3.11 LUN Mappings	26
Appendix A: MEDITECH Modules and Components	27
Acknowledgements	27
Where to Find Additional Information	28
FlexPod Design Zone	28
NetApp Technical Reports	28
ONTAP Documentation	28
Cisco Nexus, MDS, Cisco UCS, and Cisco UCS Manager Guides	28
Version History	29

LIST OF TABLES

Table 1) MEDITECH modules and components	27
--	----

LIST OF FIGURES

Figure 1) FlexPod Cooperative Support model	6
Figure 2) NetApp deduplication process	9

Figure 3) FlexPod for MEDITECH workloads..... 10
Figure 4) FlexPod for MEDITECH physical topology diagram..... 17
Figure 5) Cisco UCS Manager HTML5 UI..... 21

1 Overview

1.1 Overall Solution Benefits

By running a MEDITECH environment on the FlexPod architectural foundation, healthcare organizations can expect to see an improvement in staff productivity and a decrease in capital and operating expenses. FlexPod Datacenter with MEDITECH delivers several benefits specific to the healthcare industry:

- **Simplified operations and lowered costs.** Eliminate the expense and complexity of legacy platforms by replacing them with a more efficient and scalable shared resource capable of supporting clinicians wherever they are. This solution delivers higher resource utilization for greater return on investment (ROI).
- **Quicker deployment of infrastructure.** Whether it's an existing data center or a remote location, the integrated and tested design of FlexPod Datacenter enables customers to have the new infrastructure up and running in less time with less effort.
- **Certified storage.** NetApp ONTAP® data management software with MEDITECH enables customers to have the assurance of a tested and certified storage vendor. MEDITECH does not certify other infrastructure components.
- **Scale-out architecture.** Scale SAN and NAS from terabytes (TBs) to tens of petabytes (PBs) without reconfiguring running applications.
- **Nondisruptive operations.** Perform storage maintenance, hardware lifecycle operations, and FlexPod upgrades without interrupting the business.
- **Secure multitenancy.** Supports the increased needs of virtualized server and storage shared infrastructure, enabling secure multitenancy of facility-specific information, particularly if hosting multiple instances of databases and software.
- **Pooled resource optimization.** Help reduce physical server and storage controller counts, load balance workload demands, and boost utilization while improving performance.
- **Quality of service (QoS).** FlexPod offers QoS on the entire stack. Industry-leading QoS network, compute, and 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.** Reduce storage costs with the NetApp 7:1 storage efficiency guarantee.¹
- **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 test and training environments to analytics database replications for population health management initiatives.
- **Productivity.** Quickly deploy and scale this solution for optimal clinician end-user experiences.
- **NetApp 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 it often needs to be leveraged and shared with other locations, applications, and infrastructures. Customers want a way to manage data that is consistent and integrated. It provides a way to manage data that puts IT in control and simplifies ever-increasing IT complexity.

1.2 FlexPod

New Infrastructure Approach for MEDITECH EHR

Healthcare provider organizations remain under pressure to maximize the benefits of their substantial investments in industry-leading MEDITECH electronic health records (EHRs). For mission-critical

¹ www.netapp.com/us/media/netapp-aff-efficiency-guarantee.pdf

applications, when customers design their data centers for MEDITECH solutions, they often identify the following goals for their data center architecture:

- High availability of the MEDITECH applications
- High performance
- Ease of implementing MEDITECH in the data center
- Agility and scalability to enable growth with new MEDITECH releases or applications
- Cost effectiveness
- Alignment with MEDITECH guidance and target platforms
- Manageability, stability, and ease of support
- Robust data protection, backup, recovery, and business continuance

As MEDITECH users evolve their organizations to become accountable care organizations and adjust to tightened, bundled reimbursement models, the challenge becomes delivering the required MEDITECH infrastructure in a more efficient and agile IT delivery model.

Value of Prevalidated Converged Infrastructure

MEDITECH is prescriptive as to its customers' hardware requirements because of an overarching requirement for delivering predictable low-latency system performance and high availability.

FlexPod is 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 MEDITECH compliance and ultimately the best response time for users of the MEDITECH 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. It provides:

- **Modular architecture.** FlexPod addresses the varied needs of the MEDITECH modular architecture with purpose-configured FlexPod platforms for each specific workload. All components are connected through a clustered server and storage management fabric and a cohesive management toolset.
- **Industry-leading technology at each level of the converged stack.** Cisco, NetApp, VMware, and Windows are all ranked as number 1 or number 2 by industry analysts in their respective categories of servers, networking, storage, and operating systems.
- **Investment protection with standardized, flexible IT.** The FlexPod reference architecture anticipates new product versions and updates, with rigorous ongoing interoperability testing to accommodate future technologies as they become available.
- **Proven deployment across a broad range of environments.** Pretested and jointly validated with popular hypervisors, operating systems, applications, and infrastructure software, FlexPod has been installed in multiple MEDITECH customer organizations.

Proven FlexPod Architecture and Cooperative Support

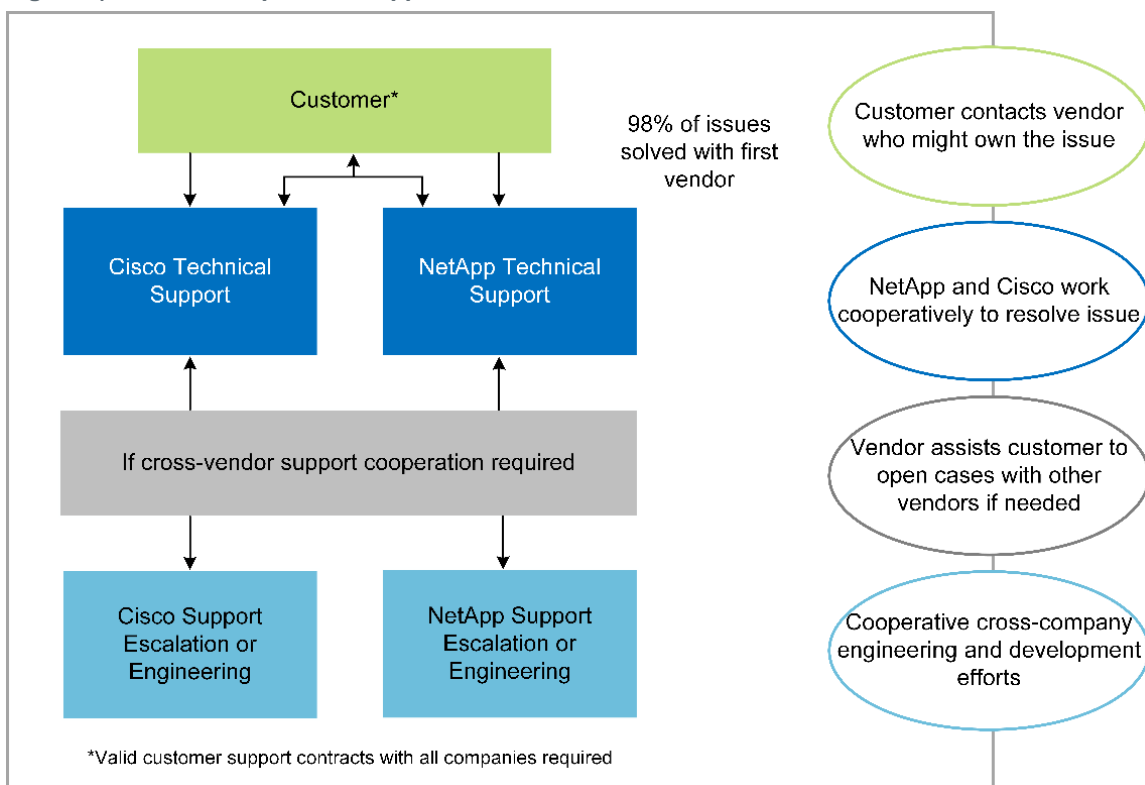
FlexPod is a proven data center solution, offering a flexible, shared infrastructure that easily scales to support growing workload demands without affecting performance. By leveraging the FlexPod architecture, this solution delivers the full benefits of FlexPod, including:

- **Performance to meet the MEDITECH workload requirements.** Depending on the Hardware Configuration Proposal requirements, different ONTAP platforms can be deployed to meet the required I/O and latency requirements.
- **Scalability to easily accommodate clinical data growth.** Dynamically scale virtual machines (VMs), servers, and storage capacity on demand, without traditional limits.
- **Enhanced efficiency.** Reduce both administration time and TCO with a converged virtualized infrastructure, which is easier to manage and stores data more efficiently while driving more performance from MEDITECH software.

