



## **FlexPod for Medical Imaging**

**FlexPod**

NetApp

October 30, 2025

This PDF was generated from [https://docs.netapp.com/us-en/flexpod/healthcare/medical-imaging\\_introduction.html](https://docs.netapp.com/us-en/flexpod/healthcare/medical-imaging_introduction.html) on October 30, 2025. Always check [docs.netapp.com](https://docs.netapp.com) for the latest.

# Table of Contents

FlexPod for Medical Imaging .....	1
TR-4865: FlexPod for Medical Imaging .....	1
Overall solution benefits .....	3
Scope .....	5
Audience .....	5
Medical imaging application .....	5
Size of the healthcare organization and platform sizing .....	6
FlexPod .....	8
Architecture .....	13
Storage architecture .....	13
Networking .....	17
Compute—Cisco Unified Computing System .....	18
Virtualization .....	18
Medical imaging system architecture .....	18
Solution infrastructure hardware and software components .....	23
Solution sizing .....	25
Storage sizing .....	25
Compute sizing .....	25
Networking and Cisco UCS infrastructure sizing .....	26
Best practices .....	27
Storage best practices .....	27
Backup best practices .....	30
Infrastructure best practices .....	30
Compute best practices .....	30
Virtualization best practices .....	31
Medical imaging system best practices .....	31
Conclusion .....	31
Additional information .....	32

# FlexPod for Medical Imaging

## TR-4865: FlexPod for Medical Imaging

Jaya Kishore Esanakula and Atul Bhalodia, NetApp

Medical imaging accounts for 70% of all data that is generated by Healthcare organizations. As digital modalities continue to advance and new modalities emerge, the amount of data will continue to increase. For example, the transition from analog to digital pathology will dramatically increase image sizes at a rate that will challenge any data management strategies currently in place.

COVID-19 has clearly reshaped the digital transformation; according to a recent [report](#), COVID-19 has accelerated digital commerce by 5 years. The technological innovation driven by problem solvers is fundamentally changing the way that we go about our daily life. This technology-driven change will overhaul many critical aspects of our life, including healthcare.

Healthcare is poised to undergo a major change in the coming years. COVID is accelerating innovation in healthcare that will propel the industry by at least several years. At the heart of this change is the need to make healthcare more flexible in handling pandemics by being more affordable, available, and accessible, without compromising reliability.

At the foundation of this healthcare change is a well-designed platform. One of the key metrics to measure the platform is the ease with which platform changes can be implemented. Speed is the new scale and data protection cannot be compromised. Some of the world's most critical data is being created and consumed by the clinical systems that support clinicians. NetApp has made critical data available for patient care where the clinicians need it, on premise, in the cloud, or in a hybrid setting. Hybrid multi- cloud environments are the current state of the art for IT architecture.

Healthcare as we know it revolves around providers (doctors, nurses, radiologists, medical device technicians, and so on) and patients. As we bring patients and providers closer together, making the geographic location a mere data point, it becomes even more important for the underlying platform to be available when providers and patients need it. The platform must be both efficient and cost-effective in the long term. In their efforts to drive patient care costs even lower, [Accountable Care Organizations](#) (ACOs) would be empowered by an efficient platform.

When it comes to health information systems used by healthcare organizations, the question of build versus purchase tends to have a single answer: purchase. This could be for many subjective reasons. Purchasing decisions made over many years can create heterogeneous information systems. Each system has a specific set of requirements for the platform that they are deployed on. The most significant issue is the large, diverse set of storage protocols and performance levels that information systems require, which makes platform standardization and optimal operational efficiency a significant challenge. Healthcare organizations cannot focus on mission critical issues because their attention is spread thin by trivial operational needs like the large set of platforms that require a diversified set of skills and thus SME retention.

The challenges can be classified into the following categories:

- Heterogeneous storage needs
- Departmental silos
- IT operational complexity
- Cloud connectivity

- Cybersecurity
- Artificial intelligence and deep learning

With FlexPod, you get a single platform that supports FC, FCoE, iSCSI, NFS/pNFS, SMB/CIFS and so on from a single platform. People, processes, and technology are part of the DNA that FlexPod is designed and built upon. FlexPod adaptive QoS helps to break down the departmental silos by supporting multiple mission critical clinical systems on the same underlying FlexPod platform. FlexPod is FedRAMP certified and FIPS 140-2 certified. Additionally, healthcare organizations are faced with opportunities such as artificial intelligence and deep learning. FlexPod and NetApp solve these challenges and make the data available where it is needed on premises or in a hybrid multi- cloud setting in a standardized platform. For more information and a series customer success stories, see [FlexPod Healthcare](#).

Typical medical imaging information and PACS systems have the following set of capabilities:

- Reception and registration
- Scheduling
- Imaging
- Transcription
- Management
- Data exchange
- Image archive
- Image viewing for image capturing and reading for technicians and image viewing for clinicians

Regarding imaging, the healthcare sector is trying to solve the following clinical challenges:

- Wider adoption of [natural language processing](#) (NLP)-based assistants by technicians and physicians for image reading. Radiology department can benefit from voice recognition to transcribe reports. NLP can be used to identify and anonymize a patient's record, specifically DICOM tags embedded in the DICOM image. NLP capabilities require high performing platforms with low latency response times for image processing. FlexPod QoS not only delivers and performance but also provides mature capacity projections for future growth.
- Wider adoption of standardized clinical pathways and protocols by ACOs and community health organizations. Historically, clinical pathways have been used as a static set of guidelines rather than an integrated workflow that guides clinical decisions. With advancements in NLP and image processing, DICOM tags in images can be integrated into clinical pathways as facts to drive clinical decisions. Therefore, these processes require high performance, low latency, and high throughput from the underlying infrastructure platform and storage systems.
- ML models that leverage convolutional neural networks enable automation of image- processing capabilities in real time and thus require infrastructure that is GPU-capable. FlexPod offers both CPU and GPU compute components built into the same system, and CPUs and GPUs can be scaled independently of each other.
- If DICOM tags are used as facts in clinical best-practice advisories, then the system must perform more reads of DICOM artifacts with low latency and high throughput.
- When evaluating images, real-time collaboration between radiologists across organizations requires high performance graphics processing in the end- user compute devices. NetApp provides industry- leading VDI solutions specifically designed and proven for high-end graphics use cases. More information can be found [here](#).
- Image and media management across ACO health organizations can uses a single platform, regardless of the system of record for the image, by using protocols such as Digital Imaging and Communications in

Medicine ( [DICOM](#)) and web access to DICOM-persistent objects ( [WADO](#))

- Health information exchange ( [HIE](#)) includes images embedded in messages.
- Mobile modalities, such as handheld, wireless scanning devices (for example, pocket handheld ultrasound scanners attached to a phone), require a robust network infrastructure with DoD-level security, reliability, and latency at the edge, the core, and in the cloud. [A data fabric enabled by NetApp](#) provide organizations with this capability at scale.
- Newer modalities have exponential storage needs; for example, CT and MRI require a few hundred MBs for each modality, but digital pathology images (including whole slide imaging) can be a few GBs in size. FlexPod is designed with [performance, reliability and scaling as foundational traits](#).

A well-architected medical imaging system platform is at the heart of innovation. The FlexPod architecture provides flexible compute and storage capabilities with industry-leading storage efficiency.

## Overall solution benefits

By running an imaging application environment on a FlexPod architectural foundation, your healthcare organization can expect to see an improvement in staff productivity and a decrease in capital and operating expenses. FlexPod provides a rigorously tested, prevalidated, converged that is engineered and designed to deliver predictable low-latency system performance and high availability. This approach results in high comfort levels and, ultimately, optimal response times for users of the medical imaging system.

Different components of the imaging system might require the storage of data in SMB/CIFS, NFS, Ext4, or NTFS file systems. That requirement means that the infrastructure must provide data access over the NFS, SMB/CIFS, and SAN protocols. A single NetApp storage system can support the NFS, SMB/CIFS, and SAN protocols, thus eliminating the need for the legacy practice of protocol- specific storage systems.

The FlexPod infrastructure is a modular, converged, virtualized, scalable (scale-out and scale- up), and cost-effective platform. With the FlexPod platform, you can independently scale out compute, network, and storage to accelerate your application deployment. And the modular architecture enables nondisruptive operations even during system scale-out and upgrade activities.


