



FlexPod Datacenter for MEDITECH Deployment Guide

FlexPod

NetApp
January 21, 2025

This PDF was generated from https://docs.netapp.com/us-en/flexpod/healthcare/ehr-meditech-deploy_overview.html on January 21, 2025. Always check docs.netapp.com for the latest.

Table of Contents

- FlexPod Datacenter for MEDITECH Deployment Guide 1
 - TR-4753: FlexPod Datacenter for MEDITECH Deployment Guide 1
 - Design 12
 - Deployment and configuration 15
 - MEDITECH modules and components 28
 - Acknowledgments 30
 - Where to find additional information. 30

FlexPod Datacenter for MEDITECH Deployment Guide

TR-4753: FlexPod Datacenter for MEDITECH Deployment Guide

Brandon Agee and John Duignan, NetApp
Mike Brennan and Jon Ebmeier, Cisco

In partnership with:



Overall solution benefits

By running a MEDITECH environment on the FlexPod architectural foundation, your healthcare organization can expect an improvement in staff productivity and a decrease in capital and operational expenditures. FlexPod Datacenter for MEDITECH delivers several benefits that are specific to the healthcare industry, including:

- **Simplified operations and lowered costs.** 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 higher resource utilization for greater return on investment (ROI).
- **Faster deployment of infrastructure.** Whether it's an existing data center or a remote location, with the integrated and tested design of FlexPod Datacenter, you can have your new infrastructure up and running in less time, with less effort.
- **Certified storage.** NetApp ONTAP data management software with MEDITECH gives you the superior reliability of a tested and certified storage vendor. MEDITECH does not certify other infrastructure components.
- **Scale-out architecture.** Scale SAN and NAS from terabytes (TB) to tens of petabytes (PB) without reconfiguring running applications.
- **Nondisruptive operations.** Perform storage maintenance, hardware lifecycle operations, and FlexPod upgrades without interrupting the business.
- **Secure multitenancy.** Support the increased needs of virtualized server and storage shared infrastructure, enabling secure multitenancy of facility-specific information, particularly if your system hosts 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](#).
- **Agility.** With the industry-leading workflow automation, orchestration, and management tools that FlexPod systems provide, your IT team can 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.

- **Increased 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 is used in multiple locations, and often you need to leverage and to share it with other locations, applications, and infrastructures. You need a way to manage your data that is consistent and integrated. The Data Fabric provides a way to manage data that puts IT in control and that simplifies ever-increasing IT complexity.

FlexPod

New infrastructure approach for MEDITECH EHRs

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

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

FlexPod is a prevalidated, rigorously tested converged infrastructure from the strategic partnership of Cisco and NetApp. It is engineered and designed specifically to deliver predictable low-latency system performance and high availability. This approach results in MEDITECH compliance and ultimately optimal 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 meets 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 a storage management fabric and a cohesive management toolset.
- **Industry-leading technology at each level of the converged stack.** Cisco, NetApp, VMware, and Microsoft Windows are all ranked as number 1 or number 2 by industry analysts in their respective categories of servers, networking, storage, and operating systems.