- **Reduced risk.** Minimize business disruption with a prevalidated platform built on a defined architecture that eliminates deployment guesswork and accommodates ongoing workload optimization.
- **FlexPod Cooperative Support.** NetApp and Cisco have established Cooperative Support, a strong, scalable, and flexible support model to address the unique support requirements of the FlexPod converged infrastructure. This model uses the combined experience, resources, and technical support expertise of NetApp and Cisco to provide a streamlined process for identifying and resolving a customer's FlexPod support issue, regardless of where the problem resides. The FlexPod Cooperative Support model helps to make sure that your FlexPod system operates efficiently and benefits from the most up-to-date technology, while providing an experienced team to help resolve integration issues.

FlexPod Cooperative Support is especially valuable to healthcare organizations running business-critical applications such as MEDITECH on the FlexPod converged infrastructure. Figure 1 illustrates the FlexPod Cooperative Support model.

Figure 1) FlexPod Cooperative Support model.



In addition to these benefits, each component of the FlexPod Datacenter stack with MEDITECH solution delivers specific benefits for MEDITECH EHR workflows.

Cisco Unified Computing System

A self-integrating, self-aware system, Cisco Unified Computing System (UCS) consists of a single management domain interconnected with a unified I/O infrastructure. Cisco UCS for MEDITECH environments has been aligned with MEDITECH infrastructure recommendations and best practices to help ensure that the infrastructure can deliver critical patient information with maximum availability.

The foundation of MEDITECH on Cisco UCS architecture is Cisco UCS technology, with its integrated systems management, Intel Xeon processors, and server virtualization. These integrated technologies solve data center challenges and enable customers to meet their goals for data center design for MEDITECH. Cisco UCS unifies LAN, SAN, and systems management into one simplified link for rack servers, blade servers, and VMs. Cisco UCS is an end-to-end I/O architecture that incorporates Cisco unified fabric and Cisco fabric extender (FEX) technology to connect every component in Cisco UCS with a single network fabric and a single network layer.

The system can be deployed as a single or multiple logical units that incorporate and scale across multiple blade chassis, rack servers, racks, and data centers. The system implements a radically simplified architecture that eliminates the multiple redundant devices that populate traditional blade server chassis and rack servers, that result in layers of complexity: Ethernet and Fibre Channel (FC) adapters and chassis management modules. Cisco UCS consists of a redundant pair of Cisco UCS Fabric Interconnects (FIs) that provide a single point of management, and a single point of control, for all I/O traffic.

Cisco UCS uses service profiles to help ensure that virtual servers in the Cisco UCS infrastructure are configured correctly. Service profiles are composed of network, storage, and compute policies that are created once by subject matter experts in each discipline. Service profiles include critical server information about the server identity such as LAN and SAN addressing, I/O configurations, firmware versions, boot order, network virtual LAN (VLAN), physical port, and QoS policies. Service profiles can be dynamically created and associated with any physical server in the system in minutes rather than hours or days. The association of service profiles with physical servers is performed as a simple, single operation and enables migration of identities between servers in the environment without requiring any physical configuration changes. It facilitates rapid bare-metal provisioning of replacements for retired servers.

Using service profiles helps ensure that servers are configured consistently throughout the enterprise. When using multiple Cisco UCS management domains, Cisco UCS Central can use global service profiles to synchronize configuration and policy information across domains. If maintenance needs to be performed in one domain, the virtual infrastructure can be migrated to another domain. This approach helps to ensure that even when a single domain is offline, applications continue to run with high availability.

Cisco UCS has been extensively tested with MEDITECH over a multi-year period to demonstrate that it meets the server configuration requirements. Cisco UCS is a supported server platform, as listed on the MEDITECH Product Resources System Support site.

Cisco Networking

Cisco Nexus switches and Cisco MDS multilayer directors provide enterprise-class connectivity and SAN consolidation. Cisco multiprotocol storage networking reduces business risk by providing flexibility and options: FC, Fibre Connection (FICON), FC over Ethernet (FCoE), SCSI over IP (iSCSI), and FC over IP (FCIP).

Cisco Nexus switches offer one of the most comprehensive data center network feature sets in a single platform. They deliver high performance and density for both data center and campus core. They also offer a full feature set for data center aggregation, end-of-row, and data center interconnect deployments in a highly resilient modular platform.

Cisco UCS integrates computing resources with Cisco Nexus switches and a unified I/O fabric that identifies and handles different types of network traffic, including storage I/O, streamed desktop traffic, management, and access to clinical and business applications:

- **Infrastructure scalability.** Virtualization, efficient power and cooling, cloud scale with automation, high density, and performance all support efficient data center growth.
- **Operational continuity.** The design integrates hardware, NX-OS software features, and management to support zero-downtime environments.
- **Network and computer QoS.** Cisco delivers policy-driven Class of Service (CoS) and QoS across the networking, storage, and compute fabric to ensure optimal performance of mission-critical applications.
- **Transport flexibility.** Incrementally adopt new networking technologies with a cost-effective solution.

Together, Cisco UCS with Cisco Nexus switches and Cisco MDS multilayer directors provide an optimal compute, networking, and SAN connectivity solution for MEDITECH.

NetApp ONTAP

NetApp storage running ONTAP software reduces overall storage costs while delivering the low-latency read and write response times and IOPS required for MEDITECH workloads. ONTAP supports both all-flash and hybrid storage configurations to create an optimal storage platform to meet MEDITECH requirements. NetApp flash-accelerated systems have received MEDITECH's validation and certification, providing MEDITECH customers with the performance and responsiveness key to latency-sensitive MEDITECH operations. NetApp can also isolate production from nonproduction by creating multiple fault domains in a single cluster. NetApp reduces performance issues by guaranteeing a minimum performance level for workloads with ONTAP minimum QoS.

The scale-out architecture of the ONTAP software can flexibly adapt to various I/O workloads. To deliver the necessary throughput and low latency required for clinical applications while providing a modular scale-out architecture, all-flash configurations are typically used in ONTAP architectures. NetApp AFF nodes can be combined in the same scale-out cluster with hybrid (HDD and flash) storage nodes suitable for storing large datasets with high throughput. Along with a MEDITECH approved backup solution, customers can clone, replicate, and back up the MEDITECH environment (from expensive solid-state drive [SSD] storage) to more economical HDD storage on other nodes, meeting or exceeding MEDITECH guidelines for SAN-based cloning and backup of production disk pools.