FlexPod delivers several benefits that are specific to the medical imaging industry:

- **Low-latency system performance.** Radiologist time is a high- value resource, and efficient use of a radiologist's time is paramount. Waiting for images or videos to load can contribute to clinician burnout and can affect clinician's efficiency and patient safety.
- **Modular architecture.** FlexPod components are connected through a clustered server, a storage management fabric, and a cohesive management toolset. As imaging facilities grow year over year and the number of studies increase, there will be a need for the underlying infrastructure to scale accordingly. FlexPod can scale compute, storage, and network independently.
- **Quicker deployment of infrastructure.** Whether it is in an existing data center or a remote location, the integrated and tested design of FlexPod Datacenter with Medical Imaging enables you to get the new infrastructure up and running in less time, with less effort.
- **Accelerated application deployment.** A prevalidated architecture reduces implementation integration time and risk for any workload, and NetApp technology automates infrastructure deployment. Whether you use the solution for an initial rollout of medical imaging, a hardware refresh, or expansion, you can shift more resources to the business value of the project.
- **Simplified operations and lower costs.** You can eliminate the expense and complexity of legacy proprietary platforms by replacing them with a more efficient and scalable shared resource that can meet the dynamic needs of your workload. This solution delivers higher infrastructure resource utilization for greater return on investment (ROI).

- **Scale-out architecture.** You can scale SAN and NAS from terabytes to tens of petabytes without reconfiguring running applications.
- **Nondisruptive operations.** You can perform storage maintenance, hardware lifecycle operations, and software upgrades without interrupting your business.
- **Secure multitenancy.** This benefit supports the increased needs of virtualized server and storage shared infrastructure, enabling secure multitenancy of facility-specific information, particularly if you are hosting multiple instances of databases and software.
- **Pooled resource optimization.** This benefit can help you 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. These industry-leading QoS storage policies enable differentiated service levels in a shared environment. These policies help optimize performance for workloads and help to isolate and control runaway applications.
- **Support for storage tier SLAs by using QoS.** You don't have to deploy different storage systems for the different storage tiers that a medical imaging environment typically requires. A single storage cluster with multiple NetApp FlexVol volumes with specific QoS policies for different tiers can serve that purpose. With this approach, storage infrastructure can be shared by dynamically accommodating the changing needs of a particular storage tier. NetApp AFF can support different SLAs for storage tiers by allowing QoS at the level of the FlexVol volume, thus eliminating the need for different storage systems for different storage tiers for the application.
- **Storage efficiency.** Medical images are typically pre-compressed by the imaging application to jpeg2k lossless compression which is around 2.5:1. However, this is imaging application and vendor specific. In larger imaging application environments (greater than 1PB), 5-10% storage savings are possible, and you can reduce storage costs with NetApp storage efficiency features. Work with your imaging application vendors and your NetApp subject matter expert to unlock potential storage efficiencies for your medical imaging system.
- **Agility.** With the industry-leading workflow automation, orchestration, and management tools that FlexPod systems offer, your IT team can be far more responsive to business requests. These business requests can range from medical imaging backup and provisioning of additional test and training environments to analytics database replications for population health- management initiatives.
- **Higher productivity.** You can quickly deploy and scale this solution for optimal clinician end-user experiences.
- **Data fabric.** Your data fabric powered by NetApp weaves data together across sites, beyond physical boundaries, and across applications. Your data fabric powered by NetApp 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. So, you want a consistent and integrated way to manage it. This solution provides a way to manage data that puts your IT team in control and that simplifies ever-increasing IT complexity.
- **FabricPool.** NetApp ONTAP FabricPool helps reduce storage costs without compromising performance, efficiency, security, or protection. FabricPool is transparent to enterprise applications and capitalizes on cloud efficiencies by lowering storage TCO without the need to rearchitect the application infrastructure. FlexPod can benefit from the storage tiering capabilities of FabricPool to make more efficient use of ONTAP flash storage. For full information, see [FlexPod with FabricPool](#).
- **FlexPod security.** Security is at the very foundation of FlexPod. In the past few years, ransomware has become a significant and increasing threat. Ransomware is malware that is based on crypto virology, the use of cryptography to build malicious software. This malware can use both symmetric and asymmetric key encryption to lock a victim's data and demand a ransom to provide the key to decrypt the data. To learn how FlexPod helps mitigate threats like ransomware, see [The Solution to Ransomware](#). FlexPod infrastructure components are also Federal Information Processing Standard ([FIPS](#)) [140-2](#) compliant.
- **FlexPod Cooperative Support.** NetApp and Cisco have established FlexPod 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. The FlexPod Cooperative Support model helps confirm 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 if your healthcare organization runs business-critical applications. The illustration below shows an overview of the FlexPod Cooperative Support model.



## Scope

This document provides a technical overview of a Cisco Unified Computing System (Cisco UCS) and NetApp ONTAP-based FlexPod infrastructure for hosting this medical imaging solution.

## Audience

This document is intended for technical leaders in the healthcare industry and for Cisco and NetApp partner solutions engineers and professional services personnel. NetApp assumes that the reader has a good understanding of compute and storage sizing concepts as well as technical familiarity with the medical imaging system, Cisco UCS, and NetApp storage systems.

## Medical imaging application

A typical medical imaging application offers a suite of applications that together make an enterprise-grade

imaging solution for small, medium, and large healthcare organizations.

At the heart of the product suite are the following clinical capabilities:

- Enterprise imaging repository
- Supports traditional image sources such as radiology and cardiology. Also supports other care areas like ophthalmology, dermatology, colonoscopy, and other medical imaging objects like photos and videos.
- [Picture archiving and communication system \(PACS\)](#), which is a computerized means of replacing the roles of conventional radiological film
- Enterprise Imaging Vendor Neutral Archive (VNA):
  - Scalable consolidation of DICOM and non-DICOM documents
  - Centralized Medical Imaging system
  - Support for document synchronization and data integrity between multiple (PACSs) in the enterprise
  - Document lifecycle management by a rules-based expert system that leverages document metadata, such as:
    - Modality type
    - Age of study
    - Patient age (current and at the time of image capture)
    - Single point of integration within and outside (HIE) of the enterprise:
      - Context- aware document linking
      - Health Level Seven International (HL7), DICOM, and WADO
      - Storage- agnostic archiving capability
- Integration with other health information systems that use HL7 and context-aware linking:
  - Enables EHRs to implement direct links to patient images from patient charts, imaging workflows, and so on.
  - Helps embed a patient's longitudinal care image history into EHRs.
- Radiology technologist workflows
- Enterprise zero footprint viewers for image viewing from anywhere on any capable device
- Analytical tools that leverage retrospective and real-time data:
  - Compliance reporting
  - Operational reports
  - Quality control and quality assurance reports

## Size of the healthcare organization and platform sizing

Healthcare organizations can be broadly classified by using standards-based methods that help programs such as ACO. One such classification uses the concept of a clinical integrated network (CIN). A group of hospitals can be called a CIN if they collaborate and adhere to proven standard clinical protocols and pathways to improve the value of care and reduce patient costs. Hospitals within a CIN have controls and practices in place to onboard physicians who follow the core values of the CIN. Traditionally, an integrated delivery networks (IDN) has been limited to hospitals and physician groups. A CIN crosses traditional IDN boundaries, and a CIN can still be part of an ACO. Following the principles of a CIN, healthcare organizations can be classified into small, medium, and large.

## **Small healthcare organizations**

A healthcare organization is small if it includes only a single hospital with ambulatory clinics and an inpatient department, but it is not part of a CIN. Physicians work as caregivers and coordinate patient care during a care continuum. These small organizations typically include physician-operated facilities. They might or might not offer emergency and trauma care as integrated care for the patient. Typically, a small-sized healthcare organization performs about 250,000 clinical imaging studies annually. Imaging centers are considered to be small healthcare organizations and they do provide imaging services. Some also provide radiology dictation services to other organizations.

## **Medium healthcare organizations**

A healthcare organization considered to be of medium size if it includes multiple hospital systems with focused organizations, such as the following:

- Adult care clinics and adult inpatient hospitals
- Labor and delivery departments
- Childcare clinics and child inpatient hospitals
- A cancer treatment center
- Adult emergency departments
- Child emergency departments
- A family medicine and primary care office
- An adult trauma care center
- A child trauma care center

In a medium-sized healthcare organization, physicians follow the principles of a CIN and operate as a single unit. Hospitals have separate hospital, physician, and pharmacy billing functions. Hospitals might be associated with academic research institutes and perform interventional clinical research and trials. A medium healthcare organization performs as many as 500,000 clinical imaging studies annually.