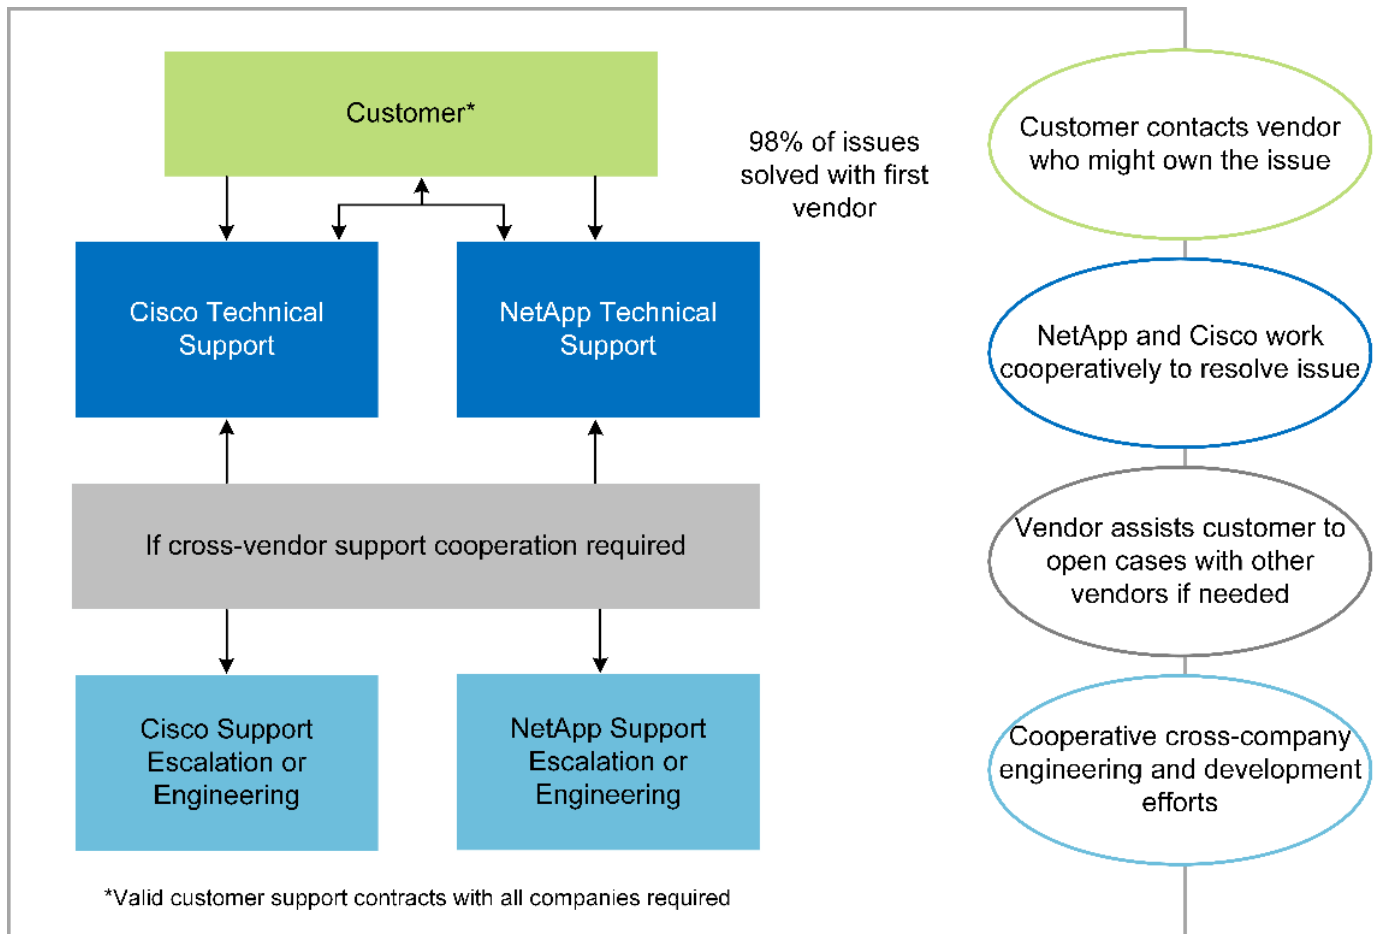
- **Investment protection with standardized, flexible IT.** The FlexPod reference architecture anticipates new product versions and updates, with rigorous ongoing interoperability testing to accommodate future technologies as they become available.
- **Proven deployment across a broad range of environments.** 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 your growing workload demands without negatively 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 your MEDITECH Hardware Configuration Proposal requirements, different ONTAP platforms can be deployed to meet your 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 which stores data more efficiently while driving more performance from MEDITECH software.
- **Reduced risk.** Minimize business disruption with a prevalidated platform that is 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 meet 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 your FlexPod support issue, regardless of where the problem resides. With the FlexPod Cooperative Support model, your FlexPod system operates efficiently and benefits from the most up-to-date technology, and you work with an experienced team to help you resolve integration issues.

FlexPod Cooperative Support is especially valuable to healthcare organizations that run business-critical applications such as MEDITECH on the FlexPod converged infrastructure. The following figure illustrates the 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 (Cisco UCS) consists of a single management domain that is interconnected with a unified I/O infrastructure. So that the infrastructure can deliver critical patient information with maximum availability, Cisco UCS for MEDITECH environments has been aligned with MEDITECH infrastructure recommendations and best practices.

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 help you meet your 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 Technology (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. In traditional systems, redundant devices such as Ethernet and FC adapters and chassis management modules result in layers of complexity. 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 in 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.

The use of service profiles helps ensure that servers are configured consistently throughout the enterprise. When multiple Cisco UCS management domains are employed, 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.

To demonstrate that it meets the server configuration requirements, Cisco UCS has been extensively tested with MEDITECH over a multiyear period. 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 cores. 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. This traffic includes storage I/O, streamed desktop traffic, management, and access to clinical and business applications. You get:

- **Infrastructure scalability.** Virtualization, efficient power and cooling, cloud scale with automation, high density, and high 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 for 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 provides an optimal compute, networking, and SAN connectivity solution for MEDITECH.