ONTAP offers features that are useful in MEDITECH environments: simplifying management, increasing availability and automation, and reducing the total amount of storage needed:

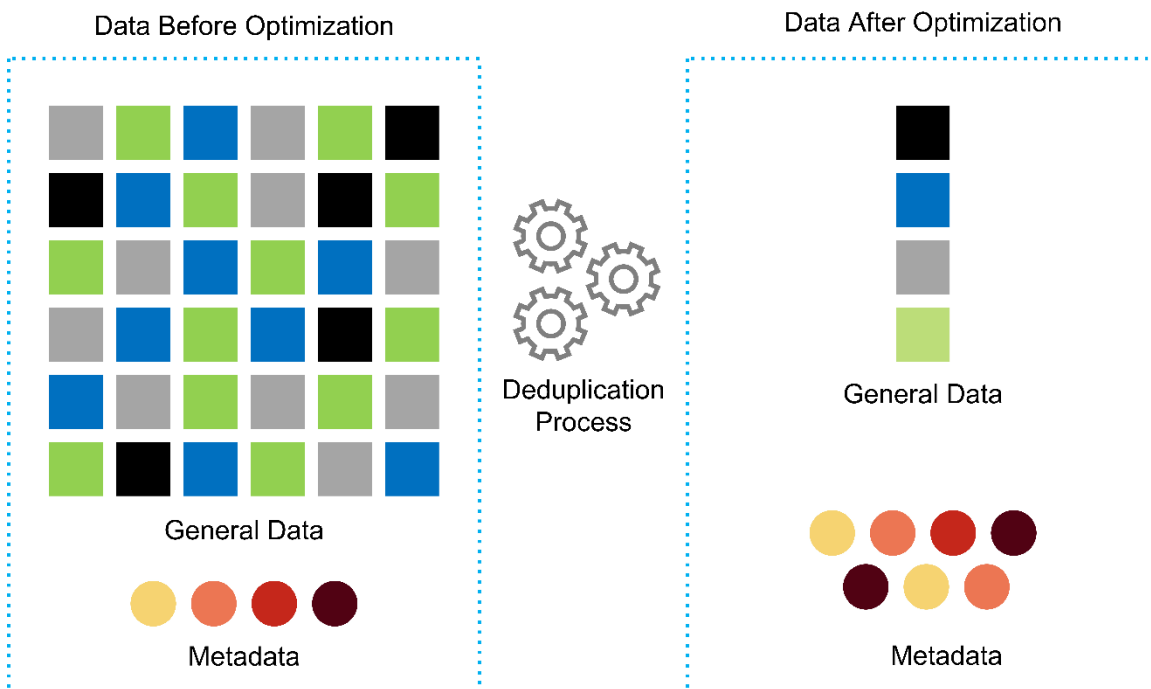
- **Outstanding performance.** The NetApp AFF solution shares the unified storage architecture, ONTAP software, management interface, rich data services, and advanced feature set as the rest of the FAS product families. This innovative combination of all-flash media with ONTAP delivers the consistent low latency and high IOPS of all-flash storage with the industry-leading ONTAP software.
- **Storage efficiency.** Reduce total capacity requirements with deduplication, NetApp FlexClone® data replication technology, inline compression, inline compaction, thin replication, thin provisioning, and aggregate deduplication.

NetApp deduplication provides block-level deduplication in a NetApp FlexVol® volume or data constituent. Essentially, deduplication removes duplicate blocks, storing only unique blocks in the FlexVol volume or data constituent.

Deduplication works with a high degree of granularity and operates on the active file system of the FlexVol volume or data constituent. It is application transparent, and therefore, it can be used to deduplicate data originating from any application that uses the NetApp system. Volume deduplication can be run as an inline process (starting in Data ONTAP 8.3.2) and/or as a background process that can be configured to run automatically, be scheduled, or run manually through the CLI, NetApp System Manager, or NetApp OnCommand® Unified Manager.

Figure 2 illustrates how NetApp deduplication works at the highest level.

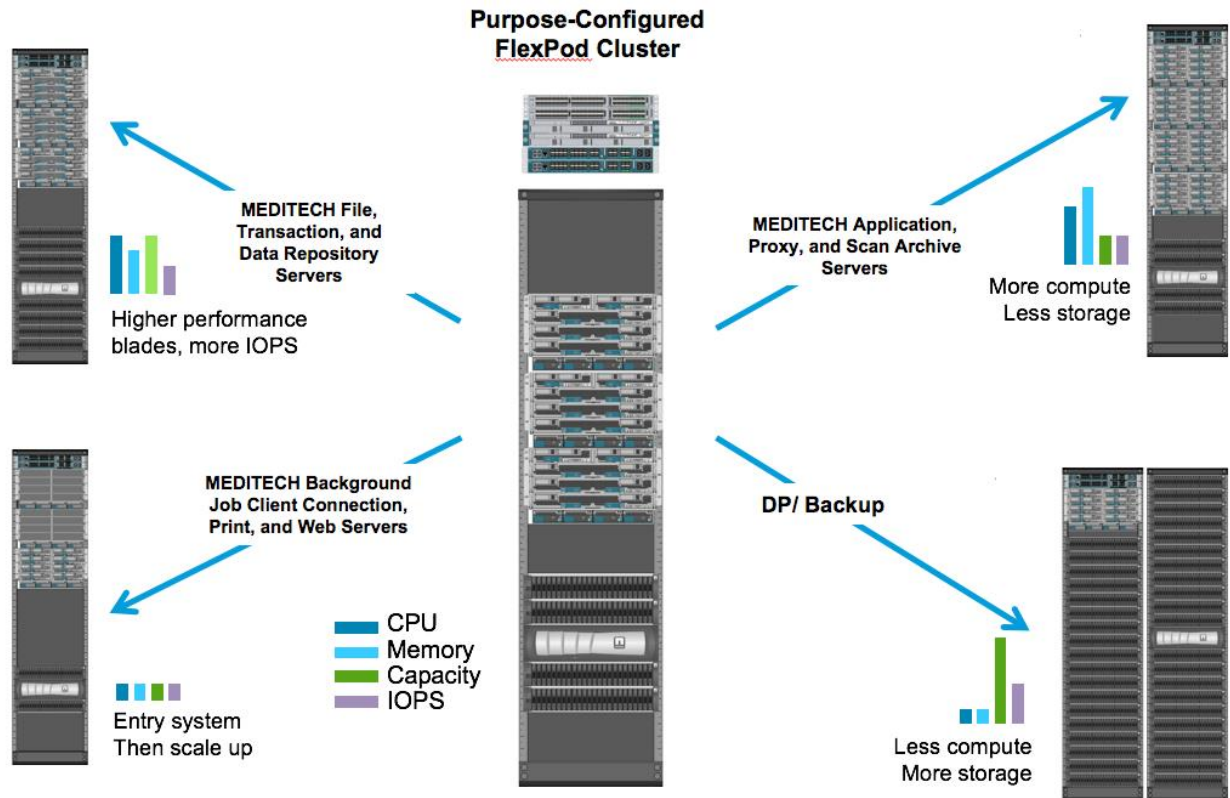
Figure 2) NetApp deduplication process.



- **Space-efficient cloning.** The FlexClone capability allows you to almost instantly create clones to support backup and test environment refresh. These clones consume more storage only as changes are made.
- **NetApp Snapshot™ and SnapMirror® technologies.** ONTAP is able to create space-efficient Snapshot copies of the logical unit numbers (LUNs) used by MEDITECH host. For dual-site deployments, SnapMirror software can be implemented for more data replication and resiliency.
- **Integrated data protection.** Full data protection and disaster recovery features help customers protect critical data assets and provide disaster recovery.
- **Nondisruptive operations.** Upgrading and maintenance can be performed without taking data offline.
- **QoS and Adaptive QoS (AQoS).** Storage QoS allows you to limit potential bully workloads. More importantly, QoS can guarantee minimum performance for critical workloads such as MEDITECH production. NetApp QoS can reduce performance-related issues by limiting contention. AQoS works with predefined policy groups, which can be applied directly to a volume. These policy groups are able to automatically scale a throughput ceiling or floor-to-volume size, maintaining the ratio of IOPS to terabytes and gigabytes as the size of the volume changes.
- **NetApp Data Fabric.** The NetApp Data Fabric simplifies and integrates data management across cloud and on-premises to accelerate digital transformation. It delivers consistent and integrated data management services and applications for data visibility and insights, data access and control, and data protection and security. NetApp is integrated with Amazon Web Services (AWS), Azure, Google Public Cloud, and IBM Cloud clouds, giving customers a wide breadth of choice.

