



Healthcare

FlexPod

NetApp
November 09, 2021

Table of Contents

- Healthcare 1
 - FlexPod for Genomics 1
 - FlexPod Datacenter for Epic Directional Sizing Guide 41
 - FlexPod Datacenter for Epic EHR Deployment Guide 56
 - FlexPod for Epic Performance Testing 93
 - FlexPod for MEDITECH Directional Sizing Guide 110
 - FlexPod Datacenter for MEDITECH Deployment Guide 121
 - FlexPod for Medical Imaging 151

Healthcare

FlexPod for Genomics

TR-4911: FlexPod Genomics

JayaKishore Esanakula, NetApp

There are few fields of medicine that are more important than genomics for healthcare and the life sciences, and genomics are fast becoming a key clinical tool for doctors and nurses. Genomics, when combined with medical imaging and digital pathology, help us understand how a patient's genes might be affected by treatment protocols. The success of genomics in healthcare increasingly depends on data interoperability at scale. The end goal is to make sense of the enormous volumes of genetic data and identify clinically relevant correlations and variants that improve diagnosis and make precision medicine a reality. Genomics help us understand the origin of disease outbreaks, how diseases evolve, and which treatments and strategies might be effective. Clearly, genomics has many benefits that span prevention, diagnosis, and treatment. Healthcare organizations are grappling with several challenges, including the following:

- Improved care quality
- Value-based care
- Data explosion
- Precision medicine
- Pandemics
- Wearables, remote monitoring, and care
- Cyber security

Standardized clinical pathways and clinical protocols are one of the critical components of modern medicine. One of the key aspects of standardization is interoperability between care providers, not just for medical records but also for genomic data. The big question is will healthcare organizations relinquish ownership of genomic data in lieu of patient ownership of their personal genomics data and related medical records?

Interoperable patient data is key for enabling precision medicine, one of the driving forces behind the recent explosion of data growth. The objective for precision medicine is to make health maintenance, disease prevention, diagnoses, and treatment solutions more effective and precise.

The rate of data growth has been exponential. In early February 2021, US laboratories sequenced approximately 8,000 COVID-19 strains per week. The number of genomes sequenced had increased to 29,000 per week by April 2021. Each fully sequenced human genome is around 125GB in size. Therefore, at a rate of 29,000 genomes sequenced per week, total genome storage at rest would be more than 180 petabytes per year. Various countries have committed resources to genomic epidemiology to improve genomic surveillance and prepare for the next wave of global health challenges.

The reduced cost of genomic research is driving genetic testing and research at an unprecedented rate. The three Ps are at an inflection point: computer power, privacy of data, and personalization of medicine. By 2025 researchers estimate that 100 million to as many as 2 billion human genomes will be sequenced. For genomics to be effective and a valuable proposition, genomics capabilities must be a seamless part of care workflows; it should be easy to access and be actionable during a patient's visit. It is also equally important that patient electronic medical-record data be integrated with patient genomics data. With the advent of state-of-the-art converged infrastructure like FlexPod, organizations can bring their genomics capabilities into the everyday workflows of physicians, nurses, and clinic managers. For the latest FlexPod platform information, see this

For a physician, the true value of genomics includes precision medicine and personalized treatment plans based on the genomic data of a patient. There has never been such synergy between clinicians and data scientists in the past, and genomics is benefiting from the technological innovations in the recent past, and also real partnerships between healthcare organizations and technology leaders in the industry.

Academic medical centers and other healthcare and life science organizations are well on their way to establishing center of excellence (COE) in genome science. According to Dr. Charlie Gersbach, Dr. Greg Crawford, and Dr. Tim E Reddy from Duke University, “We know that genes aren’t turned on or off by a simple binary switch, but instead it’s a result of multiple gene regulatory switches that work together.” They have also determined that “none of these parts of the genome work in isolation. The genome is a very complicated web that evolution has woven” ([ref](#)).

NetApp and Cisco have been hard at work implementing incremental improvements into the FlexPod platform for over 10 years. All customer feedback is heard, evaluated, and tied into the value streams and feature sets in FlexPod. It is this continuous loop of feedback, collaboration, improvement, and celebration that sets FlexPod apart as a trusted converged infrastructure platform the world over. It has been simplified and designed from the ground up to be the most reliable, robust, versatile, and agile platform for healthcare organizations.

Scope

The FlexPod converged infrastructure platform enables a healthcare organization to host one or more genomics workloads, along with other clinical and nonclinical healthcare applications. This technical report uses an open-source, industry-standard genomics tool called GATK during FlexPod platform validation. However, a deeper discussion of genomics or GATK is outside the scope of this document.