NetApp ONTAP

NetApp storage that runs ONTAP software reduces your overall storage costs while it delivers the low-latency read and write response times and IOPS that MEDITECH workloads need. ONTAP supports both all-flash and hybrid storage configurations to create an optimal storage platform that meets MEDITECH requirements. NetApp flash-accelerated systems have received MEDITECH's validation and certification, giving you as a MEDITECH customer the performance and responsiveness that are key to latency-sensitive MEDITECH operations. By creating multiple fault domains in a single cluster, NetApp systems can also isolate production from nonproduction. NetApp systems also reduce performance issues with a guaranteed performance level minimum for workloads with ONTAP 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 that clinical applications need while also 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 that are suitable for storing large datasets with high throughput. Along with a MEDITECH-approved backup solution, you can clone, replicate, and back up your MEDITECH environment from expensive solid-state drive (SSD) storage to more economical HDD storage on other nodes. This approach meets or exceeds MEDITECH guidelines for SAN-based cloning and backup of production pools.

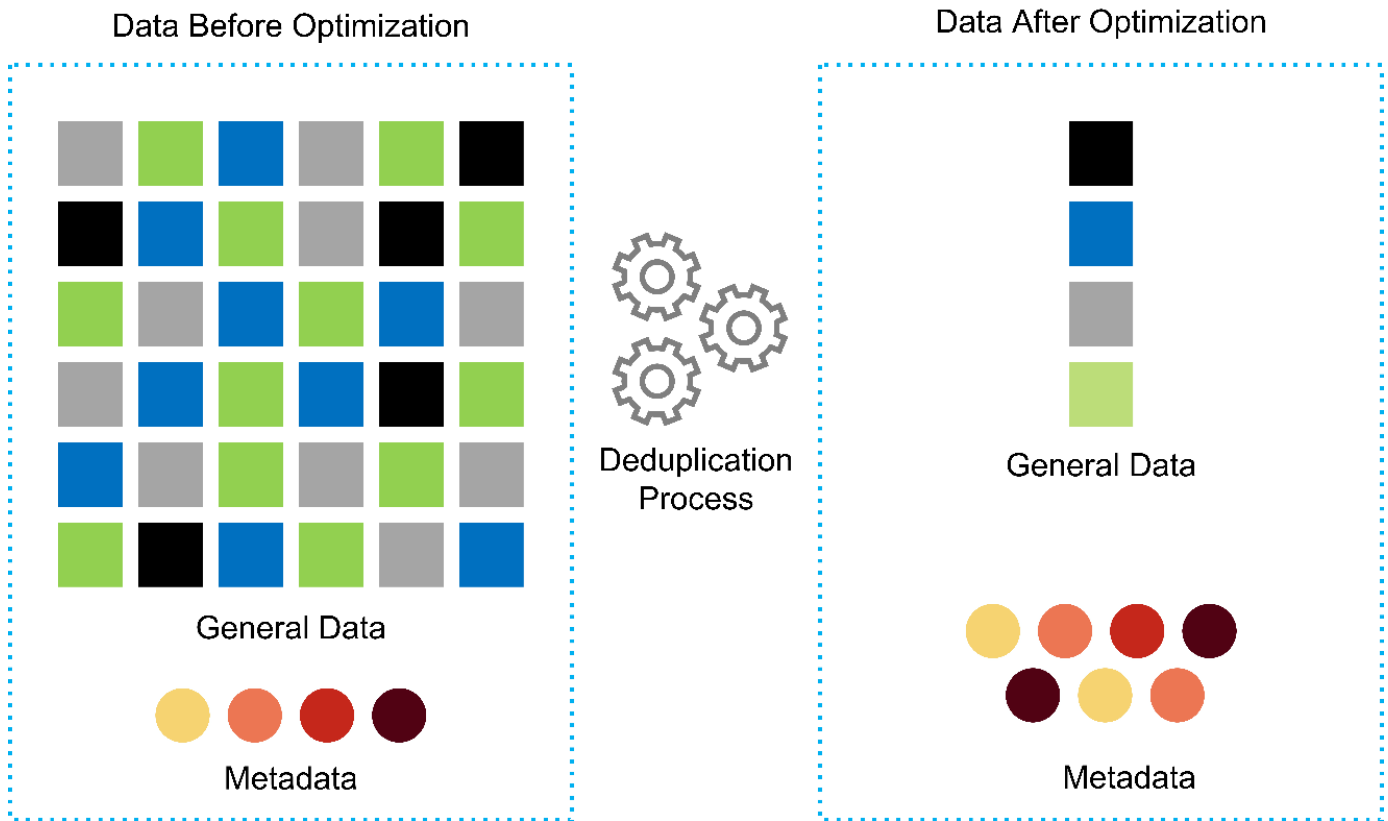
Many of the ONTAP features are especially useful in MEDITECH environments: simplifying management, increasing availability and automation, and reducing the total amount of storage needed. With these features, you get:

- **Outstanding performance.** The NetApp AFF solution shares the Unified Storage Architecture, ONTAP software, management interface, rich data services, and advanced feature set that the rest of the NetApp FAS product families have. 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 quality of 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.