Figure 3 illustrates the FlexPod for MEDITECH workloads.

Figure 3) FlexPod for MEDITECH workloads.



1.3 MEDITECH Overview

Medical Information Technology, Inc., commonly known as MEDITECH, is a Massachusetts-based software company that provides information systems that are installed in healthcare organizations. MEDITECH provides an EHR system designed to store and organize the latest patient data and provides the data to clinical staff. Patient data includes, but is not limited to, demographics, medical history, medication, laboratory test results, radiology images, and personal information such as age, height, and weight.

It is beyond the scope of this document to cover the wide span of functions supported by MEDITECH software. See the appendix for these broad sets of MEDITECH functions. MEDITECH application requires several VMs to support these functions. Refer to the recommendations provided by MEDITECH to deploy these applications.

For each deployment, from the storage system point of view, all MEDITECH software systems require a distributed patient-centric database. MEDITECH has their own proprietary database, which uses the Windows operating system.

Bridgehead and Commvault are the two backup software applications certified by both NetApp and MEDITECH. The scope of this document does not cover the deployment of these backup applications.

The primary focus of this document is to enable the FlexPod stack (servers and storage) to satisfy performance-driven requirements for the MEDITECH database and backup requirements used in the EHR environment.

Purpose-Built for Specific MEDITECH Workloads

MEDITECH does not resell server, network, or storage hardware, hypervisors, or operating systems, however, it has specific requirements for each component of the infrastructure stack. Therefore, Cisco and NetApp worked together to test and enable FlexPod Datacenter to be successfully configured, deployed, and supported to meet customers' MEDITECH production environment requirements.

MEDITECH Category

MEDITECH associates the deployment size with a category number ranging from 1 to 6. Category 1 represents the smallest MEDITECH deployments and category 6 represents the largest MEDITECH deployments.

For information about the I/O characteristics and performance requirements for a MEDITECH host in each category, see NetApp [TR-4190: NetApp Sizing Guidelines for MEDITECH Environments](#).

MEDITECH Platform

The MEDITECH Expanse platform is the latest version of the company's EHR software. Earlier MEDITECH platforms are Client/Server 6.x and MAGIC. This section describes the MEDITECH platform (applicable to Expanse, 6.x, C/S 5.x, and MAGIC) pertaining to the MEDITECH host and its storage requirements.

For all the preceding MEDITECH platforms, multiple servers run MEDITECH software, performing various tasks. Figure 1 depicts a typical MEDITECH system, including MEDITECH hosts serving as application database servers and other MEDITECH servers. Examples of other MEDITECH servers include the Data Repository application, Scanning and Archiving application, and Background Job Clients. For the complete list of other MEDITECH servers, see the "Hardware Configuration Proposal" (for new deployments) and "Hardware Evaluation Task" (for existing deployments) documents. You can obtain these documents from MEDITECH through the MEDITECH system integrator or the MEDITECH Technical Account Manager (TAM).

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.

MEDITECH hosts can be physical servers or VMs running on the Microsoft Windows Server operating system. Most commonly in the field, MEDITECH hosts are deployed as Windows VMs running in a VMware ESX server. As of this writing, VMware is the only hypervisor supported by MEDITECH. A MEDITECH host stores its program, dictionary, and data files on a Microsoft Windows drive (for example, drive E) on the Windows system.

In a virtual environment, a Windows E drive resides on a LUN attached to the VM by way of a raw device mapping (RDAM) in physical compatibility mode. The use of virtual machine disk (VMDK) files as a Windows E drive in this scenario is not supported by MEDITECH.

MEDITECH Host Workload I/O Characteristic

The I/O characteristic of each MEDITECH host and the system as a whole depends on the MEDITECH platform deployed. All MEDITECH platforms (Expanse, 6.x, C/S 5.x, and MAGIC) generate workloads that are 100% random.

The MEDITECH Expanse platform generates the most demanding workload because it has the highest percentage of write operations and overall IOPS per host, followed by 6.x, C/S 5.x and the MAGIC platforms.

For more details about the MEDITECH workload descriptions, see [TR-4190: NetApp Sizing Guidelines for MEDITECH Environments](#).

Storage Network

MEDITECH requires that the Fibre Channel Protocol (FCP) be used for data traffic between the NetApp FAS or AFF system and the MEDITECH hosts of all categories.

Storage Presentation for a MEDITECH Host

Each MEDITECH host uses two Windows drives:

- **Drive C.** This drive stores the Windows Server operating system and the MEDITECH host application files.
- **Drive E.** The MEDITECH host stores its program, dictionary, and data files on drive E of the Windows Server operating system. Drive E is a LUN mapped from the NetApp FAS or AFF system using the FCP. MEDITECH requires that the FCP be used so that the MEDITECH host's IOPS and read and write latency requirements are met.

Volume and LUN Naming Convention

MEDITECH requires that a specific naming convention be used for all LUNs.

Before any storage deployment, verify the MEDITECH Hardware Configuration Proposal (HCP) to confirm the naming convention for the LUNs. The MEDITECH backup process relies on the volume and LUN naming convention to properly identify the specific LUNs to back up.

1.4 Comprehensive Management Tools and Automation Capabilities

Cisco UCS with Cisco UCS Manager

Cisco focuses on three key elements to deliver the best data center infrastructure: simplification, security, and scalability. The Cisco UCS Manager software combined with platform modularity provides a simplified, secure, and scalable desktop virtualization platform:

- **Simplified.** Cisco UCS provides a radical new approach to industry-standard computing and provides the core of the data center infrastructure for all workloads. Among the many features and benefits of Cisco UCS are the reduction in the number of servers needed, the reduction in the number of cables used per server, and the capability to rapidly deploy or re-provision servers through Cisco UCS service profiles. With fewer servers and cables to manage and with streamlined server and application workload provisioning, operations are simplified. Scores of blade and rack servers can be provisioned in minutes with Cisco UCS Manager service profiles. Cisco UCS service profiles eliminate server integration run books and eliminate configuration drift. This approach accelerates the time to productivity for end users, improves business agility, and allows IT resources to be allocated to other tasks.

Cisco UCS Manager automates many mundane, error-prone data center operations such as configuration and provisioning of server, network, and storage access infrastructure. In addition, Cisco UCS B-Series blade servers and C-Series rack servers with large memory footprints enable high application user density, which helps reduce server infrastructure requirements.

Simplification leads to a faster, more successful MEDITECH infrastructure deployment.

- **Secure.** Although VMs are inherently more secure than their physical predecessors, they introduce new security challenges. Mission-critical web and application servers using a common infrastructure such as virtual desktops are now at a higher risk for security threats. Inter-VM traffic now poses an important security consideration that IT managers need to address, especially in dynamic environments in which VMs, using VMware vMotion, move across the server infrastructure.

Virtualization, therefore, significantly increases the need for VM-level awareness of policy and security, especially given the dynamic, and fluid nature of VM mobility across an extended computing infrastructure. The ease with which new virtual desktops can proliferate magnifies the importance of a virtualization-aware network and security infrastructure. Cisco data center infrastructure (Cisco UCS, Cisco MDS, and Cisco Nexus family solutions) for desktop virtualization provides strong data center, network, and desktop security, with comprehensive security from the desktop to the hypervisor. Security is enhanced with segmentation of virtual desktops, VM-aware policies and administration, and network security across the LAN and WAN infrastructure.

- **Scalable.** Growth of virtualization solutions is all but inevitable, so a solution must be able to scale, and scale predictably, with that growth. The Cisco virtualization solutions support high VM density (VMs per server), and more servers scale with near-linear performance. Cisco data center infrastructure provides a flexible platform for growth and improves business agility. Cisco UCS Manager service profiles allow on-demand host provisioning and make it as easy to deploy dozens of hosts as it is to deploy hundreds.