## **Large healthcare organizations**

A healthcare organization is considered to be large if it includes the traits of a medium- sized healthcare organization and offers the medium-sized clinical capabilities to the community in multiple geographical locations.

A large healthcare organization typically performs the following functions:

- Has a central office to manage the overall functions
- Participates in joint ventures with other hospitals
- Negotiates rates with payer organizations annually
- Negotiates payer rates by state and region
- Participates in Meaningful Use (MU) programs
- Performs advanced clinical research across population health cohorts by using standards-based population health management (PHM) tools
- Performs up to one million clinical imaging studies annually

Some large healthcare organizations that participate in a CIN also have AI- based imaging reading capabilities.

These organizations typically perform one to two million clinical imaging studies annually.

Before you look into how these different-sized organizations translate into an optimally sized FlexPod system, you should understand the various FlexPod components and the different capabilities of a FlexPod system.

## FlexPod

### Cisco Unified Computing System

Cisco UCS consists of a single management domain that is interconnected with a unified I/O infrastructure. Cisco UCS for medical imaging environments has been aligned with NetApp medical imaging system infrastructure recommendations and best practices so that the infrastructure can deliver critical patient information with maximum availability.

The compute foundation of enterprise medical imaging is Cisco UCS technology, with its integrated systems management, Intel Xeon processors, and server virtualization. These integrated technologies solve data center challenges and enable you to meet your goals for data center design with a typical medical imaging system. Cisco UCS unifies LAN, SAN, and systems management into one simplified link for rack servers, blade servers, and virtual machines (VMs). Cisco UCS consists of a redundant pair of Cisco UCS fabric interconnects that provide a single point of management and a single point of control for all I/O traffic.

Cisco UCS uses service profiles so that virtual servers in the Cisco UCS infrastructure are configured correctly and consistently. 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 single, simple operation that enables migration of identities between servers in the environment without requiring any physical configuration changes. It also facilitates rapid bare-metal provisioning of replacements for failed servers.

The use of service profiles helps confirm 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 must be performed in one domain, the virtual infrastructure can be migrated to another domain. With this approach, even when a single domain is offline, applications continue to run with high availability.

Cisco UCS is a next-generation solution for blade and rack server computing. The system integrates a low-latency, lossless, 40GbE unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain. Cisco UCS accelerates the delivery of new services simply, reliably, and securely through end-to-end provisioning and migration support for both virtualized and nonvirtualized systems. Cisco UCS provides the following features:

- Comprehensive management
- Radical simplification
- High performance

Cisco UCS consists of the following components:

- **Compute.** The system is based on an entirely new class of computing system that incorporates rack-mounted and blade servers based on the Intel Xeon scalable processor product family.
- **Network.** The system is integrated into a low-latency, lossless, 40Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing networks, which are

separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables and also by decreasing power and cooling requirements.

- **Virtualization.** The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.
- **Storage access.** The system provides consolidated access to both SAN storage and NAS over the unified fabric. It is also an ideal system for software-defined storage. By combining the benefits of a single framework to manage both the compute and the storage servers in a single pane, QoS can be implemented if needed to inject I/O throttling in the system. And your server administrators can preassign storage-access policies to storage resources, which simplifies storage connectivity and management and can help increase productivity. In addition to external storage, both rack and blade servers have internal storage that can be accessed through built-in hardware RAID controllers. By setting up the storage profile and disk configuration policy in Cisco UCS Manager, the storage needs of the host OS and application data are fulfilled by user-defined RAID groups. The result is high availability and better performance.
- **Management.** The system uniquely integrates all system components so that the entire solution can be managed as a single entity by Cisco UCS Manager. To manage all system configuration and operations, Cisco UCS Manager has an intuitive GUI, a CLI, and a powerful scripting library module for Microsoft Windows PowerShell that are built on a robust API.

Cisco Unified Computing System fuses access layer networking and servers. This high-performance, next-generation server system gives your data center a high degree of workload agility and scalability.

## Cisco UCS Manager

Cisco UCS Manager provides unified, embedded management for all software and hardware components in Cisco UCS. By using single-connection technology, UCS Manager manages, controls, and administers multiple chassis for thousands of VMs. Through an intuitive GUI, a CLI, or an XML API, your administrators use the software to manage the entire Cisco UCS as a single logical entity. Cisco UCS Manager resides on a pair of Cisco UCS 6300 Series Fabric Interconnects that use clustered, active-standby configuration for high availability.

Cisco UCS Manager offers a unified embedded management interface that integrates your servers, network, and storage. Cisco UCS Manager performs auto discovery to detect the inventory of, to manage, and to provision system components that you add or change. It offers a comprehensive set of XML APIs for third-party integration, and it exposes 9,000 points of integration. It also facilitates custom development for automation, for orchestration, and to achieve new levels of system visibility and control.

Service profiles benefit both virtualized and nonvirtualized environments. They increase the mobility of nonvirtualized servers, such as when you move workloads from server to server or when you take a server offline for service or upgrade. You can also use profiles in conjunction with virtualization clusters to bring new resources online easily, complementing existing VM mobility.

For more information about Cisco UCS Manager, see the [Cisco UCS Manager product page](#).

## Cisco UCS differentiators

Cisco Unified Computing System is revolutionizing the way that servers are managed in the data center. See the following unique differentiators of Cisco UCS and Cisco UCS Manager:

- **Embedded management.** In Cisco UCS, the servers are managed by the embedded firmware in the fabric interconnects, eliminating the need for any external physical or virtual devices to manage them.
- **Unified fabric.** In Cisco UCS, from blade server chassis or rack servers to fabric interconnects, a single

Ethernet cable is used for LAN, SAN, and management traffic. This converged I/O reduces the number of cables, SFPs, and adapters that you need, in turn reducing your capital and operational expenses for the overall solution.

- **Autodiscovery.** By simply inserting the blade server in the chassis or by connecting rack servers to the fabric interconnects, discovery and inventory of compute resource occurs automatically without any management intervention. The combination of unified fabric and auto discovery enables the wire-once architecture of Cisco UCS, where its compute capability can be extended easily while keeping the existing external connectivity to LAN, SAN, and management networks.
- **Policy-based resource classification.** When a compute resource is discovered by Cisco UCS Manager, it can be automatically classified to a given resource pool based on the policies that you defined. This capability is useful in multitenant cloud computing.
- **Combined rack and blade server management.** Cisco UCS Manager can manage B-Series blade servers and C-Series rack servers under the same Cisco UCS domain. This feature, along with stateless computing, makes compute resources truly hardware form factor-agnostic.
- **Model-based management architecture.** The Cisco UCS Manager architecture and management database are model-based and data-driven. The open XML API that is provided to operate on the management model enables easy and scalable integration of Cisco UCS Manager with other management systems.
- **Policies, pools, and templates.** The management approach in Cisco UCS Manager is based on defining policies, pools, and templates instead of a cluttered configuration. It enables a simple, loosely coupled, data-driven approach in managing compute, network, and storage resources.
- **Loose referential integrity.** In Cisco UCS Manager, a service profile, a port profile, or policies can refer to other policies or to other logical resources with loose referential integrity. A referred policy cannot exist at the time of authoring the referring policy, but a referred policy can be deleted even though other policies are referring to it. This feature enables different subject-matter experts to work independently from each other. You gain great flexibility by enabling different experts from different domains—such as network, storage, security, server, and virtualization—to work together to accomplish a complex task.
- **Policy resolution.** In Cisco UCS Manager, you can create a tree structure of organizational unit hierarchy that mimics the real-life tenants and organizational relationships. You can define various policies, pools, and templates at different levels of your organizational hierarchy. A policy that refers to another policy by name is resolved in the organizational hierarchy with the closest policy match. If no policy with a specific name is found in the hierarchy of the root organization, then a special policy named “default” is searched. This policy resolution practice enables automation-friendly management APIs and provides great flexibility to the owners of the different organizations.
- **Service profiles and stateless computing.** A service profile is a logical representation of a server, carrying its various identities and policies. You can assign this logical server to any physical compute resource, as long as it meets the resource requirements. Stateless computing enables procurement of a server within minutes, which used to take days in legacy server management systems.
- **Built-in multitenancy support.** The combination of policies, pools, templates, a loose referential integrity, policy resolution in organizational hierarchy, and a service profiles- based approach to compute resources makes Cisco UCS Manager inherently friendly to multitenant environments that are typically observed in private and public clouds.
- **Extended memory.** The enterprise-class Cisco UCS B200 M5 Blade Server extends the capabilities of the Cisco Unified Computing System portfolio in a half-width blade form factor. The Cisco UCS B200 M5 harnesses the power of the latest Intel Xeon scalable- processor CPUs with up to 3TB of RAM. This feature enables the huge VM-to-physical- server ratio that many deployments need or enables certain architectures to support large memory operations, such as big data.
- **Virtualization- aware network.** Cisco Virtual Machine Fabric Extender (VM-FEX) technology makes the access network layer aware of host virtualization. This awareness prevents pollution of compute and

network domains with virtualization when a virtual network is managed by port profiles that are defined by your network administrator team. VM-FEX also offloads hypervisor CPU by performing switching in the hardware, thus enabling the hypervisor CPU to perform more virtualization-related tasks. To simplify cloud management, VM-FEX technology is well integrated with VMware vCenter, Linux Kernel-Based Virtual Machine (KVM), and Microsoft Hyper-V SR-IOV.