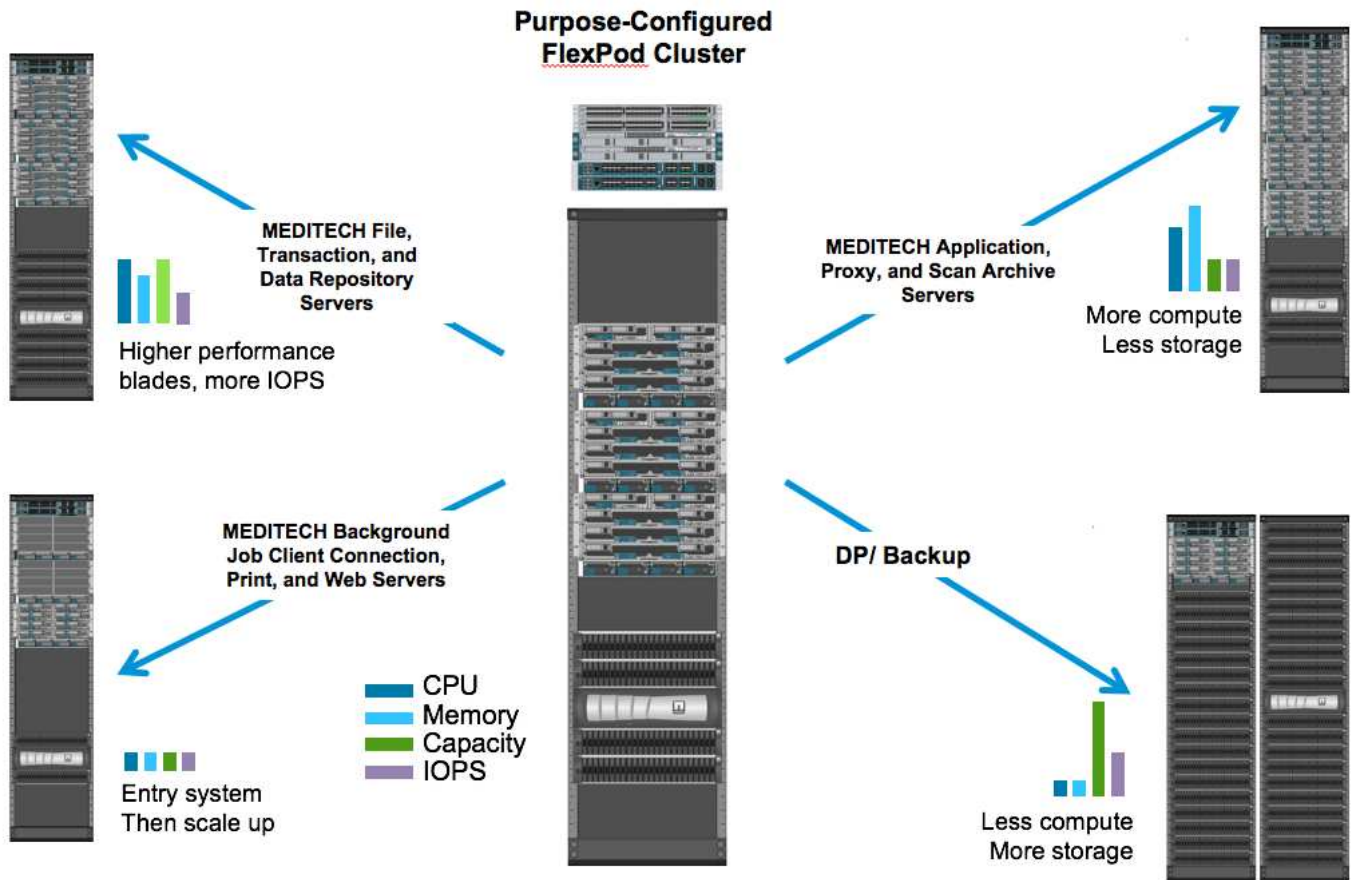
Deduplication works with a high degree of granularity and operates on the active file system of the FlexVol volume or data constituent. It is application transparent; therefore, you can use it to deduplicate data that originates from any application that uses the NetApp system. You can run volume deduplication as an inline process (starting in ONTAP 8.3.2). You can also run it as a background process that you can configure to run automatically, to be scheduled, or to run manually through the CLI, NetApp ONTAP System Manager, or NetApp Active IQ Unified Manager.

The following figure illustrates how NetApp deduplication works at the highest level.