Cisco UCS Servers provide near-linear performance and scale. Cisco UCS implements the patented Cisco Extended Memory Technology to offer large memory footprints with fewer sockets (with scalability to up to 1TB of memory with 2- and 4-socket servers). Using unified fabric technology as a building block, Cisco UCS Server aggregate bandwidth can scale to up to 80Gbps per server, and the northbound Cisco UCS fabric interconnect can output 2Tbps at line rate, helping prevent desktop virtualization I/O and memory bottlenecks. Cisco UCS, with its high-performance, low-latency unified fabric-based networking architecture, supports high volumes of virtual desktop traffic, including high-resolution video and communications traffic. In addition, ONTAP helps to maintain data availability and optimal performance during boot and login storms as part of the FlexPod virtualization solutions.

Cisco UCS, Cisco MDS, and Cisco Nexus data center infrastructure designs provide an excellent platform for growth, with transparent scaling of server, network, and storage resources to support desktop virtualization, data center applications, and cloud computing.

VMware vCenter Server

VMware vCenter Server provides a centralized platform for managing MEDITECH environments so that healthcare organizations can automate and deliver a virtual infrastructure with confidence:

- **Simple deployment.** Quickly and easily deploy vCenter Server using a virtual appliance.
- **Centralized control and visibility.** Administer the entire VMware vSphere infrastructure from a single location.
- **Proactive optimization.** Allocate and optimize resources for maximum efficiency.
- **Management.** Use powerful plug-ins and tools to simplify management and extend control.

Virtual Storage Console for VMware vSphere

Virtual Storage Console (VSC), vSphere API for Storage Awareness (VASA) Provider, and VMware Storage Replication Adapter (SRA) for VMware vSphere from NetApp are a virtual appliance. This product suite includes capabilities of VSC, VASA Provider, and SRA. The product suite includes SRA and VASA Provider as plug-ins to vCenter Server, which provides end-to-end lifecycle management for VMs in VMware environments using NetApp storage systems.

The virtual appliance for VSC, VASA Provider, and SRA integrates smoothly with the VMware vSphere Web Client and enables you to use SSO services. In an environment with multiple VMware vCenter Server instances, each vCenter Server instance that you want to manage must have its own registered instance of VSC. The VSC dashboard page enables you to quickly check the overall status of your datastores and VMs.

By deploying the virtual appliance for VSC, VASA Provider, and SRA, you can perform the following tasks:

- **Use VSC to deploy and manage storage and configure the ESXi host.** You can use VSC to add credentials, remove credentials, assign credentials, and set up permissions for storage controllers in your VMware environment. In addition, you can manage ESXi servers that are connected to NetApp storage systems. You can set recommended best practice values for host timeouts, NAS, and multipathing for all the hosts with a couple of clicks. You can also view storage details and collect diagnostic information.
- **Use VASA Provider to create storage capability profiles and set alarms.** VASA Provider for ONTAP is registered with VSC when you enable the VASA Provider extension. You can create and use storage capability profiles and virtual datastores. You can also set alarms to alert you when the thresholds for volumes and aggregates are almost full. You can monitor the performance of VMDKs and the VMs that are created on virtual datastores.
- **Use SRA for disaster recovery.** You can use SRA to configure protected and recovery sites in your environment for disaster recovery during failures.

NetApp OnCommand Insight and ONTAP

NetApp OnCommand Insight integrates infrastructure management into the MEDITECH service delivery chain. This approach provides healthcare organizations with better control, automation, and

analysis of the storage, network, and compute infrastructure. IT can optimize the current infrastructure for maximum benefit while simplifying the process of determining what and when to buy. It also mitigates the risks associated with complex technology migrations. Because it requires no agents, installation is straightforward and nondisruptive. Installed storage and SAN devices are continually discovered, and detailed information is collected for full visibility of your entire storage environment. You can quickly identify misused, misaligned, underused, or orphaned assets and reclaim them to fuel future expansion:

- **Optimize existing resources.** Identify misused, underused, or orphaned assets by using established best practices to avoid problems and meet service levels.
- **Make better decisions.** Real-time data helps resolve capacity problems more quickly to accurately plan future purchases, avoid overspending, and defer capital expenditures.
- **Accelerate IT initiatives.** Better understand virtual environments in order to manage risks, minimize downtime, and speed cloud deployment.

2 Design

The architecture of FlexPod for MEDITECH is based on guidance from MEDITECH, Cisco, and NetApp and from 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 small or large and whether centralized, distributed, or multitenant.

The correct storage architecture can be determined by the overall size with the total IOPS. Performance alone is not the only factor, and you might decide to go with a larger node count based on more customer requirements. The advantage of using NetApp is that the cluster can easily be scaled up nondisruptively as requirements change. You can also nondisruptively remove nodes from the cluster to repurpose or during equipment refreshes.

Here are some of the benefits of the NetApp ONTAP storage architecture:

- **Easy nondisruptive scale up and scale out.** Disks and nodes can be upgraded, added, or removed by using ONTAP nondisruptive operations. Customers can start with four nodes and move to six nodes or upgrade to larger controllers nondisruptively.
- **Storage efficiencies.** Reduce total capacity requirements with deduplication, FlexClone, inline compression, inline compaction, thin replication, thin provisioning, and aggregate deduplication. The FlexClone capability allows you to almost instantly create clones to support backup and test environment refreshes. These clones consume more storage only as changes are made.
- **DR shadow database server.** The DR shadow database server is part of a customer's business continuity strategy (used to support storage read-only [SRO] functionality and potentially configured to be a storage read/write [SRW] instance). Therefore, the placement and sizing of the third storage system are usually the same as in the production database storage system.
- **Database consistency (requires some consideration).** If SnapMirror backup copies are used in relation to business continuity, see [TR-3446: SnapMirror Async Overview and Best Practices Guide](#).

2.1 Storage Layout

Dedicated Aggregates for MEDITECH Hosts

The first step toward satisfying MEDITECH's high-performance and high-availability requirements is properly designing the storage layout for the MEDITECH environment to isolate the MEDITECH host production workload onto dedicated, high-performance storage.

One dedicated aggregate should be provisioned on each storage controller for storing the program, dictionary, and data files of the MEDITECH hosts. To eliminate the possibility of other workloads using the same disks and affecting performance, no other storage is provisioned from these aggregates.

Note: Storage provisioned for the other MEDITECH servers should not be placed on the dedicated aggregate for the LUNs used by the MEDITECH hosts. The storage for other MEDITECH servers should be placed on a separate aggregate. Storage requirements for other

MEDITECH servers are available in the “Hardware Configuration Proposal” (for new deployments) and “Hardware Evaluation Task” (for existing deployments) documents. You can obtain these documents from MEDITECH through the MEDITECH system integrator or the MEDITECH Technical Account Manager (TAM). NetApp system engineers might consult with the NetApp MEDITECH Independent Software Vendor (ISV) team to facilitate a proper and complete NetApp storage sizing configuration.

Spread MEDITECH Host Workload Evenly Across All Storage Controllers

NetApp FAS and AFF storage systems are deployed as one or more high-availability pairs. NetApp recommends spreading the MEDITECH Expanse and 6.x workloads evenly across each storage controller to apply the compute, network, and caching resources on each storage controller.

Use the following guidelines to spread the MEDITECH workloads evenly across each storage controller:

- If you know the IOPS for each MEDITECH host, you can spread the MEDITECH Expanse and 6.x workload evenly across all storage controllers by confirming that each controller services a similar number of IOPS from the MEDITECH hosts.
- If you do not know the IOPS for each MEDITECH host, you can spread the MEDITECH Expanse and 6.x workload evenly across all storage controllers. Complete this task by confirming that the capacity of the aggregates for the MEDITECH hosts is evenly distributed across all storage controllers. By doing so, the number of disks is the same across all data aggregates dedicated to the MEDITECH hosts.
- Use similar disk types and identical RAID groups to create the storage aggregates of both controllers for distributing the workloads equally. Refer to the platform sizing guide before creating the storage aggregate.