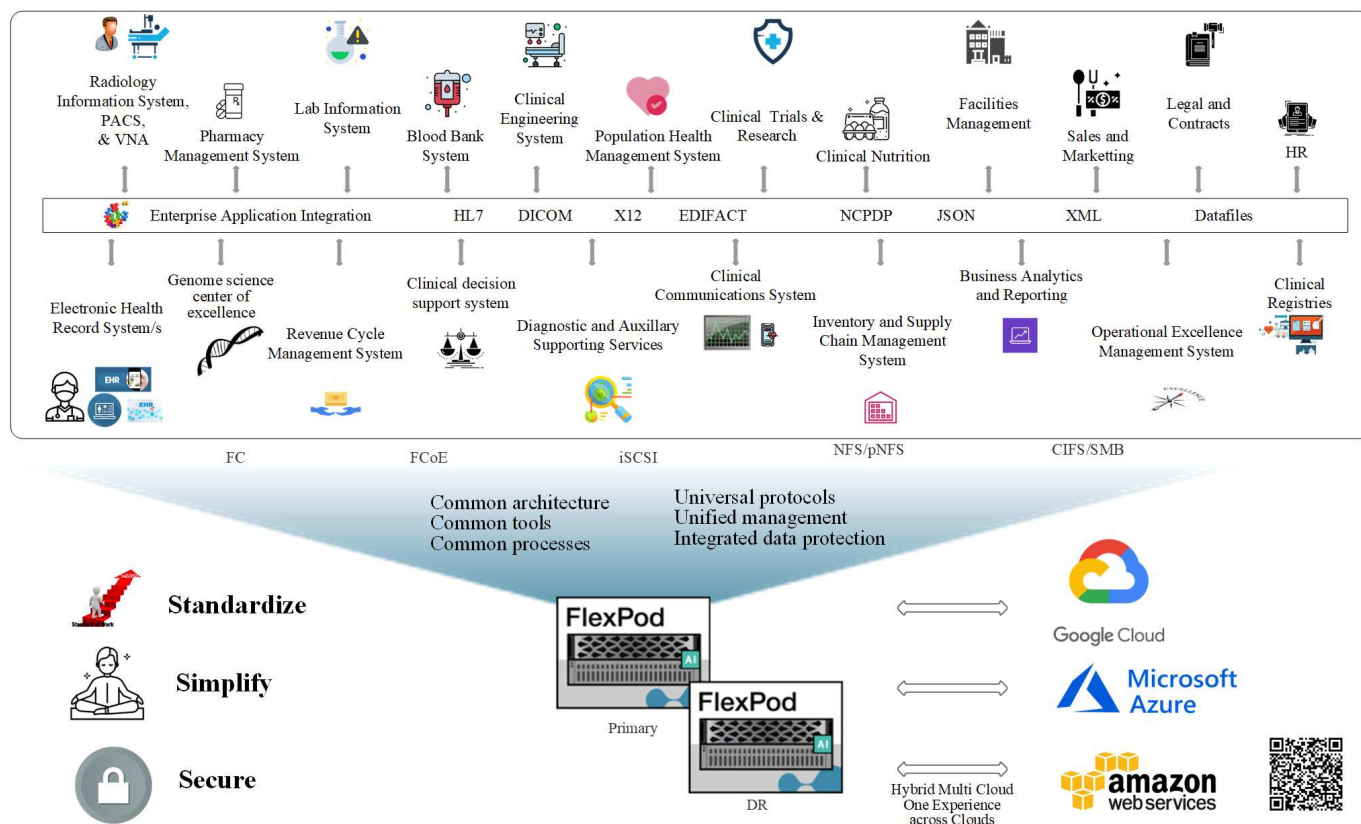
Audience

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 a technical familiarity with healthcare threats, healthcare security, healthcare IT systems, Cisco UCS, and NetApp storage systems.

Hospital capabilities deployed on FlexPod

A typical hospital has a diversified set of IT systems. The majority of such systems are purchased from a vendor, whereas very few are built by the hospital system in house. Therefore, the hospital system must manage a diverse infrastructure environment in their data centers. When hospitals unify their systems into a converged infrastructure platform like FlexPod, organizations can standardize their data center operations. With FlexPod, healthcare organizations can implement clinical and non-clinical systems on the same platform, thereby unifying data center operations.

Hospital capabilities deployed on a FlexPod



[Next: Benefits of deploying genomic workloads on FlexPod.](#)

Benefits of deploying genomic workloads on FlexPod

[Previous: Introduction.](#)

This section provides a brief list of benefits for running a genomics workload on a FlexPod converged infrastructure platform. Let's quickly describe the capabilities of a hospital. The following business architecture view shows a hospital's capabilities deployed on a hybrid-cloud-ready FlexPod converged infrastructure platform.

- **Avoid siloes in healthcare.** Silos in healthcare are a very real concern. Departments are often siloed into their own set of hardware and software not by choice but organically by evolution. For example, radiology, cardiology, EHR, genomics, analytics, revenue cycle, and other departments end up with their individual set of dedicated software and hardware. Healthcare organizations maintain a limited set of IT professionals to manage their hardware and software assets. The inflection point comes when this set of individuals are expected to manage a very diversified set of hardware and software. Heterogeneity is made worse by an incongruent set of processes brought to the healthcare organization by vendors.
- **Start small and grow.** The GATK tool kit is tuned for CPU execution, which best suites platforms like FlexPod. FlexPod enables independent scalability of network, compute, and storage. Start small and scale as your genomics capabilities and the environment grows. Healthcare organizations don't have to invest in specialized platforms to run genomic workloads. Instead, organizations can leverage versatile platforms like a FlexPod to run genomics and non-genomics workloads on the same platform. For example, if the pediatrics department wants to implement genomics capability, IT leadership can provision compute, storage, and networking on an existing FlexPod instance. As the genomics business unit grows, healthcare organization can scale their FlexPod platform as needed.

- **Single control pane and unparalleled flexibility.** Cisco Intersight significantly simplifies IT operations by bridging applications with infrastructure, providing visibility and management from bare-metal servers and hypervisors to serverless applications, thereby reducing costs and mitigating risk. This unified SaaS platform uses a unified Open API design that natively integrates with third-party platforms and tools. Moreover, it allows management to occur from your data center operations team on site or from anywhere by using a mobile app.

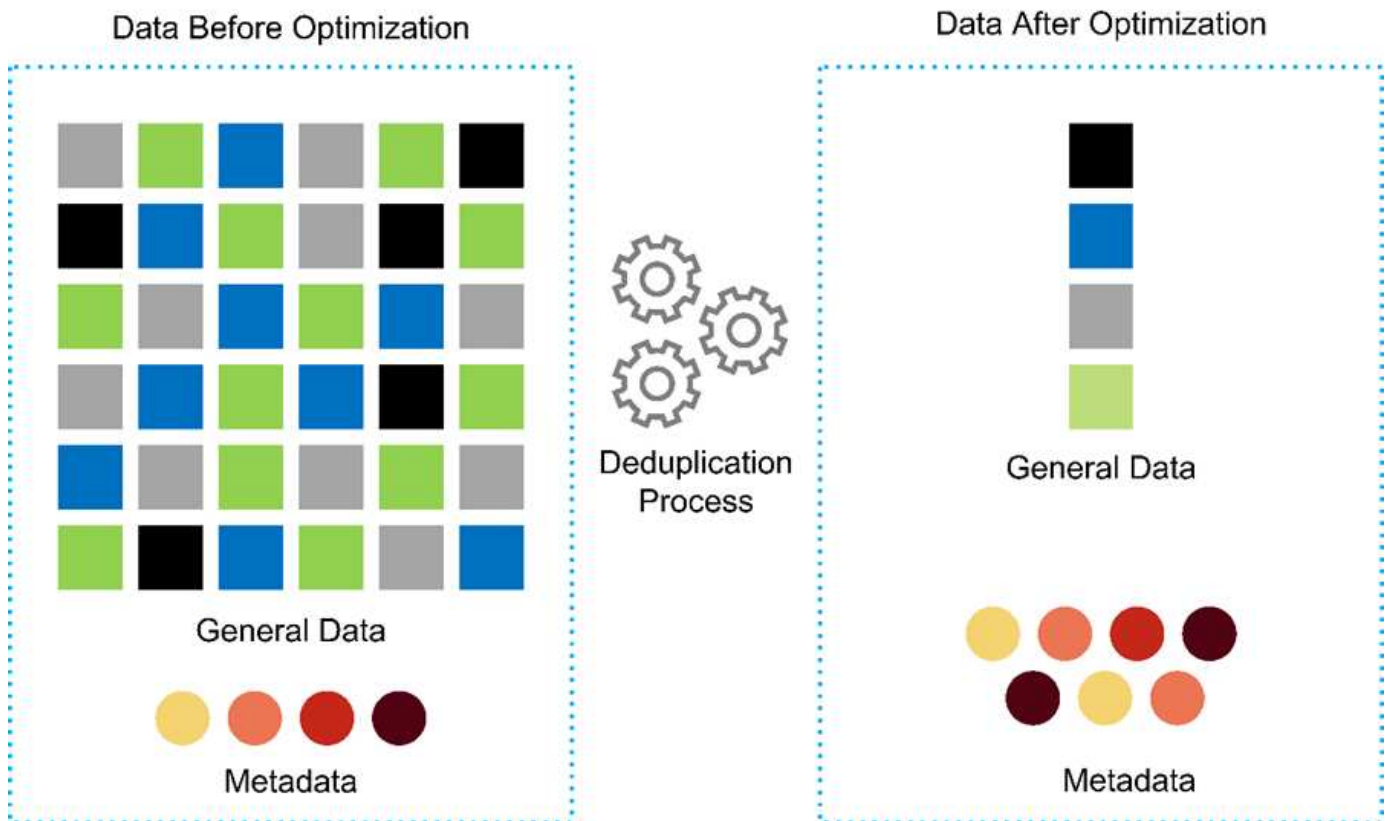
Users quickly unlock tangible value in their environment by leveraging Intersight as their management platform. Enabling automation for many daily manual tasks, Intersight removes errors and simplifies your daily operations. Moreover, advanced support capabilities facilitated by Intersight allow adopters to stay ahead of problems and accelerate issue resolution. Taken in combination, organizations spend far less time and money on their application infrastructure and more time on their core business development.