- **Space-efficient cloning.** The FlexClone capability enables you to almost instantly create clones to support backup and testing environment refresh. These clones consume more storage only as changes are made.
- **NetApp Snapshot and SnapMirror technologies.** ONTAP can create space-efficient Snapshot copies of the logical unit numbers (LUNs) that the MEDITECH host uses. For dual-site deployments, you can implement SnapMirror software for more data replication and resiliency.
- **Integrated data protection.** Full data protection and disaster recovery features help you protect critical data assets and provide disaster recovery.
- **Nondisruptive operations.** You can perform upgrades and maintenance without taking data offline.
- **QoS and adaptive QoS (AQoS).** Storage QoS enables you to limit potential bully workloads. More important, QoS can guarantee a performance minimum for critical workloads such as MEDITECH production. By limiting contention, NetApp QoS can reduce performance-related issues. AQoS works with predefined policy groups, which you can apply directly to a volume. These policy groups can 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 environments 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 Cloud Platform, and IBM Cloud clouds, giving you a wide breadth of choice.

The following figure illustrates the FlexPod architecture for MEDITECH workloads.



MEDITECH overview

Medical Information Technology, Inc., commonly known as MEDITECH, is a Massachusetts-based software company that provides information systems for healthcare organizations. MEDITECH provides an EHR system that is designed to store and to 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 that MEDITECH software supports. Appendix A provides more information about these broad sets of MEDITECH functions. MEDITECH applications require several VMs to support these functions. To deploy these applications, see the recommendations from MEDITECH.

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

BridgeHead and Commvault are the two backup software applications that are 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 meet the performance-driven requirements for the MEDITECH database and the backup requirements 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 to enable FlexPod Datacenter to be successfully configured, deployed, and supported to meet the MEDITECH production environment requirements of customers like you.

MEDITECH categories

MEDITECH associates the deployment size with a category number that ranges 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 5.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. The previous figure 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, the 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 from your 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 as a MAGIC machine (for the MAGIC platform). This document uses the term MEDITECH host to refer to a MEDITECH file server or a MAGIC machine.

MEDITECH hosts can be physical servers or VMs that run on the Microsoft Windows Server operating system. Most commonly in the field, MEDITECH hosts are deployed as Windows VMs that run on a VMware ESXi server. As of this writing, VMware is the only hypervisor that MEDITECH supports. 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 that is attached to the VM by way of a raw device mapping (RDM) 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 that you deploy. 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 FC Protocol 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 that is mapped from the NetApp FAS or AFF system by using the FC Protocol. MEDITECH requires that the FC Protocol 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 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.

Comprehensive management tools and automation capabilities

Cisco UCS with Cisco UCS Manager

Cisco focuses on three key elements to deliver a superior 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. Cisco UCS offers many features and benefits, including reduction in the number of servers that you need and reduction in the number of cables that are used per server. Another important feature is the capability to rapidly deploy or to reprovision 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 runbooks 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 that use 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 your IT managers must 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 to 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 hundreds of hosts as it is to deploy dozens.

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 of up to 1TB of memory with 2- and 4-socket servers). By using Unified Fabric technology as a building block, Cisco UCS Server aggregate bandwidth can scale up to 80Gbps per server, and the northbound Cisco UCS Fabric Interconnect can output 2Tbps at line rate. This capability helps 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. You get 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 your healthcare organization can automate and deliver a virtual infrastructure with confidence:

- **Simple deployment.** Quickly and easily deploy vCenter Server by 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 to 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 make up a single virtual appliance. 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 that use 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 to configure the ESXi host.** You can use VSC to add credentials, to remove credentials, to assign credentials, and to set up permissions for storage controllers in your VMware environment. In addition, you can manage ESXi servers that are connected to NetApp storage systems. With a couple clicks, you can set recommended best practice values for host timeouts, NAS, and multipathing for all the hosts. You can also view storage details and collect diagnostic information.
- **Use VASA Provider to create storage capability profiles and to 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 gives your healthcare organization better control, automation, and analysis of your storage, network, and compute infrastructure. IT can optimize your current infrastructure for maximum benefit while simplifying the process of determining what and when to buy. It also mitigates the risks that are 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. OnCommand Insight helps you:

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

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 your data center strategy; the size of your organization; and whether your system is 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 use a larger node count based on additional customer requirements. The advantage of using NetApp storage is that you can easily and nondisruptively scale up the cluster as your requirements change. You can also nondisruptively remove nodes from the cluster to repurpose equipment or during equipment refreshes.

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

- **Easy, nondisruptive scale-up and scale-out.** You can upgrade, add, or remove disks and nodes by using ONTAP nondisruptive operations. You can start with four nodes and move to six nodes or upgrade to larger controllers nondisruptively.