Note: According to MEDITECH, there are two hosts in the MEDITECH system that generate higher IOPS than the rest of the hosts. The LUNs for these two hosts should be placed on separate storage controllers. These two hosts should be identified with the assistance of the MEDITECH team before deployment.

2.2 Storage Placement

Database Storage for MEDITECH Hosts

The database storage for a MEDITECH host is presented as a block device (that is, a LUN) from the NetApp FAS or AFF system. The LUN is typically mounted to the Windows operating system as the E drive.

Other Storage

The MEDITECH host operating system and the database application normally generate a considerable amount of IOPS on the storage. Storage provisioning for the MEDITECH host VMs and their VMDK files, if necessary, is considered independent from the storage requirements needed to meet the MEDITECH performance thresholds.

Storage provisioned for the other MEDITECH servers should not be placed on the dedicated aggregate for the LUNs used by the MEDITECH hosts. Place the storage for other MEDITECH servers on a separate aggregate.

2.3 Storage Controller Configuration

High Availability

Storage systems should be configured with controllers in a high-availability pair in the high-availability mode to mitigate the effect of controller failure and enable nondisruptive upgrades of the storage system.

With the high-availability controller pair configuration, disk shelves should be connected to controllers by multiple paths. This connection increases storage resiliency by protecting against a single-path failure while providing improved performance consistency if there is a controller failover.

Storage Performance During Storage Controller Failover

For storage systems configured with controllers in a high-availability pair, in the unlikely event of a controller failure, the partner controller takes over the failed controller's storage resource and workload. It is important to consult the customer to determine the performance requirements that must be met if there is a controller failure and to size the system accordingly.

Hardware-Assisted Takeover

NetApp recommends turning on the hardware-assisted takeover feature on both storage controllers.

Hardware-assisted takeover is designed to minimize the storage controller failover time. It enables one controller's Remote LAN Module or Service Processor module to notify its partner of a controller failure faster than a heartbeat timeout trigger, reducing the time elapsed before failover. The hardware-assisted takeover feature is enabled by default for storage controllers in a high-availability configuration.

For more information about hardware-assisted takeover, see the [Clustered Data ONTAP 8.3 High-Availability Configuration Guide](#) or the [Data ONTAP 8.2 High-Availability and MetroCluster Configuration Guide for 7-Mode](#). As of this writing, the latest Data ONTAP version is 8.3. Refer to the Data ONTAP manuals corresponding to the Data ONTAP version being deployed. Refer to the [ONTAP 9 Documentation Center](#) for all other ONTAP documentation.

Disk Type

NetApp recommends using a high-performance SSD disk type for aggregates on AFF systems that are dedicated for the MEDITECH hosts to support MEDITECH workloads' low read latency requirement.

NetApp AFF

NetApp offers high-performance AFF arrays to address MEDITECH workloads that demand high throughput and that have random data access patterns and low latency requirements. For MEDITECH workloads, AFF arrays offer performance advantages over systems based on HDDs. The combination of flash technology and enterprise data management delivers advantages in three major areas: performance, availability, and storage efficiency.

NetApp Support Tools and Services

NetApp offers a complete set of support tools and services. The NetApp AutoSupport™ tool should be enabled and configured on NetApp FAS systems to call home if there is a hardware failure or system misconfiguration. Calling home alerts the NetApp Support team to remediate any issues in a timely manner.

3 Deployment and Configuration

The NetApp storage FlexPod deployment guidance 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 \(CVD\)](#).
- 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).

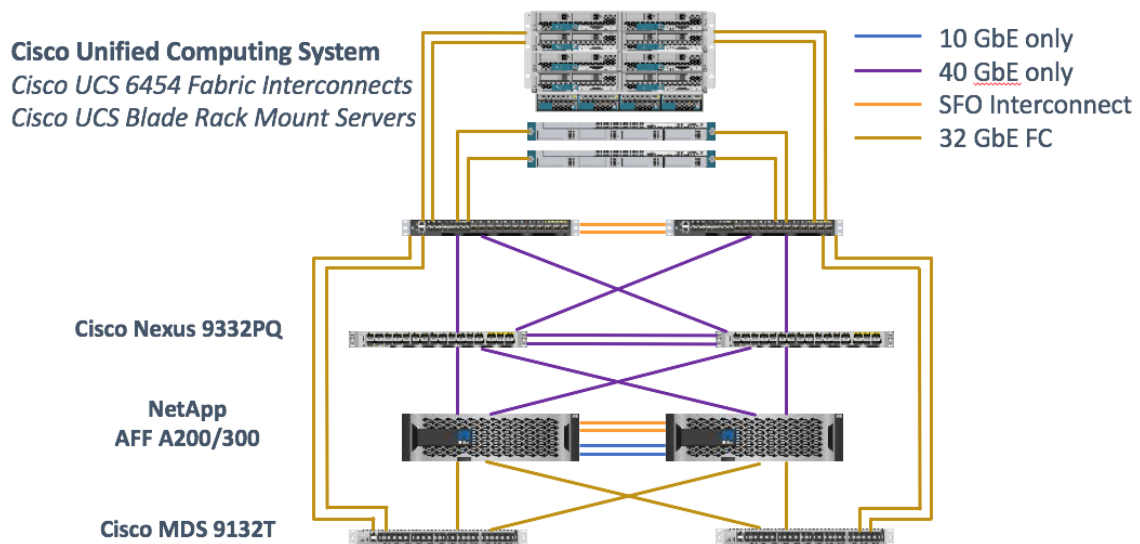
- Quantitative performance requirements and sizing guidance.
For more information, see [TR-4190: NetApp Sizing Guidelines for MEDITECH Environments](#).
- Use of NetApp SnapMirror technologies to meet backup and disaster recovery requirements.
- Generic NetApp storage deployment guidance.

This section provides an example configuration with infrastructure deployment best practices and lists the various infrastructure hardware and software components and the versions that could potentially be used.

3.1 Cabling Diagram

Figure 4 illustrates the 32Gb FC/40GbE topology diagram for an MEDITECH deployment.

Figure 4) FlexPod for MEDITECH physical topology diagram.



Always use the [Interoperability Matrix Tool \(IMT\)](#) to validate that all versions of software and firmware are supported. Table 1 lists the infrastructure hardware and software components that were used in the solution testing.

3.2 Base Infrastructure Configuration

Network Connectivity

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

- Link aggregation using 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 vNICs with redundant connectivity to the unified fabric. Network interface card (NIC) failover is used between FI for redundancy.
 - Each server has 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 configuring the infrastructure:

- Storage ports ifgroups (vPC)
- 10G link to switch N9k-A
- 10G link to switch N9k-B
- In-band management (active-passive bond):
 - 1G link to management switch N9k-A
 - 1G link to management switch N9k-B
- 32G 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 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

3.3 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. Cisco recommends the use of hot spare blade servers for optimal performance.

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 FC Storage.

Cisco Nexus Switch Configuration

A fault-tolerant pair of Cisco Nexus 9300 Series Ethernet switches is deployed for the solution. These switches should be cabled as described in section 3.1, “Cabling Diagram.” The Cisco Nexus configuration makes sure that Ethernet traffic flows are optimized for the MEDITECH application.

1. After the initial setup and licensing are completed, 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 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 FC Cisco Validated Design](#).

Cisco MDS 9132T Configuration

The Cisco MDS 9100 Series FC switches provide redundant 32Gb FC connectivity between the NetApp AFF A200/300 controllers and the Cisco UCS compute fabric. The cables should be connected as described in section 3.1, “Cabling Diagram.”