Leveraging Intersight management and FlexPod's easily scalable architecture enables organizations to run several genome workloads on a single FlexPod platform, increasing utilization and reducing total cost of ownership (TCO). FlexPod allows for flexible sizing, with choices starting with our small FlexPod Express and scaling into large FlexPod Datacenter implementations. With role-based access control capabilities built into Cisco Intersight, healthcare organizations can implement robust access control mechanisms, avoiding the need for separate infrastructure stacks. Multiple business units within the healthcare organization can leverage genomics as a key core competency.


Ultimately FlexPod helps simplify IT operations and lower operating costs, and it allows IT infrastructure admins to focus on tasks that help clinicians innovate rather than being relegated to keeping the lights on.

- **Validated design and guaranteed outcomes.** FlexPod design and deployment guides are validated to be repeatable, and they cover comprehensive configuration details and industry best practices that are needed to deploy a FlexPod with confidence. Cisco and NetApp validated design guides, deployment guides, and architectures help your healthcare or life science organization remove guesswork from the implementation of a validated and trusted platform from the beginning. With FlexPod, you can speed up deployment times and reduce cost, complexity, and risk. FlexPod validated designs and deployment guides establish FlexPod as the ideal platform for a variety of genomics workloads.
- **Innovation and agility.** FlexPod is recommended as an ideal platform by EHRs like Epic, Cerner, Meditech and imaging systems like Agfa, GE, Philips. For more information on [Epic honor roll](#) and target platform architecture, see the Epic userweb. Running genomics on [FlexPod](#) enables healthcare organizations to continue their journey of innovation with agility. With FlexPod, implementing organizational change comes naturally. When organizations standardize on a FlexPod platform, healthcare IT experts can provision their time, effort, and resources to innovate and thus be as agile as the ecosystem demands.
- **Data liberated.** With the FlexPod converged infrastructure platform and a NetApp ONTAP storage system, genomics data can be made available and accessible using a wide variety of protocols at scale from a single platform. FlexPod with NetApp ONTAP offers a simple, intuitive, and powerful hybrid cloud platform. Your data fabric powered by NetApp ONTAP weaves data together across sites, beyond physical boundaries, and across applications. Your data fabric is built for data-driven enterprises in a data-centric world. Data is created and used in multiple locations, and it often needs to be leveraged and shared with other locations, applications, and infrastructures. Therefore, you need a consistent and integrated way to manage it. FlexPod puts your IT team in control and simplifies ever-increasing IT complexity.
- **Secure multitenancy.** FlexPod uses FIPS 140-2 compliant cryptographic modules, hence enabling organizations to implement security as a foundational element, not an afterthought. FlexPod enables organizations implement secure multitenancy from a single converged infrastructure platform irrespective of the size of the platform. FlexPod with secured multitenancy and QoS help with workload separation and maximize utilization. This helps avoid capital being locked into specialized platforms that is potentially underutilized and requires a specialized skill set to manage.
- **Storage efficiency.** Genomics requires that the underlying storage have industry- leading storage efficiency capabilities. You can reduce storage costs with NetApp storage efficiency features such as

deduplication (inline and on demand), data compression, and data compaction ([ref](#)). NetApp deduplication provides block-level deduplication in a FlexVol volume. Essentially, deduplication removes duplicate blocks, storing only unique blocks in the FlexVol volume. Deduplication works with a high degree of granularity and operates on the active file system of the FlexVol volume. The following figure shows an overview of how NetApp deduplication works. Deduplication is application transparent. Therefore, it can be used to deduplicate data originating from any application that uses the NetApp system. You can run volume deduplication as an inline process and as a background process. You can configure it to run automatically, to be scheduled, or to run manually through the CLI, NetApp ONTAP System Manager, or NetApp Active IQ Unified Manager.



- **Enable genomics interoperability.** ONTAP FlexCache is a remote caching capability that simplifies file distribution, reduces WAN latency, and lowers WAN bandwidth costs, ([ref](#)). One of the key activities during genomic variant identification and annotation is collaboration between clinicians. ONTAP FlexCache technology increases data throughput even when collaborating clinicians are in different geographic locations. Given the typical size of a *.BAM file (1GB to 100s of GB), it is critical that the underlying platform can make files available to clinicians in different geographic locations. FlexPod with ONTAP FlexCache makes genomic data and applications truly multisite ready, which makes collaboration between researchers located around the world seamless with low latency and high throughput. Healthcare organizations running genomics applications in a multisite setting can scale-out using the data fabric to balance manageability with cost and speed.
- **Intelligent use of storage platform.** FlexPod with ONTAP auto-tiering and NetApp Fabric Pool technology simplifies data management. 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 more information, see [FlexPod with FabricPool](#). The following diagram provides a high-level overview of FabricPool and its benefits.



Automatic tiering
 Zero-touch management
 Preserves file system
 Lower cost of ownership
 Choice of object tier locations



- **Faster variant analysis and annotation.** The FlexPod platform is faster to deploy and operationalize. The FlexPod platform enables clinician collaboration by making data available at scale with low latency and increased throughput. Increased interoperability enables innovation. Healthcare organizations can run their genomic and non-genomic workloads side by side, which means organizations do not need specialized platforms to start their genomics journey.


FlexPod ONTAP routinely adds cutting edge features to the storage platform. FlexPod Datacenter is the optimal shared infrastructure foundation for deploying FC- NVMe to allow high-performance storage access to applications that need it. As FC- NVMe evolves to include high availability, multipathing, and additional operating system support, FlexPod is well suited as the platform of choice, providing the scalability and reliability needed to support these capabilities. ONTAP with faster I/O with end-to-end NVMe allows genomics analyses to completed faster ([ref](#)).

Sequenced raw genome data produces large file sizes, and it is important that these files are made available to the variant analyzers to reduce the total time it takes from sample collection to variant annotation. NVMe (nonvolatile memory express) when used as a storage access and data transport protocol provides unprecedented levels of throughput and the fastest response times. FlexPod deploys the NVMe protocol while accessing flash storage via the PCI express bus (PCIe). PCIe enables implementation of tens of thousands of command queues, increasing parallelization and throughput. One single protocol from storage to memory makes data access fast.

- **Agility for clinical research from the ground up.** Flexible, expandable storage capacity and performance allows the healthcare research organizations to optimize the environment in an elastic or just-in-time (JIT) manner. By decoupling storage from compute and network infrastructure, FlexPod platform can be scaled up and out without disruption. Using Cisco Intersight, the FlexPod platform can be managed with both built-in and custom automated workflows. Cisco Intersight workflows enable healthcare organizations to reduce application life-cycle management times. When an academic medical center requires that patient data be anonymized and made available to their center for research informatics and/or center for quality, their IT organization can leverage Cisco Intersight FlexPod workflows to take secure data backups, clone, and the

restore in a matter of seconds, not hours. With NetApp Trident and Kubernetes, IT organizations can provision new data scientists and make clinical data available for model development in a matter of minutes, sometimes even in seconds.