- **Storage efficiencies.** Reduce your total capacity requirements with deduplication, NetApp FlexClone, inline compression, inline compaction, thin replication, thin provisioning, and aggregate deduplication. The FlexClone capability enables you to almost instantly create clones to support backup and testing environment refreshes. These clones consume more storage only as changes are made.
- **Disaster recovery shadow database server.** The disaster recovery shadow database server is part of your business continuity strategy (used to support storage read-only functionality and potentially configured to be a storage read/write instance). Therefore, the placement and sizing of the third storage system are usually the same as in your production database storage system.
- **Database consistency (requires some consideration).** If you use NetApp SnapMirror backup copies in relation to business continuity, see [TR-3446: SnapMirror Async Overview and Best Practices Guide](#).

Storage layout

Dedicated aggregates for MEDITECH hosts

The first step toward meeting MEDITECH's high-performance and high-availability requirements is to properly design 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.



Storage that you provision for the other MEDITECH servers should not be placed on the dedicated aggregate for the LUNs that are used by the MEDITECH hosts. You should place the storage for other MEDITECH servers 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 from your MEDITECH Technical Account Manager (TAM). NetApp solutions 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 systems are deployed as one or more high-availability pairs. NetApp recommends that you spread 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 workloads 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 still spread the MEDITECH Expanse and 6.x workloads 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 that are 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. Before you create the storage aggregate, contact a NetApp Certified Integrator.



According to MEDITECH, two hosts in the MEDITECH system generate higher IOPS than the rest of the hosts. The LUNs for these two hosts should be placed on separate storage controllers. You should identify these two hosts with the assistance of the MEDITECH team before you deploy your system.

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 that is required to meet the MEDITECH performance thresholds.

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

Storage controller configuration

High availability

To mitigate the effect of controller failure and to enable nondisruptive upgrades of the storage system, you should configure your storage system with controllers in a high-availability pair in the high-availability mode.

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, and it improves performance consistency if a controller failover occurs.

Storage performance during storage controller failover

For storage systems that are 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 resources and workloads. 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 that you turn 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 about a controller failure faster than a heartbeat timeout trigger can, reducing the time that it takes to 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 [ONTAP 9 Documentation Center](#).

Disk type

To support the low read latency requirement of MEDITECH workloads, NetApp recommends that you use a high-performance SSD for aggregates on AFF systems that are dedicated for the MEDITECH hosts.

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 that are 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 AFF/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. NetApp Active IQ is a web-based application that is based on AutoSupport information from your NetApp systems providing predictive and proactive insight to help improve availability, efficiency, and performance.

Deployment and configuration

Overview

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

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

This document does not cover:

- Detailed deployment of the FlexPod Datacenter environment

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

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

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

- Quantitative performance requirements and sizing guidance.

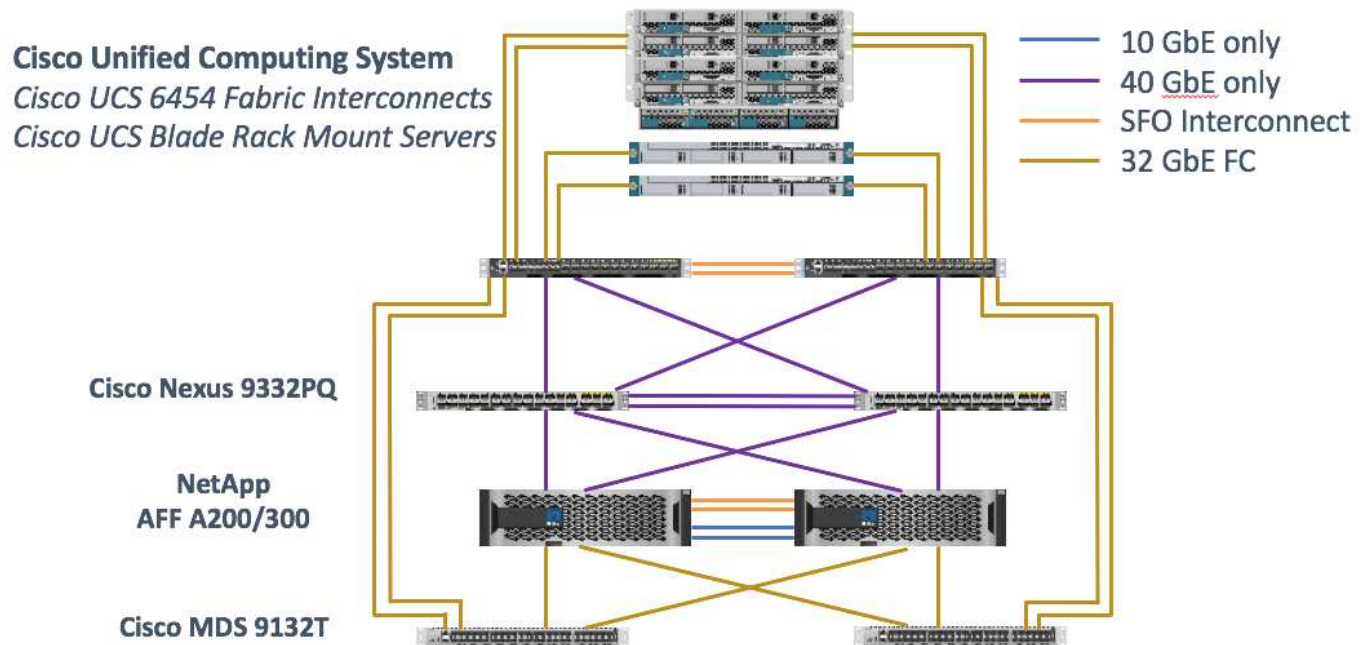
For more information, see [TR-4190: NetApp Sizing Guidelines for MEDITECH Environments](#).