- **Simplified QoS.** Even though FC and Ethernet are converged in the Cisco UCS, built-in support for QoS and lossless Ethernet make it seamless. By representing all system classes in one GUI panel, network QoS is simplified in Cisco UCS Manager.

## Cisco Nexus IP and MDS switches

Cisco Nexus switches and Cisco MDS multilayer directors give you enterprise-class connectivity and SAN consolidation. Cisco multiprotocol storage networking helps reduce your business risk by providing flexibility and options: FC, Fiber Connection (FICON), FC over Ethernet (FCoE), 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 the data center and the 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 compute resources with Cisco Nexus switches and a unified 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 the following capabilities:

- **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, Cisco NX-OS software features, and management to support zero-downtime environments.
- **Transport flexibility.** You can incrementally adopt new networking technologies with this cost-effective solution.

Together, Cisco UCS with Cisco Nexus switches and MDS multilayer directors provide a compute, networking, and SAN connectivity solution for an enterprise medical Imaging system.

## NetApp all-flash storage

NetApp storage that runs ONTAP software reduces your overall storage costs while delivering the low-latency read and write response times and high IOPS that medical imaging system workloads need. To create an optimal storage system that meets a typical medical imaging system requirement, ONTAP supports both all-flash and hybrid storage configurations. NetApp flash storage gives medical imaging system customers like you the key components of high performance and responsiveness to support latency-sensitive medical imaging system operations. By creating multiple fault domains in a single cluster, NetApp technology can also isolate your production environments from your nonproduction environments. And by guaranteeing that system performance do not drop below a certain level for workloads with ONTAP minimum QoS, NetApp reduces performance issues for your system.

The scale-out architecture of ONTAP software can flexibly adapt to your various I/O workloads. To deliver the necessary throughput and low latency that clinical applications need and to provide 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. You can clone, replicate, and back up your medical imaging system environment from expensive SSD storage to more economical HDD storage on other nodes. With NetApp cloud-enabled storage and a data fabric delivered by NetApp, you can back up to object storage on premises

or in the cloud.

For medical imaging, ONTAP has been validated by most leading medical imaging systems. That means it has been tested to deliver fast and reliable performance for medical imaging. Additionally, the following features simplify management, increase availability and automation, and reduce the total amount of storage that you need.

- **Outstanding performance.** The NetApp AFF solution shares the same unified storage architecture, ONTAP software, management interface, rich data services, and advanced feature set as the rest of the NetApp FAS product families. This innovative combination of all-flash media with ONTAP gives you the consistent low latency and high IOPS of all-flash storage with industry-leading ONTAP software.
- **Storage efficiency.** You can reduce your total capacity requirements work with your NetApp SME to understand how this applied your specific medical imaging system.
- **Space-efficient cloning.** With the FlexClone capability, your system can almost instantly create clones to support backup and testing environment refresh. These clones consume additional storage only as changes are made.
- **Integrated data protection.** Full data protection and disaster recovery features help you protect your critical data assets and provide disaster recovery.
- **Nondisruptive operations.** You can perform upgrades and maintenance without taking data offline.
- **QoS.** Storage QoS helps you limit potential bully workloads. More importantly, QoS creates a minimum performance guarantee that your system performance will not drop below a certain level for critical workloads such as a medical imaging system's production environment. And by limiting contention, NetApp QoS can also reduce performance-related issues.
- **Data fabric.** To accelerate digital transformation, your data fabric delivered by NetApp simplifies and integrates data management across cloud and on-premises environments. It delivers consistent and integrated data management services and applications for superior data visibility and insights, data access and control, and data protection and security. NetApp is integrated with large public clouds, such AWS, Azure, Google Cloud, and IBM Cloud, giving you a wide breadth of choice.

## Host virtualization — VMware vSphere

FlexPod architectures are validated with VMware vSphere 6.x, which is the industry-leading virtualization platform. VMware ESXi 6.x is used to deploy and run the VMs. vCenter Server Appliance 6.x is used to manage the ESXi hosts and VMs. Multiple ESXi hosts that run on Cisco UCS B200 M5 blades are used to form a VMware ESXi cluster. The VMware ESXi cluster pools the compute, memory, and network resources from all the cluster nodes and provides a resilient platform for the VMs that are running on the cluster. The VMware ESXi cluster features, vSphere high availability, and Distributed Resource Scheduler (DRS) all contribute to the vSphere cluster's tolerance to withstand failures, and they help distribute the resources across the VMware ESXi hosts.

The NetApp storage plug-in and the Cisco UCS plug-in integrate with VMware vCenter to enable operational workflows for your required storage and compute resources.

The VMware ESXi cluster and vCenter Server give you a centralized platform for deploying medical imaging environments in VMs. Your healthcare organization can realize all the benefits of an industry-leading virtual infrastructure with confidence, such as the following:

- **Simple deployment.** Quickly and easily deploy vCenter Server by using a virtual appliance.
- **Centralized control and visibility.** Administer the entire vSphere infrastructure from a single location.
- **Proactive optimization.** Allocate, optimize, and migrate resources for maximum efficiency.


- **Management.** Use powerful plug-ins and tools to simplify management and to extend control.

## Architecture

The FlexPod architecture is designed to provide high availability if a component or a link fails in your entire compute, network, and storage stack. Multiple network paths for client access and storage access provide load balancing and optimal resource utilization.

The following figure illustrates the 16Gb FC/40Gb Ethernet (40GbE) topology for the medical imaging system solution deployment.

### FlexPod Infrastructure for an Enterprise Medical Imaging System



## Storage architecture

Use the storage architecture guidelines in this section to configure your storage infrastructure for an enterprise medical imaging system.

### Storage tiers

A typical enterprise medical imaging environment consists of several different storage tiers. Each tier has specific performance and storage protocol requirements. NetApp storage supports various RAID technologies; more information can be found [here](#). Here is how NetApp AFF storage systems serve the needs of different

storage tiers for the imaging system:

- **Performance Storage (tier 1).** This tier offers high performance and high redundancy for databases, OS drives, VMware Virtual Machine File System (VMFS) datastores, and so on. Block I/O moves over fiber to a shared storage array of SSD, as is configured in ONTAP. The minimum latency is 1ms to 3ms, with an occasional peak of 5ms. This storage tier is typically used for short-term storage cache, typically 6 to 12 months of image storage for fast access to online DICOM images. This tier offers high performance and high redundancy for image caches, database backup, and so on. NetApp all-flash arrays provide <1ms latency at a sustained bandwidth, which is far lower than the service times that are expected by a typical enterprise medical imaging environment. NetApp ONTAP supports both RAID-TEC (triple parity RAID to sustain three disk failures) and RAID DP (double-parity RAID to sustain two disk failures).
- **Archive storage (tier 2).** This tier is used for typical cost-optimized file access, RAID 5 or RAID 6 storage for larger volumes, and long-term lower-cost/performance archiving. NetApp ONTAP supports both RAID-TEC (triple parity RAID to sustain three disk failures) and RAID DP (double-parity RAID to sustain two disk failures). NetApp FAS in FlexPod enables imaging application I/O over NFS/SMB to a SAS disk array. NetApp FAS systems provide ~10ms latency at sustained bandwidth, which is far lower than the service times that are expected for storage tier 2 in an enterprise medical imaging system environment.