- **Protecting genome data.** NetApp SnapLock provides a special-purpose volume in which files can be stored and committed to a non-erasable, non-rewritable state. The user's production data residing in a FlexVol volume can be mirrored or vaulted to a SnapLock volume through NetApp SnapMirror or SnapVault technology. The files in the SnapLock volume, the volume itself, and its hosting aggregate cannot be deleted until the end of the retention period. Using ONTAP FPolicy software organizations can prevent ransomware attacks by disallowing operations on files with specific extensions. An FPolicy event can be triggered for specific file operations. The event is tied to a policy, which calls out the engine it needs to use. You might configure a policy with a set of file extensions that could potentially contain ransomware. When a file with a disallowed extension tries to perform an unauthorized operation, FPolicy prevents that operation from executing ([ref](#)).
- **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 offer a streamlined process for identifying and resolving FlexPod support issues, regardless of where the problem resides. The following figure provides an overview of the FlexPod Cooperative Support model. The customer contacts the vendor who might own the issue, and both Cisco and NetApp work cooperatively to resolve it. Cisco and NetApp have cross-company engineering and development teams that work hand in hand to resolve issues. This support model reduces loss of information during translation, enables trust, and reduces downtime.

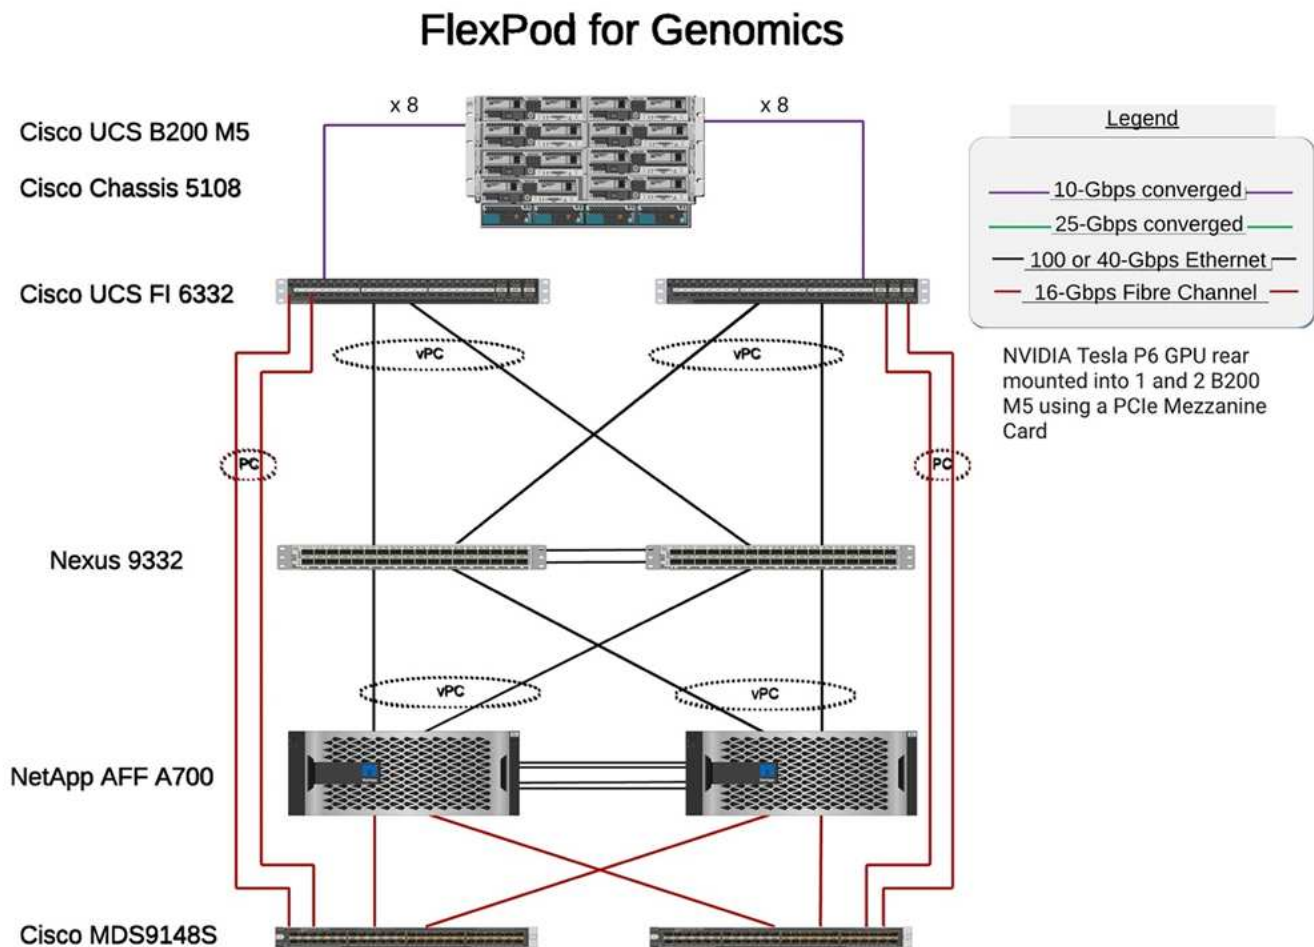


Next: Solution infrastructure hardware and software components.