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

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

Cabling diagram

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



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

Next: [Base infrastructure Configuration](#).

Base infrastructure configuration

Network connectivity

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

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

Storage connectivity

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

- Storage port interface groups (ifgroups, vPC)
- 10Gb link to switch N9K-A

- 10Gb link to switch N9K-B
- In- band management (active-passive bond):
 - 1Gb link to management switch N9K-A
 - 1Gb link to management switch N9K-B
- 32Gb FC end-to-end connectivity through Cisco MDS switches; single initiator zoning configured
- FC SAN boot to fully achieve stateless computing; servers are booted from LUNs in the boot volume that is hosted on the AFF storage cluster
- All MEDITECH workloads are hosted on FC LUNs, which are spread across the storage controller nodes

Host software

The following software must be installed:

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

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

Cisco UCS blade server and switch configuration

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

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

Cisco Nexus switch configuration

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

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

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

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

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

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

Cisco MDS 9132T configuration

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

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

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

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

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

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

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

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

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

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

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

```

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

```

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

```

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

```

Cisco UCS configuration guidance

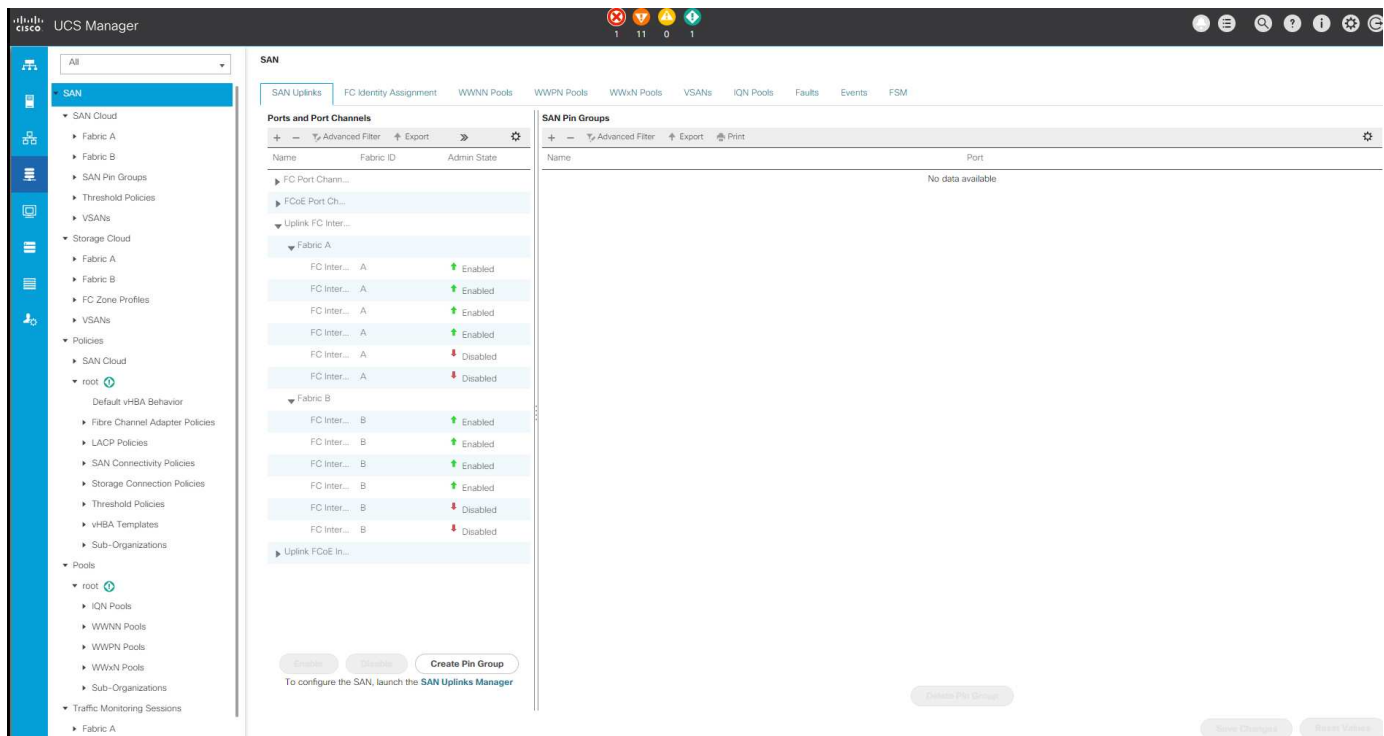
Cisco UCS enables you as a MEDITECH customer to leverage your subject-matter experts in network, storage, and compute to create policies and templates that tailor the environment to your specific needs. After