Cloud-based archiving in a hybrid-cloud environment can be used for archiving to a public cloud storage provider using S3 or similar protocols. NetApp SnapMirror technology enables replication of imaging data from all-flash or FAS arrays to slower disk-based storage arrays or to Cloud Volumes ONTAP for AWS, Azure, or Google Cloud.

NetApp SnapMirror provides industry leading data replication capabilities that help protect your medical imaging system with unified data replication. Simplify data-protection management across the data fabric with cross-platform replication—from flash to disk to cloud:

- Transport data seamlessly and efficiently between NetApp storage systems to support both backup and disaster recovery with the same target volume and I/O stream.
- Failover to any secondary volume. Recover from any point-in-time Snapshot on the secondary storage.
- Safeguard your most critical workloads with available zero-data-loss synchronous replication (RPO=0).
- Cut network traffic. Shrink your storage footprint through efficient operations.
- Reduce network traffic by transporting only changed data blocks.
- Preserve storage-efficiency benefits on the primary storage during transport—including deduplication, compression, and compaction.
- Deliver additional inline efficiencies with network compression.

More information can be found [here](#).

The table below lists each tier that a typical medical imaging system requires for specific latency and the throughput performance characteristics.

Storage tier	Requirements	NetApp recommendation
1	1–5ms latency 35–500MBps throughput	AFF with <1ms latency AFF A300 high-availability (HA) pair with two disk shelves can handle throughput of up to ~1.6GBps

Storage tier	Requirements	NetApp recommendation
2	On premises archive	FAS with up to 30ms latency
	Archive to cloud	SnapMirror replication to Cloud Volumes ONTAP or backup archiving with NetApp StorageGRID software

## Storage network connectivity

### FC fabric

- The FC fabric is for host OS I/O from compute to storage.
- Two FC fabrics (Fabric A and Fabric B) are connected to Cisco UCS Fabric A and UCS Fabric B, respectively.
- A storage virtual machine (SVM) with two FC logical interfaces (LIFs) is on each controller node. On each node, one LIF is connected to Fabric A and the other is connected to Fabric B.
- 16Gbps FC end-to-end connectivity is through Cisco MDS switches. A single initiator, multiple target ports, and zoning are all configured.
- FC SAN boot is used to create fully stateless computing. Servers are booted from LUNs in the boot volume that is hosted on the AFF storage cluster.

### IP network for storage access over iSCSI, NFS, and SMB/CIFS

- Two iSCSI LIFs are in the SVM on each controller node. On each node, one LIF is connected to Fabric A and the second is connected to Fabric B.
- Two NAS data LIFs are in the SVM on each controller node. On each node, one LIF is connected to Fabric A and the second is connected to Fabric B.
- Storage port interface groups (virtual port channel [vPC]) for 10Gbps link to switch N9k-A and for 10Gbps link to switch N9k-B.
- Workload in Ext4 or NTFS file systems from VM to storage:
  - iSCSI protocol over IP.
- VMs hosted in NFS datastore:
  - VM OS I/O goes over multiple Ethernet paths through Nexus switches.

### In-band management (active-passive bond)

- 1Gbps link to management switch N9k-A, and 1Gbps link to management switch N9k-B.

## Backup and recovery

FlexPod Datacenter is built on a storage array that is managed by NetApp ONTAP data management software. ONTAP software has evolved over 20 years to provide many data management features for VMs, Oracle databases, SMB/CIFS file shares, and NFS. It also provides protection technology such as NetApp Snapshot technology, SnapMirror technology, and NetApp FlexClone data replication technology. NetApp SnapCenter software has a server and a GUI client to use ONTAP Snapshot, SnapRestore, and FlexClone features for VM, SMB/CIFS file shares, NFS, and Oracle database backup and recovery.

NetApp SnapCenter software employs [patented](#) Snapshot technology to create a backup of an entire VM or

Oracle database on a NetApp storage volume instantaneously. Compared with Oracle Recovery Manager (RMAN), Snapshot copies do not require a full baseline backup copy, because they are not stored as physical copies of blocks. Snapshot copies are stored as pointers to the storage blocks as they existed in the ONTAP WAFL file system when the Snapshot copies were created. Because of this tight physical relationship, the Snapshot copies are maintained on the same storage array as the original data. Snapshot copies can also be created at the file level to give you more granular control for the backup.

Snapshot technology is based on a redirect-on-write technique. It initially contains only metadata pointers and does not consume much space until the first data change to a storage block. If an existing block is locked by a Snapshot copy, a new block is written by the ONTAP WAFL file system as an active copy. This approach avoids the double-writes that occur with the change-on-write technique.

For Oracle database backup, Snapshot copies yield incredible time savings. For example, a backup that took 26 hours to complete by using RMAN alone can take less than 2 minutes to complete by using SnapCenter software.


And because data restoration does not copy any data blocks but instead flips the pointers to the application-consistent Snapshot block images when the Snapshot copy was created, a Snapshot backup copy can be restored almost instantaneously. SnapCenter cloning creates a separate copy of metadata pointers to an existing Snapshot copy and mounts the new copy to a target host. This process is also fast and storage efficient.

The following table summarizes the primary differences between Oracle RMAN and NetApp SnapCenter software.


	<b>Backup</b>	<b>Restore</b>	<b>Clone</b>	<b>Need Full Backup</b>	<b>Space usage</b>	<b>Off-site copy</b>
RMAN	Slow	Slow	Slow	Yes	High	Yes
SnapCenter	Fast	Fast	Fast	No	Low	Yes

The following figure presents the SnapCenter architecture.

## VMware Integration



## NetApp® SnapCenter® Servers



NetApp MetroCluster configurations are used by thousands of enterprises worldwide for high availability (HA), zero data loss, and nondisruptive operations both within and beyond the data center. MetroCluster is a free feature of ONTAP software that synchronously mirrors data and configuration between two ONTAP clusters in separate locations or failure domains. MetroCluster provides continuously available storage for applications by automatically handling two objectives: Zero recovery point objective (RPO) by synchronously mirroring data written to the cluster. Near zero recovery time objective (RTO) by mirroring configuration and automating access to data at the second site. MetroCluster provides simplicity with automatic mirroring of data and configuration between the two independent clusters located in the two sites. As storage is provisioned within one cluster, it is automatically mirrored to the second cluster at the second site. NetApp SyncMirror technology provides a complete copy of all data with a zero RPO. Therefore, workloads from one site can switch over at any time to the opposite site and continue serving data without data loss. More information can be found [here](#).

## Networking

A pair of Cisco Nexus switches provides redundant paths for IP traffic from compute to storage, and for external clients of the medical imaging system image viewer:

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

- An FC storage network is provided by a pair of Cisco MDS switches.

## Compute—Cisco Unified Computing System

Two Cisco UCS fabrics through different fabric interconnects provide two failure domains. Each fabric is connected to both IP networking switches and to different FC networking switches.

Identical service profiles for each Cisco UCS blade are created as per FlexPod best practices to run VMware ESXi. Each service profile should have the following components:

- Two vNICs (one on each fabric) to carry NFS, SMB/CIFS, and client or management traffic
- Additional required VLANs to the vNICs for NFS, SMB/CIFS, and client or management traffic
- Two vNICs (one on each fabric) to carry iSCSI traffic
- Two storage FC HBAs (one on each fabric) for FC traffic to storage
- SAN boot

## Virtualization

The VMware ESXi host cluster runs workload VMs. The cluster comprises ESXi instances running on Cisco UCS blade servers.

Each ESXi host includes the following network components:

- SAN boot over FC or iSCSI
- Boot LUNs on NetApp storage (in a dedicated FlexVol for boot OS)
- Two VMNICs (Cisco UCS vNIC) for NFS, SMB/CIFS, or management traffic
- Two storage HBAs (Cisco UCS FC vHBA) for FC traffic to storage
- Standard switch or distributed virtual switch (as needed)
- NFS datastore for workload VMs
- Management, client traffic network, and storage network port groups for VMs
- Network adapter for management, client traffic, and storage access (NFS, iSCSI, or SMB/CIFS) for each VM
- VMware DRS enabled
- Native multipathing enabled for FC or iSCSI paths to storage
- VMware snapshots for VM turned off
- NetApp SnapCenter deployed for VMware for VM backups


## Medical imaging system architecture

In healthcare organizations, medical imaging systems are critical applications and well-integrated into the clinical workflows that begin from patient registration and end with billing related activities in the revenue cycle.