1. From the switch 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 \(CVD\)](#).
3. To create the necessary virtual SANs (VSANs) for the solution, complete the following steps while in global configuration mode:
 - a. For 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
```

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

- b. For 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
```

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

4. For each FC switch, create device alias names that make identifying each device intuitive for ongoing operations using the details in the reference document.
5. Finally, create the FC zones by using the device alias names created in step 4 for each MDS switch as follows:
 - a. For 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>

```

– For 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

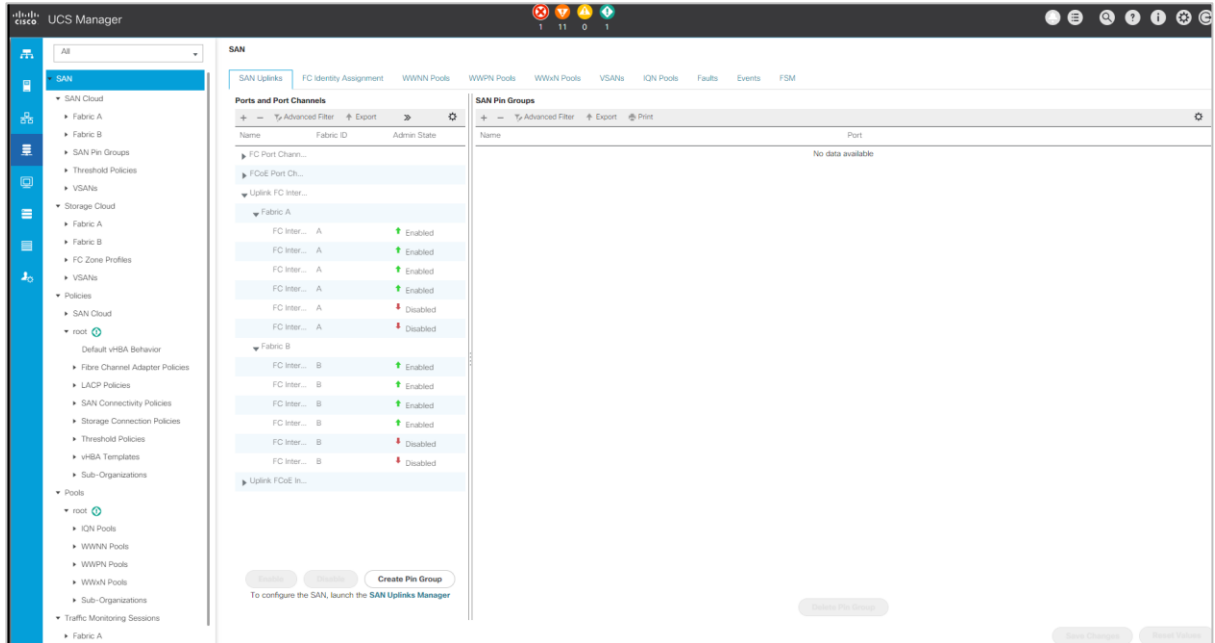
Cisco UCS allows MEDITECH customers to leverage their subject matter experts in network, storage, and compute to create policies and templates that tailor the environment to their specific needs. After being created, these policies and templates can be combined into service profiles that deliver consistent, 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 HTML 5 GUI
- Cisco UCS CLI
- Cisco UCS Central for multidomain environments

Figure 5 shows a sample screenshot of the SAN node in Cisco UCS Manager.

Figure 5) Cisco UCS Manager HTML5 UI.



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.

Complete the following procedures in order to set up the Cisco UCS compute platform. Perform these procedures after the Cisco UCS B200 M5 blade servers are installed in the Cisco UCS 5108AC blade chassis. Also, the cabling requirements must be completed as described in section 3.1, “Cabling Diagram.”

1. Upgrade the Cisco UCS Manager firmware to version 3.2(2f) or later.
2. Configure the reporting, 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, UUIDs, MAC address, servers, WWNN, and 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 virtual host bus adapter (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 FC Cisco Validated Design \(CVD\)](#) document.

3.4 ESXi Configuration Best Practices

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

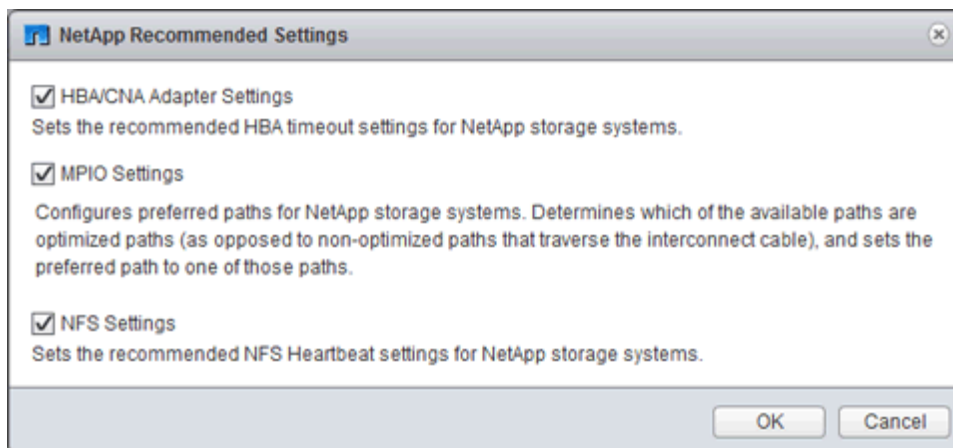
- VSC for VMware vSphere checks and sets the ESXi host multipathing settings and HBA timeout settings that work best with NetApp storage systems. The values that VSC sets are based on rigorous internal testing by NetApp.
- For the best storage performance, customers should consider using VMware vStorage APIs for Array Integration (VAAI)–capable storage hardware. 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.

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.

3.5 NetApp Configuration

NetApp storage deployed for MEDITECH software environments uses storage controllers in a high-availability pair configuration. Storage is required to be presented from both controllers to MEDITECH database servers over the FC Protocol (FCP). 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 using the relevant ONTAP commands. The emphasis is to show how storage is provisioned to implement the storage layout recommended by NetApp, which uses an HA controller pair. One of the major advantages with ONTAP is the ability to scale out without disturbing the existing HA pairs.

ONTAP Licenses

After the storage controllers are set up, apply licenses to enable ONTAP features recommended by NetApp. The licenses for MEDITECH workloads would be FC, CIFS, Snapshot, SnapRestore®, FlexClone, and SnapMirror technologies.

1. Open NetApp 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 AutoSupport tool sends summary support information to NetApp through HTTPS. To configure AutoSupport, complete the following step:

1. Run the following ONTAP commands in order to configure AutoSupport:

```
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 required to initiate a takeover following the unlikely failure of a controller. To configure hardware-assisted takeover, complete the following steps:

1. Run the following ONTAP command.

Note: 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:

Note: 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 `prod1-01` and `prod1-02` HA controller pair.

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

3.6 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 in a specific site. For the number of disks that NetApp recommends using in a RAID group, see [NetApp TR-3838: Storage Subsystem Configuration Guide](#).

The RAID group size is important for storage expansion because NetApp recommends adding 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. The number of data disks required is determined by using the NetApp System Performance Modeler (SPM) sizing tool. After deciding on 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 using the SPM sizing tool for MEDITECH environments, see [NetApp TR-4190: NetApp Sizing Guidelines for MEDITECH Environments](#).

Storage Expansion Considerations

When expanding aggregates with more disks, add disks in groups equal to the aggregate RAID group size. Doing so helps provide performance consistency throughout the aggregate.

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

After aggregate expansion, 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.

Note: Reallocation of schedules should be planned during nonproduction hours because it is a high-CPU and disk-intensive task.

For more information about using reallocate 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.