Solution infrastructure hardware and software components

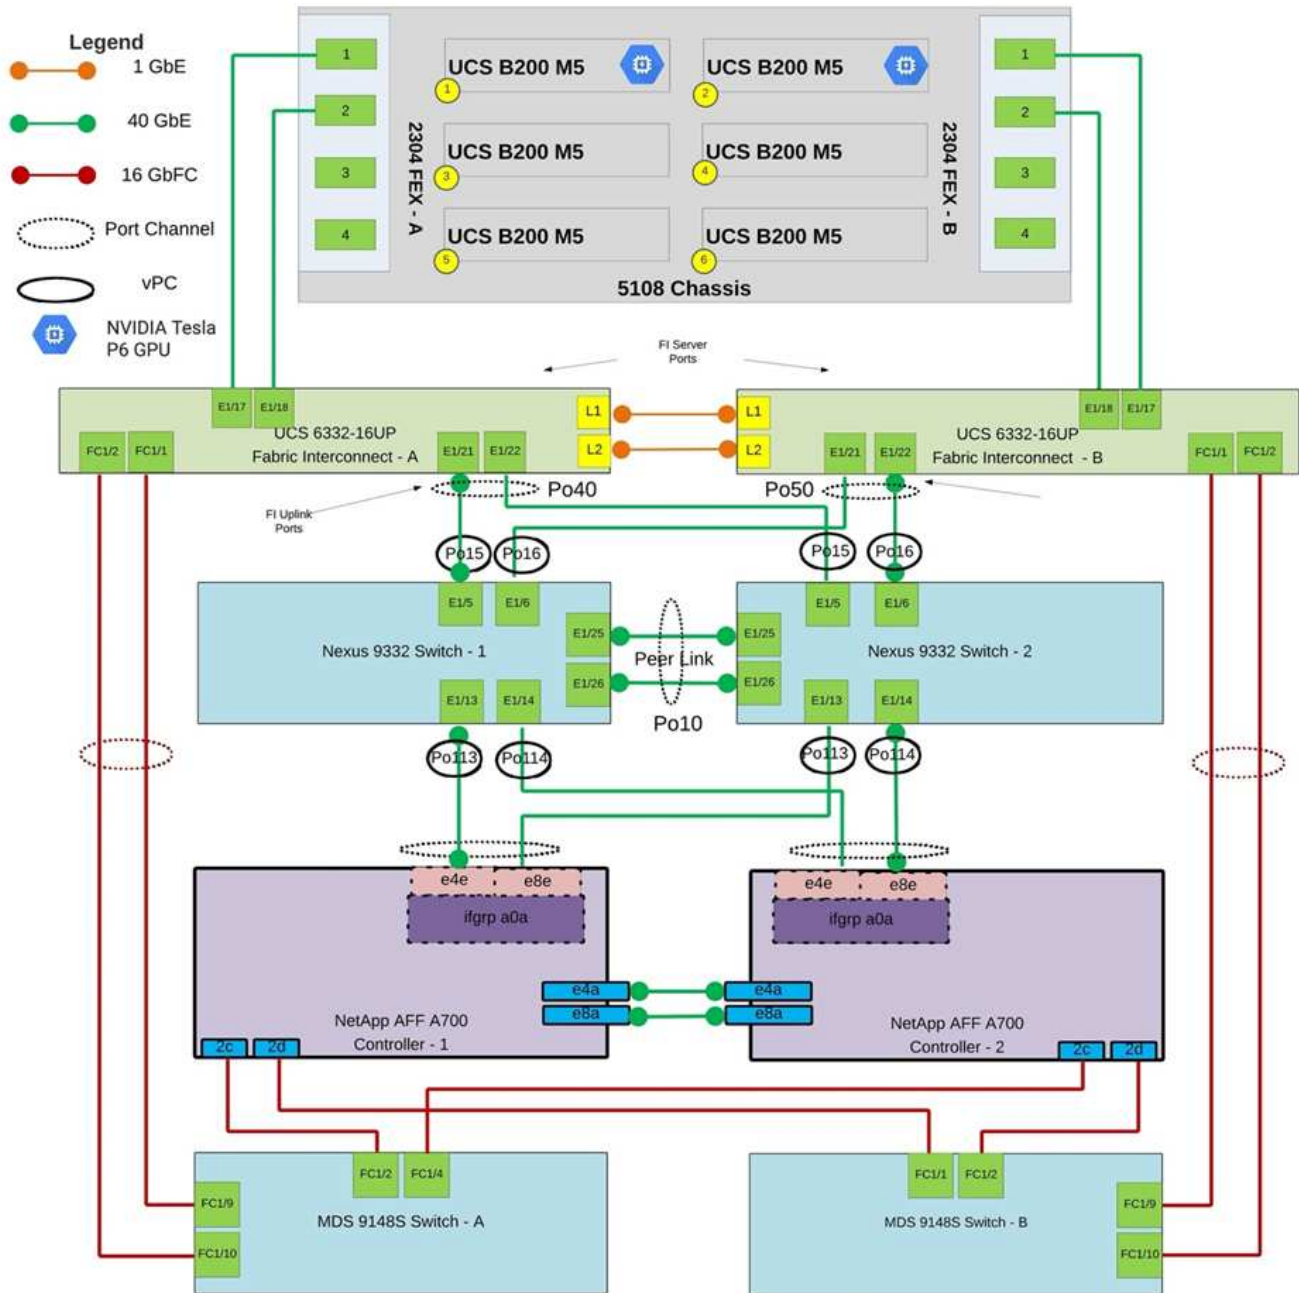
Previous: [Benefits of deploying genomic workloads on FlexPod.](#)

The following figure depicts the FlexPod system used for GATK setup and validation. We used [FlexPod Datacenter with VMware vSphere 7.0 and NetApp ONTAP 9.7 Cisco Validated Design \(CVD\)](#) during the setup process.



The following diagram depicts the FlexPod cabling details.

FlexPod for Genomics



The following table lists the hardware components used during the GATK testing enabling on a FlexPod. Here is the [NetApp Interoperability Matrix Tool \(IMT\)](#) and [Cisco Hardware Compatibility List \(HCL\)](#).

Layer	Product family	Quantity and model	Details
Compute	Cisco UCS 5108 chassis	1 or 2	
	Cisco UCS blade servers	6 B200 M5	Each with 2x 20 or more cores, 2.7GHz, and 128-384GB RAM

Layer	Product family	Quantity and model	Details
	Cisco UCS Virtual Interface Card (VIC)	Cisco UCS 1440	See the
	2x Cisco UCS Fabric Interconnects	6332	-
Network	Cisco Nexus switches	2x Cisco Nexus 9332	-
Storage network	IP network for storage access over SMB/CIFS, NFS, or iSCSI protocols	Same network switches as above	-
	Storage access over FC	2x Cisco MDS 9148S	-
Storage	NetApp AFF A700 all-flash storage system	1 Cluster	Cluster with two nodes
	Disk shelf	One DS224C or NS224 disk shelf	Fully populated with 24 drives
	SSD	24, 1.2TB or larger capacity	-

This table lists the infrastructure software.

Software	Product family	Version or release	Details
Various	Linux	RHEL 8.3	-
	Windows	Windows Server 2012 R2 (64 bit)	-
	NetApp ONTAP	ONTAP 9.8 or later	-
	Cisco UCS Fabric Interconnect	Cisco UCS Manager 4.1 or later	-
	Cisco Ethernet 3000 or 9000 series switches	For 9000 series, 7.0(3)I7(7) or later For 3000 series, 9.2(4) or later	-
	Cisco FC: Cisco MDS 9132T	8.4(1a) or later	-
	Hypervisor	VMware vSphere ESXi 7.0	-
Storage	Hypervisor management system	VMware vCenter Server 7.0 (vCSA) or later	-
Network	NetApp Virtual Storage Console (VSC)	VSC 9.7 or later	-
	NetApp SnapCenter	SnapCenter 4.3 or later	-
	Cisco UCS Manager	4.1(3c) or later	
Hypervisor	ESXi		

Software	Product family	Version or release	Details
Management	Hypervisor management systemVMware vCenter Server 7.0 (vCSA) or later		
	NetApp Virtual Storage Console (VSC)	VSC 9.7 or later	
	NetApp SnapCenter	SnapCenter 4.3 or later	
	Cisco UCS Manager	4.1(3c) or later	

Next: [Genomics - GATK setup and execution](#).

Genomics - GATK setup and execution

Previous: [Solution infrastructure hardware and software components](#).

According to the National Human Genome Research Institute ([NHGRI](#)), “Genomics is the study of all of a person’s genes (the genome), including interactions of these genes with each other and with a person’s environment. ”

According to the [NHGRI](#), “Deoxyribonucleic acid (DNA) is the chemical compound that contains the instructions needed to develop and direct the activities of nearly all living organisms. DNA molecules are made of two twisting, paired strands, often referred to as a double helix.” “An organism’s complete set of DNA is called its genome.”