The following diagram shows the various systems involved in a typical large hospital; this diagram is intended to provide architectural context to a medical imaging system before we zoom into the architectural components of a typical medical imaging system. Workflows vary widely and are hospital and use- case specific.

The figure below shows the medical imaging system in the context of a patient, a community clinic, and a large

hospital.



1. The patient visits the community clinic with symptoms. During the consultation, the community physician places an imaging order that is sent to the larger hospital in the form of a HL7 order message.
2. The community physician's EHR system sends the HL7 order/ORD message to the large hospital.
3. The enterprise interoperability system (also known as the Enterprise Service Bus [ESB]) processes the order message and sends the order message to the EHR system.
4. The EHR processes the order message. If a patient record does not exist, a new patient record is created.
5. The EHR sends an imaging order to the medical imaging system.
6. The patient calls the large hospital for an imaging appointment.
7. The imaging reception and registration desk schedules patient for an imaging appointment using a radiology information or similar system.
8. The patient arrives for the imaging appointment, and the images or video is created and sent to the PACS.
9. The radiologist reads the images and annotates the images in the PACS using a high-end/GPU graphics-enabled diagnostic viewer. Certain imaging systems have artificial intelligence (AI)- enabled efficiency improvement capabilities built into the imaging workflows.
10. Image order results are sent to the EHR in the form of an order results HL7 ORU message via the ESB.
11. The EHR processes the order results into the patient's record, places thumbnail image with a context-aware link to the actual DICOM image. Physicians can launch the diagnostic viewer if a higher resolution image is needed from within the EHR.
12. The physician reviews the image and enters physician notes into the patient's record. The physician could use the clinical decision support system to enhance the review process and aid in proper diagnosis for the patient.
13. The EHR system then sends the order results in the form of an order results message to the community hospital. At this point, if the community hospital could receive the complete image, then the image is sent either via WADO or DICOM.
14. The community physician completes the diagnosis and provides next steps to the patient.

A typical medical imaging system uses an N- tiered architecture. The core component of a medical imaging system is an application server to host various application components. Typical application servers are either Java runtime- based or C# .Net CLR- based. Most enterprise medical imaging solutions use an Oracle

database Server or MS SQL Server or Sybase as the primary database. Additionally, some enterprise medical imaging systems also use databases for content acceleration and caching over a geographic region. Some enterprise medical imaging systems also use NoSQL databases like MongoDB, Redis, and so on in conjunction with enterprise integration servers for DICOM interfaces and or APIs.

A typical medical imaging system provides access to images for two distinct set of users: diagnostic user/radiologist, or the clinician or physician that ordered the imaging.

Radiologists typically use high- end, graphics- enabled diagnostic viewers that are running on high- end compute and graphics workstations that are either physical or part of a virtual desktop infrastructure. If you are about to start your virtual desktop infrastructure journey, more information can be found [here](#).

When hurricane Katrina destroyed two of Louisiana's major teaching hospitals, leaders came together and built a resilient electronic health record system that included over 3000 virtual desktops in record time. More information on use cases reference architecture and FlexPod reference bundles can be found [here](#).

Clinicians access images in two primary ways:


- **Web- based access.** which is typically used by EHR systems to embed PACS images as context- aware links into the electronic medical record (EMR) of the patient, and links that can be placed into imaging workflows, procedure workflows, progress notes workflows, and so on. Web based links are also used to provide image access to the patients via patient portals. Web based access uses a technology pattern called context aware links. Context aware links can either be static links/URIs to the DICOM media directly or dynamically generated links/URIs using custom macros.
- **Thick client.** Some enterprise medical systems also allow you to use a thick- client- based approach to view the images. You can launch a thick client from within the EMR of the patient or as a standalone application.

The medical imaging system can provide image access to a community of physicians or to CIN-participating physicians. Typical medical imaging systems include components that enable image interoperability with other health IT systems within and outside of your healthcare organization. Community physicians can either access images via a web-based application or leverage an image exchange platform for image interoperability. Image- exchange platforms typically use either WADO or DICOM as the underlying image exchange protocol.


Medical imaging systems can also support academic medical centers that need PACS or imaging systems for use in a classroom. To support academic activities, a typical medical imaging system can have the capabilities of a PACS system in a smaller footprint or a teaching- only imaging environment. Typical vendor- neutral archiving systems and some enterprise- class medical imaging systems offer DICOM image tag morphing capabilities to anonymize the images that are used for teaching purposes. Tag morphing enables healthcare organization to exchange DICOM images between different vendor medical imaging systems in a vendor- neutral fashion. Also, tag morphing enables medical imaging systems to implement an enterprise- wide, vendor- neutral archival capability for medical images.

Medical imaging systems are starting to use [GPU-based compute capabilities](#) to enhance human workflows by preprocessing the images and thus improving efficiencies. Typical enterprise medical imaging systems take advantage of industry- leading NetApp storage efficiency capabilities. Enterprise medical imaging systems typically use RMAN for backup, recovery, and restore activities. For better performance and to reduce the time that it takes to create backups, Snapshot technology is available for backup operations and SnapMirror technology is available for replication.

The figure below shows the logical application components in a layered architectural view.



The figure below shows the physical application components.



The logical application components require that the infrastructure support a diverse set of protocols and file systems. NetApp ONTAP software supports an industry-leading set of protocols and file systems.

The table below lists the application components, storage protocol, and file system requirements.

Application component	SAN/NAS	File system type	Storage tier	Replication type
VMware host prod DB	local	SAN	VMFS	Tier 1
Application	VMware host prod DB	REP	SAN	VMFS
Tier 1	Application	VMware host prod application	local	SAN
VMFS	Tier 1	Application	VMware host prod application	REP

Application component	SAN/NAS	File system type	Storage tier	Replication type
SAN	VMFS	Tier 1	Application	Core database server
SAN	Ext4	Tier 1	Application	Backup database server
SAN	Ext4	Tier 1	None	Image cache server
NAS	SMB/CIFS	Tier 1	None	Archive server
NAS	SMB/CIFS	Tier 2	Application	Web server
NAS	SMB/CIFS	Tier 1	None	WADO Server
SAN	NFS	Tier 1	Application	Business intelligence server
SAN	NTFS	Tier 1	Application	Business intelligence backup
SAN	NTFS	Tier 1	Application	Interoperability server
SAN	Ext4	Tier 1	Application	Interoperability database server

## Solution infrastructure hardware and software components

The following tables list the hardware and software components, respectively, of the FlexPod infrastructure for the medical imaging system.

Layer	Product family	Quantity and model	Details
Compute	Cisco UCS 5108 chassis	1 or 2	Based on the number of blades required to support the number of annual studies
	Cisco UCS blade servers	B200 M5	Number of blades based on the number of studies annually Each with 2 x 20 or more cores, 2.7GHz, and 128-384GB RAM
	Cisco UCS Virtual Interface Card (VIC)	Cisco UCS 1440	See the
	2 x Cisco UCS fabric interconnects	6454 or later	—
Network	Cisco Nexus switches	2 x Cisco Nexus 3000 Series or 9000 Series	—

Layer	Product family	Quantity and model	Details
Storage network	IP network for storage access over SMB/CIFS, NFS, or iSCSI protocols	Same network switches as above	–
	Storage access over FC	2 x Cisco MDS 9132T	–
Storage	NetApp AFF A400 all-flash storage system	1 or more HA pair	Cluster with two or more nodes
	Disk shelf	1 or more DS224C or NS224 disk shelves	Fully populated with 24 drives
	SSD	>24, 1.2TB or larger capacity	–

Software	Product family	Version or release	Details
Enterprise medical imaging system	MS SQL or Oracle Database Server	As suggested by the medical imaging system vendor	
	No SQL DBs like MongoDB Server	As suggested by the medical imaging system vendor	
	Application Servers	As suggested by the medical imaging system vendor	
	Integration Server (MS Biztalk, MuleSoft, Rhapsody, Tibco)	As suggested by the medical imaging system vendor	
	VMs	Linux (64 bit)	
	VMs	Windows Server (64 bit)	
Storage	ONTAP	ONTAP 9.7 or later	
Network	Cisco UCS Fabric Interconnect	Cisco UCS Manager 4.1 or later	
	Cisco Ethernet switches	9.2(3)I7(2) or later	
	Cisco FC: Cisco MDS 9132T	8.4(2) or later	
Hypervisor	Hypervisor	VMware vSphere ESXi 6.7 U2 or later	
Management	Hypervisor management system	VMware vCenter Server 6.7 U1 (vCSA) or later	
	NetApp Virtual Storage Console (VSC)	VSC 9.7 or later	
	SnapCenter	SnapCenter 4.3 or later	

# Solution sizing

## Storage sizing

This section describes the number of studies and the corresponding infrastructure requirements.