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

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

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

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



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

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

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

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

and default maintenance.

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

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

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

ESXi configuration best practices

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

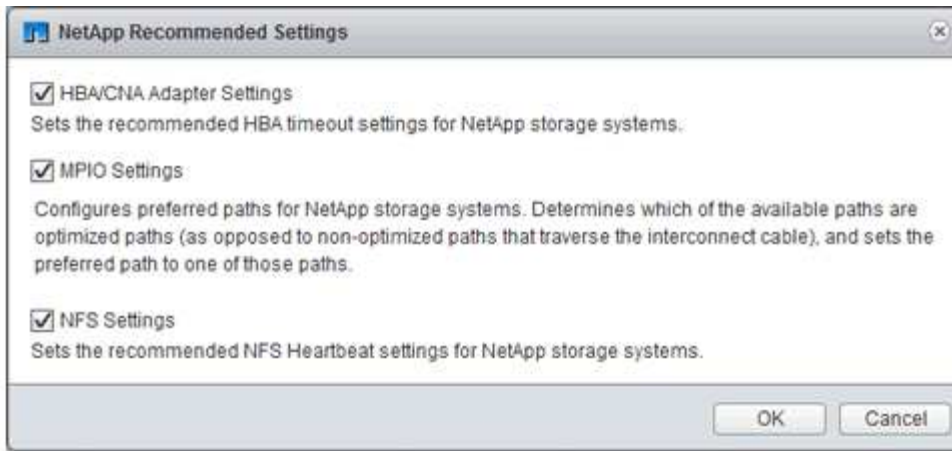
- VSC for VMware vSphere checks and sets the ESXi host multipathing settings and HBA timeout settings that work best with NetApp storage systems. The values that VSC sets are based on rigorous internal testing by NetApp.
- For optimal storage performance, consider using storage hardware that supports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

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

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

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

The standard recommended values are set by default.



4. Click OK.

[Next: NetApp Configuration.](#)

NetApp configuration

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

ONTAP configuration

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

ONTAP licenses

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

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

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

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

AutoSupport configuration

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

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

Hardware-assisted takeover configuration

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

1. Run the following ONTAP command to xxx.

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

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

2. Run the following ONTAP command to xxx:

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

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

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

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

[Next: Aggregate Configuration.](#)

Aggregate configuration

NetApp RAID DP

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

RAID group size and number of RAID groups

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

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

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

Storage expansion considerations

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

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

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



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

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

Aggregate-level Snapshot copies

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

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

Storage virtual machine configuration

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



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

SVM for MEDITECH host LUNs

You should create one dedicated SVM per ONTAP storage cluster to own and to manage the aggregates that contain the LUNs for the MEDITECH hosts.

SVM language encoding setting

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

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

[Next: Volume Configuration.](#)

Volume configuration

Volume provisioning

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

Default volume-level Snapshot copies

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

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

Number of volumes per aggregate

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

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

Volume-level reallocate schedule

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

Best practice

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



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

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

Next: [LUN Configuration](#).

LUN configuration

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

LUN provisioning

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

LUN operating system type

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

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

LUN size

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

LUN presentation

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

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

Refer to the MEDITECH Hardware Configuration Proposal when you consult with the MEDITECH and backup system installer to devise a consistent naming convention for the LUNs that the MEDITECH hosts use.

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

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

[Next: Initiator Group Configuration.](#)

Initiator group configuration

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

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

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



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

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

[Next: LUN Mappings.](#)

LUN mappings

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

MEDITECH modules and components

The MEDITECH application covers several modules and components. The following table lists the functions that are covered by these modules. For additional information about setting up and deploying these modules, see the MEDITECH documentation.

Function	Type
Connectivity	<ul style="list-style-type: none"> • Web server • Live application server (WI – Web Integration) • 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 – Client Access License)
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)

Function	Type
Third-party requirement	<ul style="list-style-type: none"> • First Databank (FDB) MedKnowledge Framework v4.3

Acknowledgments

The following people contributed to the creation of this guide.

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

- [FlexPod Design Zone](#)
- [FlexPod Data Center 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](#)

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—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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