3.7 Storage Virtual Machine Configuration

This section pertains to deployment on clustered Data ONTAP 8.3 and later.

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

SVM for MEDITECH Host LUNs

One dedicated SVM per Data ONTAP storage cluster should be created to own and manage the aggregates that contain the LUNs for the MEDITECH hosts.

SVM Language Encoding Setting

NetApp recommends setting the language encoding for all SVMs. If no language encoding setting is specified at the time 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 associated with the SVM inherit the SVM language encoding setting unless explicitly specified when the volumes are created. The language encoding setting should be used consistently in all volumes of a site to enable certain operations to work. For example, SnapMirror requires the source and destination SVM to have the same language encoding setting.

3.8 Volume Configuration

Volume Provisioning

MEDITECH volumes dedicated for MEDITECH hosts 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, NetApp recommends disabling 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 that the volume has enough space for the MEDITECH data LUNs and FlexClone volumes that will be created by the backup servers. FlexClone volumes do not occupy more space the way 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 using Flash Pool or NetApp Flash Cache™, 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 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. Volume-level reallocate can be used 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 eight hours.

Best Practice

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

Note: NetApp highly recommends running the reallocate task on one volume at a time per controller.

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

3.9 LUN Configuration

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

LUN Provisioning

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

LUN Operating System Type

You must correctly set the LUNs' operating system type to properly align the LUNs that are created. 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 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, refer to 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 FCP. In the VMware virtual environment, the LUNs are presented to the VMware ESXi servers hosting the MEDITECH hosts. Then each LUN presented to the VMware ESXi server is mapped to each MEDITECH host VM using raw device mappings in the physical compatibility mode.

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

Refer to the MEDITECH HCP when consulting with the MEDITECH and backup system installer to devise a consistent naming convention for the LUNs used by the MEDITECH hosts.

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

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

3.10 Initiator Group Configuration

When using FC as the data network protocol, create two initiator groups (igroups) on each storage controller. The first initiator group contains the worldwide port names (WWPNs) of the FC host interface cards on the VMware ESXi servers hosting the MEDITECH host VMs (igroup for MEDITECH).

The MEDITECH igroup operating system type must be set 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.

Note: The operating system type for an igroup might be different from the operating system type for a LUN. As an example, the igroup operating system type should be set to `VMware` for virtualized MEDITECH hosts. The operating system type setting for the LUNs used by the virtualized MEDITECH hosts should be set to `Windows 2008 or later`. That is because the MEDITECH host operating system is the Windows 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](#).

3.11 LUN Mappings

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

Appendix A: MEDITECH Modules and Components

The MEDITECH application covers several modules and components. Functions covered by these modules are listed Table 1. For additional information about setting up and deploying these modules, see MEDITECH documentation.

Table 1) MEDITECH modules and components.

Function	Type
Connectivity	<ul style="list-style-type: none"> • Web server • Live application server (WI) • TEST application server (WI) • SAML authentication server (WI) • SAML proxy server (WI) • Database server
Infrastructure	<ul style="list-style-type: none"> • File server • Background job client • Connection server • Transaction server
Scanning and archiving	<ul style="list-style-type: none"> • Image server
Data repository	<ul style="list-style-type: none"> • SQL server
Business and clinical analytics	<ul style="list-style-type: none"> • LIVE intelligence server (BCA) • TEST intelligence server (BCA) • Database server (BCA)
Home care	<ul style="list-style-type: none"> • Remote site solution • Connectivity • Infrastructure • Printing • Field devices • Scanning • Hosted site requirements • Firewall configuration
Support	<ul style="list-style-type: none"> • Background Job Client (CALs)
User devices	<ul style="list-style-type: none"> • Tablets • Fixed devices
Printing	<ul style="list-style-type: none"> • LIVE Network Print Server (required; might already exist) • TEST Network Print Server (required; might already exist)
Third-party requirement	<ul style="list-style-type: none"> • First Databank MedKnowledge Framework v4.3

Acknowledgements

- Brandon Agee, Technical Marketing Engineer, NetApp
- Atul Bhalodia, Technical Marketing Engineer, NetApp
- Ketan Mota, Product Manager, NetApp
- John Duignan, Solutions Architect, NetApp

- Jon Ebmeier, Cisco
- Mike Brennan, Cisco

Where to Find Additional Information

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

FlexPod Design Zone

- [NetApp FlexPod Design Zone](#)
- [FlexPod DC with FC Storage \(MDS Switches\) Using NetApp AFF, vSphere 6.5U1, and Cisco UCS Manager](#)

NetApp Technical Reports

- [TR-3929: Reallocate Best Practices Guide](#)
- [TR-3987: Snap Creator Framework Plug-In for InterSystems Caché](#)
- [TR-4300i: NetApp FAS and All-Flash Storage Systems for MEDITECH Environments Best Practices Guide](#)
- [TR-4017: FC SAN Best Practices](#)
- [TR-3446: SnapMirror Async Overview and Best Practices Guide](#)

ONTAP Documentation

- [NetApp Product Documentation](#)
- [Virtual Storage Console \(VSC\) for vSphere documentation](#)
- [ONTAP 9 Documentation Center](#)
 - [FC Express Guide for ESXi](#)
- [All ONTAP 9.3 Documentation:](#)
 - [Software Setup Guide](#)
 - [Disks and Aggregates Power Guide](#)
 - [SAN Administration Guide](#)
 - [SAN Configuration Guide](#)
 - [FC Configuration for Windows Express Guide](#)
 - [FC SAN Optimized AFF Setup Guide](#)
 - [High-Availability Configuration Guide](#)
 - [Logical Storage Management Guide](#)
 - [Performance Management Power Guide](#)
 - [SMB/CIFS Configuration Power Guide](#)
 - [SMB/CIFS Reference](#)
 - [Data Protection Power Guide](#)
 - [Data Protection Tape Backup and Recovery Guide](#)
 - [NetApp Encryption Power Guide](#)
 - [Network Management Guide](#)
 - [Commands: Manual Page Reference for ONTAP 9.3](#)

Cisco Nexus, MDS, Cisco UCS, and Cisco UCS Manager Guides

- [Cisco UCS Servers Overview](#)
- [Cisco UCS Blade Servers Overview](#)
- [Cisco UCS B200 M5 Datasheet](#)

- [Cisco UCS Manager Overview](#)
- [Cisco UCS Manager 3.2\(3a\) Infrastructure Bundle](#) (requires Cisco.com authorization)
- [Cisco Nexus 9300 Platform Switches](#)
- [Cisco MDS 9132T FC Switch](#)

Version History

Version	Date	Document Version History
Version 1.0	February 2019	Initial release.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2019 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.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS DOCUMENT ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO, ALL PRODUCT VENDORS OR MANUFACTURERS IDENTIFIED OR REFERENCED HEREIN ("PARTNERS") AND THEIR RESPECTIVE SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, OR WITH RESPECT TO ANY RESULTS THAT MAY BE OBTAINED THROUGH USE OF THE DESIGNS OR RELIANCE UPON THIS DOCUMENT, EVEN IF CISCO, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS AND USE OR RELIANCE UPON THIS DOCUMENT. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO OR ITS PARTNERS.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS DOCUMENT ARE PRESENTED "AS IS," WITH ALL FAULTS. NETAPP, ALL PRODUCT VENDORS OR MANUFACTURERS IDENTIFIED OR REFERENCED HEREIN ("PARTNERS") AND THEIR RESPECTIVE SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL NETAPP, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, OR WITH RESPECT TO ANY RESULTS THAT MAY BE OBTAINED THROUGH USE OF THE DESIGNS OR RELIANCE UPON THIS DOCUMENT, EVEN IF NETAPP, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS AND USE OR RELIANCE UPON THIS DOCUMENT. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF NETAPP, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY NETAPP OR ITS PARTNERS.