The storage requirements that are listed in the following table assume that existing data is 1 year's worth plus projected growth for 1 year of study in the primary system (tier 1, 2). Additional storage needs for projected growth for 3 years beyond the first 2 years are listed separately.

	<b>Small</b>	<b>Medium</b>	<b>Large</b>
Annual studies	<250K studies	250K–500K studies	500K–1 million studies
<b>Tier 1 Storage</b>			
IOPS (average)	1.5K–5K	5K–15K	15K–40K
IOPS (peak)	5K	20K	65K
Throughput	50–100MBps	50–150MBps	100–300MBps
Capacity data center 1 (1 year of old data and 1 year of new study)	70TB	140TB	260TB
Capacity data center 1 (additional need for 4 years for new study)	25TB	45TB	80TB
Capacity data center 2 (1 year of old data and 1 year of new study)	45TB	110TB	165TB
Capacity data center 2 (additional need for 4 years for new study)	25TB	45TB	80TB
<b>Tier 2 Storage</b>			
IOPS (average)	1K	2K	3K
Capacity data center 1	320TB	800TB	2000TB

## Compute sizing

The table below lists the compute requirements for small, medium, and large medical imaging systems.

	<b>Small</b>	<b>Medium</b>	<b>Large</b>
Annual studies	<250K studies	250K–500K studies	500K–1 million studies
<b>Data Center 1</b>			
Number of VMs	21	27	35
Total virtual CPU (vCPU) count	56	124	220

	<b>Small</b>	<b>Medium</b>	<b>Large</b>
Total memory requirement	225GB	450GB	900GB
Physical server (blades) specs (assume 1 vCPU = 1 core)	4 x servers with 20 cores and 192GB RAM each	8 x servers with 20 cores and 128GB RAM each	14 x servers with 20 cores and 128GB RAM each
<b>Data Center 2</b>			
Number of VMs	15	17	22
Total vCPU count	42	72	140
Total memory requirement	179GB	243GB	513GB
Physical server (blades) specs (assume 1 vCPU = 1 core)	3 x servers with 20 cores and 168GB RAM each	6 x servers with 20 cores and 128GB RAM each	8 x servers with 24 cores and 128GB RAM each

## **Networking and Cisco UCS infrastructure sizing**

The table below lists the networking and Cisco UCS infrastructure requirements for small, medium, and large medical imaging systems.

	<b>Small</b>	<b>Medium</b>	<b>Large</b>
<b>Data Center 1</b>			
Number of storage node ports	2 converged network adapters (CNAs); 2 FCs	2 CNAs; 2 FCs	2 CNAs; 2 FCs
IP network switch ports (Cisco Nexus 9000)	48-port switch	48-port switch	48-port switch
FC switch (Cisco MDS)	32-port switch	32-port switch	48-port switch
Cisco UCS chassis count	1 x 5108	1 x 5108	2 x 5108
Cisco UCS Fabric Interconnect	2 x 6332	2 x 6332	2 x 6332
<b>Data Center 2</b>			
Cisco UCS chassis count	1 x 5108	1 x 5108	1 x 5108
Cisco UCS Fabric Interconnect	2 x 6332	2 x 6332	2 x 6332
Number of storage node ports	2 CNAs; 2 FCs	2 CNAs; 2 FCs	2 CNAs; 2 FCs
IP network switch ports (Cisco Nexus 9000)	48-port switch	48-port switch	48-port switch
FC switch (Cisco MDS)	32-port switch	32-port switch	48-port switch

# Best practices

## Storage best practices

### High availability

The NetApp storage cluster design provides high availability at every level:

- Cluster nodes
- Back-end storage connectivity
- RAID TEC that can sustain three disk failures
- RAID DP that can sustain two disk failures
- Physical connectivity to two physical networks from each node
- Multiple data paths to storage LUNs and volumes

### Secure multitenancy

NetApp storage virtual machines (SVMs) provide a virtual storage array construct to separate your security domain, policies, and virtual networking. NetApp recommends that you create separate SVMs for each tenant organization that hosts data on the storage cluster.

### NetApp storage best practices


Consider the following NetApp storage best practices:

- Always enable NetApp AutoSupport technology, which sends support summary information to NetApp through HTTPS.
- For maximum availability and mobility, make sure that a LIF is created for each SVM on each node in the NetApp ONTAP cluster. Asymmetric logical unit access (ALUA) is used to parse paths and to identify active optimized (direct) paths versus active nonoptimized paths. ALUA is used for both FC or FCoE and iSCSI.
- A volume that contains only LUNs does not need to be internally mounted, nor is a junction path required.
- If you use the Challenge-Handshake Authentication Protocol (CHAP) in ESXi for target authentication, you must also configure it in ONTAP. Use the CLI (`vserver iscsi security create`) or NetApp ONTAP System Manager (edit Initiator Security under Storage > SVMs > SVM Settings > Protocols > iSCSI).

### SAN boot

NetApp recommends that you implement SAN boot for Cisco UCS Servers in the FlexPod Datacenter solution. This step enables the operating system to be safely secured by the NetApp AFF storage system, providing better performance. The design that is outlined in this solution uses iSCSI SAN boot.

In iSCSI SAN boot, each Cisco UCS Server is assigned two iSCSI vNICs (one for each SAN fabric), which provide redundant connectivity all the way to the storage. The storage ports in this example, e2a and e2e, which are connected to the Cisco Nexus switches, are grouped together to form one logical port called an interface group (ifgrp) (in this example, a0a). The iSCSI VLANs are created on the igroup, and the iSCSI LIFs are created on iSCSI port groups (in this example, a0a-<iSCSI-A-VLAN>). The iSCSI boot LUN is exposed to the servers through the iSCSI LIF by using igroups. This approach enables only the authorized server to have access to the boot LUN. For the port and LIF layout, see the figure below.



Unlike NAS network interfaces, the SAN network interfaces are not configured to fail over during a failure. Instead, if a network interface becomes unavailable, the host chooses a new optimized path to an available network interface. ALUA, a standard supported by NetApp, provides information about SCSI targets, which enables a host to identify the best path to the storage.

### Storage efficiency and thin provisioning

NetApp has led the industry in storage efficiency innovation, such as with the first deduplication for primary workloads and with inline data compaction, which enhances compression and stores small files and I/Os efficiently. ONTAP supports both inline and background deduplication, as well as inline and background compression.

To realize the benefits of deduplication in a block environment, the LUNs must be thin-provisioned. Although the LUN is still seen by your VM administrator as taking the provisioned capacity, the deduplication savings are returned to the volume to be used for other needs. NetApp recommends that you deploy these LUNs in FlexVol volumes that are also thin-provisioned with a capacity that is two times the size of the LUN. When you deploy the LUN that way, the FlexVol volume acts merely as a quota. The storage that the LUN consumes is reported in the FlexVol volume and its containing aggregate.

For maximum deduplication savings, consider scheduling background deduplication. These processes use system resources when they're running, however. So, ideally, you should schedule them during less active times (such as weekends) or run them more frequently to reduce the amount of changed data to be processed. Automatic background deduplication on AFF systems has much less of an effect on foreground activities. Background compression (for hard disk-based systems) also consumes resources, so you should consider it only for secondary workloads with limited performance requirements.

### Quality of service

Systems that run ONTAP software can use the ONTAP storage QoS feature to limit throughput in megabits per second (MBps) and to limit IOPS for different storage objects such as files, LUNs, volumes, or entire SVMs. Adaptive QoS is used to set an IOPS floor (QoS minimum) and ceiling (QoS maximum), which dynamically adjust based on the datastore capacity and used space.

Throughput limits are useful for controlling unknown or test workloads before a deployment to confirm that they don't affect other workloads. You might also use these limits to constrain a bully workload after it has been identified. Minimum levels of service based on IOPS are also supported to provide consistent performance for SAN objects in ONTAP.