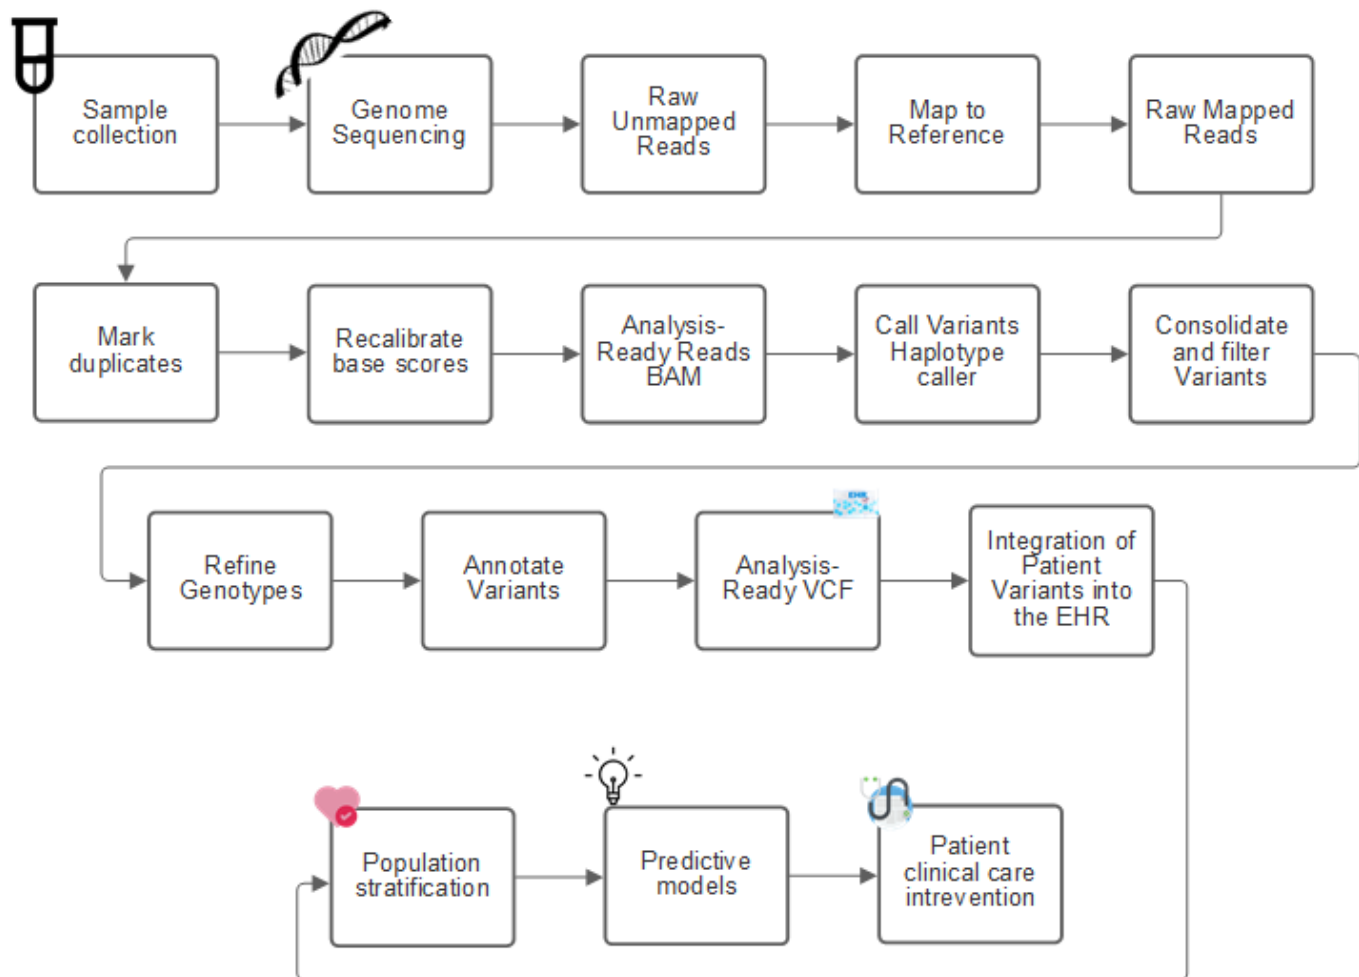
Sequencing is the process of determining the exact order of the bases in a strand of DNA. One of the most common types of sequencing used today is called sequencing by synthesis. This technique uses the emission of fluorescent signals to order the bases. Researchers can use DNA sequencing to search for genetic variations and any mutations that might play a role in the development or progression of a disease while a person is still in the embryonic stage.

From sample to variant identification, annotation, and prediction

At a high level, genomics can be classified into the following steps. This is not an exhaustive list:

1. Sample collection.
2. [Genome sequencing](#) using a sequencer to generate the raw data.
3. Preprocessing. For example, [deduplication](#) using [Picard](#).
4. Genomic analysis.
 - a. Mapping to a reference genome.
 - b. [Variant](#) identification and annotation typically performed using GATK and similar tools.
5. Integration into the electronic health record (EHR) system.
6. [Population stratification](#) and identification of genetic variation across geographical location and ethnic background.
7. [Predictive models](#) using significant single- nucleotide polymorphism.
8. [Validation](#).

The following figure shows the process from sampling to variant identification, annotation, and prediction.



The Human Genome project was completed in April 2003, and the project made a very high-quality simulation of the human genome sequence available in the public domain. This reference genome initiated an explosion in research and development of genomics capabilities. Virtually every human ailment has a signature in that human's genes. Up until recently, physicians were leveraging genes to predict and determine birth defects like sickle cell anemia, which is caused by a certain inheritance pattern caused by a change in a single gene. The treasure trove of data being made available by the human genome project led to the advent of the current state of genomics capabilities.

Genomics has a broad set of benefits. Here is a small set of benefits in the healthcare and life sciences domains:

- Better diagnosis at point of care
- Better prognosis
- Precision medicine
- Personalized treatment plans
- Better disease monitoring
- Reduction in adverse events
- Improved access to therapies
- Improved disease monitoring
- Effective clinical trial participation and better selection of patients for clinical trials based on genotypes.

Genomics is a "four-headed beast," because of the computational demands across the lifecycle of a dataset: acquisition, storage, distribution, and analysis.

Genome Analysis Toolkit (GATK)

GATK was developed as a data science platform at the [Broad Institute](#). GATK is a set of open-source tools that enable genome analysis, specifically variant discovery, identification, annotation, and genotyping. One of the benefits of GATK is that the set of tools and or commands can be chained to form a complete workflow. The primary challenges that Broad institute tackles are the following:

- Understand the root causes and biological mechanisms of diseases.
- Identify therapeutic interventions that act at the fundamental cause of a disease.
- Understand the line of sight from variants to function in human physiology.
- Create standards and policy [frameworks](#) for genome data representation, storage, analysis, security, and so on.
- Standardize and socialize interoperable genome aggregation databases (gnomAD).
- Genome-based monitoring, diagnosis, and treatment of patients with greater precision.
- Help implement tools that predict diseases well before symptoms appear.
- Create and empower a community of cross-disciplinary collaborators to help tackle the toughest and most important problems in biomedicine.

According to GATK and the Broad institute, genome sequencing should be treated as a protocol in a pathology lab; every task is well documented, optimized, reproducible, and consistent across samples and experiments. The following is a set of steps recommended by the Broad Institute, for more information, see the [GATK website](#).

FlexPod setup

Genomics workload validation includes a from-scratch setup of a FlexPod infrastructure platform. The FlexPod platform is highly available and can be independently scaled; for example, network, storage and compute can be scaled independently. We used the following Cisco validated design guide as the reference architecture document to set up the FlexPod environment: [FlexPod Datacenter with VMware vSphere 7.0 and NetApp ONTAP 9.7](#). See the following FlexPod platform set up highlights:

To perform FlexPod lab setup, complete the following steps:

1. FlexPod lab setup and validation uses the following IP4 reservations and VLANs.

IP Reservations

VLAN	IP Range	Subnet Mask	Purpose
3281	172.21.25 /24	255.255.255.0	IB-MGMT
3282	172.21.26 /24	255.255.255.0	vMotion
3283	172.21.27 /24	255.255.255.0	VM
3284	172.21.28 /24	255.255.255.0	NFS
3285	172.21.29 /24	255.255.255.0	iSCSI-A
3286	172.21.30 /24	255.255.255.0	iSCSI-B

2. Configure iSCSI-based boot LUNs on the ONTAP SVM.

ONTAP System Manager

LUNs

[+ Add](#)

	Name	Storage VM
▼	ESXi_Boot_Lun_1	Healthcare_SVM
▼	ESXi_Boot_Lun_2	Healthcare_SVM
▼	ESXi_Boot_Lun_3	Healthcare_SVM
▼	ESXi_Boot_Lun_4	Healthcare_SVM
▼	ESXi_Boot_Lun_5	Healthcare_SVM
▼	ESXi_Boot_Lun_6	Healthcare_SVM

3. Map LUNs to iSCSI initiator groups.

	Name	Storage VM	Volume	Size	IOPS	Latency (ms)	Throughput (MB/s)
▲	ESXi_Boot_Lun_1	Healthcare_SVM	ESXi_Boot_Vol	20 GB	3	0.16	0.01

STATUS	VOLUME	DESCRIPTION	SNAPSHOT COPIES (LOCAL)	SNAPMIRROR (LOCAL OR REMOTE)
Online	ESXi_Boot_Vol	-	Protected	Unprotected

SERIAL NUMBER	QOS POLICY GROUP	MAPPED TO INITIATORS	ID
80A4X+R8rAhP	-	GenomicsESXi_1 (1) iqn.1992-08.com.cisco:ucs-...	0

CAPACITY (AVAILABLE % TOTAL)	LUN FORMAT
95% 20 GB	VMware

PATH
/vol/ESXi_Boot_Vol/ESXi_Boot_Lun_1


	Name	Storage VM	Volume	Size	IOPS	Latency (ms)	Throughput (MB/s)
▼	ESXi_Boot_Lun_1	Healthcare_SVM	ESXi_Boot_Vol	20 GB	1	0.25	0.01
▲	ESXi_Boot_Lun_2	Healthcare_SVM	ESXi_Boot_Vol	20 GB	4	0.18	0.02

STATUS ✔ Online	VOLUME ESXi_Boot_Vol	DESCRIPTION -	SNAPSHOT COPIES (LOCAL) STATUS ✔ Protected	SNAPMIRROR (LOCAL OR REMOTE) STATUS ✔ Unprotected
SERIAL NUMBER 80A4X+R8rAhU	QOS POLICY GROUP -	MAPPED TO INITIATORS ▲ GenomicsESXi_2 (1) iqn.1992-08.com.cisco:ucs-...	ID 0	SNAPSHOT POLICY default
CAPACITY (AVAILABLE % TOTAL) 96% 20 GB	LUN FORMAT VMware			


4. Install vSphere 7.0 with iSCSI boot.
5. Register ESXi hosts with the vCenter.




6. Provision an NFS datastore infra_datastore_nfs on the ONTAP storage.



7. Add the datastore to the vCenter.



8. Using vCenter, add an NFS datastore to the ESXi hosts.



9. Using the vCenter, create a Red Hat Enterprise Linux (RHEL) 8.3 VM to run GATK.
10. An NFS datastore is presented to the VM and mounted at `/mnt/genomics`, which is used to store GATK executables, scripts, Binary Alignment Map (BAM) files, reference files, index files, dictionary files, and out files for variant calling.

```
[root@genomics1 genomics]# df | grep genomics
/dev/sdb                308587328  5699492 287142812   2% /mnt/genomics
[root@genomics1 genomics]#
```

GATK setup and execution

Install the following prerequisites on the RedHat Enterprise 8.3 Linux VM:

- Java 8 or SDK 1.8 or later
- Download GATK 4.2.0.0 from the Broad Institute [GitHub site](#). Genome sequence data is generally stored in the form of a series of tab-delimited ASCII columns. However, ASCII takes too much space to store. Therefore, a new standard evolved called a BAM (*.bam) file. A BAM file stores the sequence data in a compressed, indexed, and binary form. We [downloaded](#) a set of publicly available BAM files for GATK execution from the [public domain](#). We also downloaded index files (*.bai), dictionary files (*.dict), and reference data files (*.fasta) from the same public domain.

After downloading, the GATK tool kit has a jar file and a set of support scripts.

- `gatk-package-4.2.0.0-local.jar` executable
- `gatk` script file.

We downloaded the BAM files and the corresponding index, dictionary, and reference genome files for a family that consisted of father, mother, and son *.bam files.

Cromwell engine

Cromwell is an open-source engine geared towards scientific workflows that enables workflow management. The Cromwell engine can be run in two [modes](#), Server mode or a single-workflow Run mode. The behavior of the Cromwell engine can be controlled using the [Cromwell engine configuration file](#).

- **Server mode.** Enables [RESTful](#) execution of workflows in Cromwell engine.
- **Run mode.** Run mode is best suited for executing single workflows in Cromwell, [ref](#) for a complete set of available options in Run mode.

We use the Cromwell engine to execute the workflows and pipelines at scale. The Cromwell engine uses a user-friendly [workflow description language](#) (WDL)-based scripting language. Cromwell also supports a second

workflow scripting standard called the common workflow language (CWL). Throughout this technical report, we used WDL. WDL was originally developed by the Broad Institute for genome analysis pipelines. Using the WDL workflows can be implemented using several strategies, including the following:

- **Linear chaining.** As the name suggests, output from task#1 is sent to task #2 as input.
- **Multi-in/out.** This is similar to linear chaining in that each task can have multiple outputs being sent as input to subsequent tasks.
- **Scatter-gather.** This is one of the most powerful enterprise application integration (EAI) strategies available, especially when used in event-driven architecture. Each task executes in a decoupled fashion, and the output for each task is consolidated into the final output.

There are three steps when WDL is used to run GATK in a standalone mode:

1. Validate syntax using `womtool.jar`.

```
[root@genomics1 ~]# java -jar womtool.jar validate ghplo.wdl
```

2. Generate inputs JSON.

```
[root@genomics1 ~]# java -jar womtool.jar inputs ghplo.wdl > ghplo.json
```

3. Run the workflow using the Cromwell engine and `Cromwell.jar`.

```
[root@genomics1 ~]# java -jar cromwell.jar run ghplo.wdl --inputs  
ghplo.json
```

The GATK can be executed by using several methods; this document explores three of these methods.