With an NFS datastore, a QoS policy can be applied to the entire FlexVol volume or to individual Virtual Machine Disk (VMDK) files within it. With VMFS datastores (Cluster Shared Volumes [CSV] in Hyper-V) that use ONTAP LUNs, you can apply the QoS policies to the FlexVol volume that contains the LUNs or to the individual LUNs. However, because ONTAP has no awareness of the VMFS, you cannot apply the QoS policies to individual VMDK files. When you use VMware Virtual Volumes (VVols) with VSC 7.1 or later, you can set maximum QoS on individual VMs by using the storage capability profile.

To assign a QoS policy to a LUN, including VMFS or CSV, you can obtain the ONTAP SVM (displayed as vserver), LUN path, and serial number from the Storage Systems menu on the VSC home page. Select the storage system (SVM), then Related Objects > SAN. Use this approach when you specify QoS by using one of the ONTAP tools.

You can set the QoS maximum throughput limit on an object in MBps and in IOPS. If you use both, the first limit that is reached is enforced by ONTAP. A workload can contain multiple objects, and a QoS policy can be applied to one or more workloads. When you apply a policy to multiple workloads, the workloads share the total limit of the policy. Nested objects are not supported (for example, for a file within a volume, they cannot each have their own policy). QoS minimums can be set only in IOPS.

## Storage layout

This section provides best practices for layout of LUNs, volumes, and aggregates on storage.

### Storage LUNs

For optimal performance, management, and backup, NetApp recommends the following LUN-design best practices:

- Create a separate LUN to store database data and log files.
- Create a separate LUN for each instance to store Oracle database log backups. The LUNs can be part of the same volume.
- Provision LUNs with thin provisioning (disable the Space Reservation option) for database files and log files.
- All imaging data is hosted in FC LUNs. Create these LUNs in FlexVol volumes that are spread across the aggregates that are owned by different storage controller nodes.

For placement of the LUNs in a storage volume, follow the guidelines in the next section.

### Storage volumes

For optimal performance and management NetApp recommends the following volume design best practices:

- Isolate databases with I/O-intensive queries on separate storage volumes.
- The datafiles could be placed on a single LUN or a volume, but multiple volumes/LUNs are recommended for higher throughput.
- I/O parallelism can be attained by using any supported filesystem when multiple LUNs are used.
- Place database files and transaction logs on separate volumes to increase the recovery granularity.
- Consider using volume attributes like auto size, Snapshot reserve, QoS, and so on.

### Aggregates

Aggregates are the primary storage containers for NetApp storage configurations and contain one or more RAID groups that consist of both data disks and parity disks.

NetApp performed various I/O workload characterization tests by using shared and dedicated aggregates with data files and transaction log files separated. The tests show that one large aggregate with more RAID groups and drives (HDDs or SSDs) optimizes and improves storage performance and is easier for administrators to manage for two reasons:

- One large aggregate makes the I/O abilities of all drives available to all files.

- One large aggregate enables the most efficient use of disk space.

For effective disaster recovery, NetApp recommends that you place the asynchronous replica on an aggregate that is part of a separate storage cluster in your disaster recovery site and use SnapMirror technology to replicate content.

For optimal storage performance, NetApp recommends that you have at least 10% free space available in an aggregate.

Storage aggregate layout guidance for AFF A300 systems (with two disk shelves with 24 drives) includes:

- Keep two spare drives.
- Use Advanced Disk Partitioning to create three partitions on each drive: root and data.
- Use a total of 20 data partitions and two parity partitions for each aggregate.

## Backup best practices

NetApp SnapCenter is used for VM and database backups. NetApp recommends the following backup best practices:

- When SnapCenter is deployed to create Snapshot copies for backups, turn off the Snapshot schedule for the FlexVol that host VMs and application data.
- Create a dedicated FlexVol for host boot LUNs.
- Use a similar or a single backup policy for VMs that serve the same purpose.
- Use a similar or a single backup policy per workload type; for example, use a similar policy for all database workloads. Use different policies for databases, web servers, end-user virtual desktops, and so on.
- Enable verification of the backup in SnapCenter.
- Configure archiving of the backup Snapshot copies to the NetApp SnapVault backup solution.
- Configure retention of the backups on primary storage based on the archiving schedule.

## Infrastructure best practices

### Networking best practices

NetApp recommends the following networking best practices:

- Make sure that your system includes redundant physical NICs for production and storage traffic.
- Separate VLANs for iSCSI, NFS, and SMB/CIFS traffic between compute and storage.
- Make sure that your system includes a dedicated VLAN for client access to the medical imaging system.

You can find additional networking best practices in the FlexPod infrastructure design and deployment guides.

## Compute best practices

NetApp recommends the following compute best practice:

- Make sure that each specified vCPU is supported by a physical core.

## Virtualization best practices

NetApp recommends the following virtualization best practices:

- Use VMware vSphere 6 or later.
- Set the ESXi host server BIOS and OS layer to Custom Controlled–High Performance.
- Create backups during off-peak hours.

## Medical imaging system best practices

See the following best practices and some requirements from a typical medical imaging system:

- Do not overcommit virtual memory.
- Make sure that the total number of vCPUs equals the number of physical CPUs.
- If you have a large environment, dedicated VLANs are required.
- Configure database VMs with dedicated HA clusters.
- Make sure that the VM OS VMDKs are hosted in fast tier 1 storage.
- Work with the medical imaging system vendor to identify the best approach to prepare VM templates for quick deployment and maintenance.
- Management, storage, and production networks require LAN segregation for the database, with isolated VLANs for VMware vMotion.
- Use the NetApp storage-array-based replication technology called [SnapMirror](#) instead of vSphere-based replication.
- Use backup technologies that leverage VMware APIs; backup windows should be outside the normal production hours.

## Conclusion

By running a medical imaging environment on FlexPod, your healthcare organization can expect to see an improvement in staff productivity and a decrease in capital and operating expenses. FlexPod provides a prevalidated, rigorously tested converged infrastructure from the strategic partnership of Cisco and NetApp. It is engineered and designed specifically to deliver predictable low-latency system performance and high availability. This approach results in a superior user experience and optimal response time for users of the medical imaging system.

Different components of a medical imaging system require data storage in SMB/CIFS, NFS, Ext4, and NTFS file systems. Therefore, your infrastructure must provide data access over NFS, SMB/CIFS, and SAN protocols. NetApp storage systems support these protocols from a single storage array.

High availability, storage efficiency, Snapshot copy-based scheduled fast backups, fast restore operations, data replication for disaster recovery, and the FlexPod storage infrastructure capabilities all provide an industry-leading data storage and management system.

# Additional information

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

- FlexPod Datacenter for AI/ML with Cisco UCS 480 ML for Deep Learning Design Guide

[https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/flexpod\\_c480m5l\\_aiml\\_design.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_c480m5l_aiml_design.html)

- FlexPod Datacenter Infrastructure with VMware vSphere 6.7 U1, Cisco UCS 4th Generation, and NetApp AFF A-Series

[https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/flexpod\\_datacenter\\_vmware\\_ne\\_tappaffa.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_datacenter_vmware_ne_tappaffa.html)

- FlexPod Datacenter Oracle Database Backup with SnapCenter Solution Brief

<https://www.netapp.com/us/media/sb-3999.pdf>

- FlexPod Datacenter with Oracle RAC Databases on Cisco UCS and NetApp AFF A-Series

[https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/flexpod\\_orc12cr2\\_affaseries.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_orc12cr2_affaseries.html)

- FlexPod Datacenter with Oracle RAC on Oracle Linux

[https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/flexpod\\_orcrac\\_12c\\_bm.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_orcrac_12c_bm.html)

- FlexPod for Microsoft SQL Server

<https://flexpod.com/solutions/use-cases/microsoft-sql-server/>

- FlexPod from Cisco and NetApp

<https://flexpod.com/>

- NetApp Solutions for MongoDB Solution Brief (NetApp login required)

<https://fieldportal.netapp.com/content/734702>

- TR-4700: SnapCenter Plug-In for Oracle Database

<https://www.netapp.com/us/media/tr-4700.pdf>

- NetApp Product Documentation

<https://www.netapp.com/us/documentation/index.aspx>

- FlexPod for Virtual Desktop Infrastructure (VDI) Solutions

<https://flexpod.com/solutions/use-cases/virtual-desktop-infrastructure/>

## 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.