Execution of GATK using the jar file

Let's look at a single variant call pipeline execution using the Haplotype variant caller.

```
[root@genomics1 ~]# java -Dsamjdk.use_async_io_read_samtools=false \  
-Dsamjdk.use_async_io_write_samtools=true \  
-Dsamjdk.use_async_io_write_tribble=false \  
-Dsamjdk.compression_level=2 \  
-jar /mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-local.jar \  
HaplotypeCaller \  
--input /mnt/genomics/GATK/TEST\ DATA/bam/workshop_1906_2-  
germline_bams_father.bam \  
--output workshop_1906_2-germline_bams_father.validation.vcf \  
--reference /mnt/genomics/GATK/TEST\ DATA/ref/workshop_1906_2-  
germline_ref_ref.fasta
```

In this method of execution, we use the GATK local execution jar file, we use a single java command to invoke the jar file, and we pass several parameters to the command.

1. This parameter indicates that we are invoking the `HaplotypeCaller` variant caller pipeline.
2. `--input` specifies the input BAM file.
3. `--output` specifies the variant output file in variant call format (*.vcf) ([ref](#)).
4. With the `--reference` parameter, we are passing a reference genome.

Once executed, output details can be found in the section "[Output for execution of GATK using the jar file.](#)"

Execution of GATK using `./gatk` script

GATK tool kit can be executed using the `./gatk` script. Let's examine the following command:

```
[root@genomics1 execution]# ./gatk \
--java-options "-Xmx4G" \
HaplotypeCaller \
-I /mnt/genomics/GATK/TEST\ DATA/bam/workshop_1906_2-
germline_bams_father.bam \
-R /mnt/genomics/GATK/TEST\ DATA/ref/workshop_1906_2-
germline_ref_ref.fasta \
-O /mnt/genomics/GATK/TEST\ DATA/variants.vcf
```

We pass several parameters to the command.

- This parameter indicates that we are invoking the `HaplotypeCaller` variant caller pipeline.
- `-I` specifies the input BAM file.
- `-O` specifies the variant output file in variant call format (*.vcf) ([ref](#)).
- With the `-R` parameter, we are passing a reference genome.

Once executed, output details can be found in the section "[Output for execution of GATK using the `./gatk` script.](#)"

Execution of GATK using Cromwell engine

We use the Cromwell engine to manage GATK execution. Let's examine the command line and its parameters.

```
[root@genomics1 genomics]# java -jar cromwell-65.jar \
run /mnt/genomics/GATK/seq/ghplo.wdl \
--inputs /mnt/genomics/GATK/seq/ghplo.json
```

Here, we invoke the Java command by passing the `-jar` parameter to indicate that we intend to execute a jar file, for example, `Cromwell-65.jar`. The next parameter passed (`run`) indicates that the Cromwell engine is running in Run mode, the other possible option is Server mode. The next parameter is `*.wdl` that the Run

mode should use to execute the pipelines. The next parameter is the set of input parameters to the workflows being executed.

Here's what the contents of the `ghplo.wdl` file look like:

```
[root@genomics1 seq]# cat ghplo.wdl
workflow helloHaplotypeCaller {
  call haplotypeCaller
}
task haplotypeCaller {
  File GATK
  File RefFasta
  File RefIndex
  File RefDict
  String sampleName
  File inputBAM
  File bamIndex
  command {
    java -jar ${GATK} \
      HaplotypeCaller \
      -R ${RefFasta} \
      -I ${inputBAM} \
      -O ${sampleName}.raw.indels.snps.vcf
  }
  output {
    File rawVCF = "${sampleName}.raw.indels.snps.vcf"
  }
}
[root@genomics1 seq]#
```

Here's the corresponding JSON file with the inputs to the Cromwell engine.


```
[root@genomics1 seq]# cat ghplo.json
{
  "helloHaplotypeCaller.haplotypeCaller.GATK": "/mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-local.jar",
  "helloHaplotypeCaller.haplotypeCaller.RefFasta": "/mnt/genomics/GATK/TEST DATA/ref/workshop_1906_2-germline_ref_ref.fasta",
  "helloHaplotypeCaller.haplotypeCaller.RefIndex": "/mnt/genomics/GATK/TEST DATA/ref/workshop_1906_2-germline_ref_ref.fasta.fai",
  "helloHaplotypeCaller.haplotypeCaller.RefDict": "/mnt/genomics/GATK/TEST DATA/ref/workshop_1906_2-germline_ref_ref.dict",
  "helloHaplotypeCaller.haplotypeCaller.sampleName": "fatherbam",
  "helloHaplotypeCaller.haplotypeCaller.inputBAM": "/mnt/genomics/GATK/TEST DATA/bam/workshop_1906_2-germline_bams_father.bam",
  "helloHaplotypeCaller.haplotypeCaller.bamIndex": "/mnt/genomics/GATK/TEST DATA/bam/workshop_1906_2-germline_bams_father.bai"
}
[root@genomics1 seq]#
```

Please note that Cromwell uses an in-memory database for the execution. Once executed, the output log can be seen in the section ["Output for execution of GATK using the Cromwell engine."](#)

For a comprehensive set of steps on how to execute GATK, see the [GATK documentation](#).

[Next: Output for execution of GATK using the jar file.](#)

Output for execution of GATK using the jar file

[Previous: Genomics - GATK setup and execution.](#)

Execution of GATK using the jar file produced the following sample output.

```
[root@genomics1 execution]# java -Dsamjdk.use_async_io_read_samtools=false \
-Dsamjdk.use_async_io_write_samtools=true \
-Dsamjdk.use_async_io_write_tribble=false \
-Dsamjdk.compression_level=2 \
-jar /mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-local.jar \
HaplotypeCaller \
--input /mnt/genomics/GATK/TEST\ DATA/bam/workshop_1906_2-germline_bams_father.bam \
--output workshop_1906_2-germline_bams_father.validation.vcf \
--reference /mnt/genomics/GATK/TEST\ DATA/ref/workshop_1906_2-germline_ref_ref.fasta \
22:52:58.430 INFO NativeLibraryLoader - Loading libgkl_compression.so
from jar:file:/mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-local.jar!/com/intel/gkl/native/libgkl_compression.so
```

```

Aug 17, 2021 10:52:58 PM
shaded.cloud_nio.com.google.auth.oauth2.ComputeEngineCredentials
runningOnComputeEngine
INFO: Failed to detect whether we are running on Google Compute Engine.
22:52:58.541 INFO HaplotypeCaller -
-----
22:52:58.542 INFO HaplotypeCaller - The Genome Analysis Toolkit (GATK)
v4.2.0.0
22:52:58.542 INFO HaplotypeCaller - For support and documentation go to
https://software.broadinstitute.org/gatk/
22:52:58.542 INFO HaplotypeCaller - Executing as
root@genomics1.healthylife.fp on Linux v4.18.0-305.3.1.el8_4.x86_64 amd64
22:52:58.542 INFO HaplotypeCaller - Java runtime: OpenJDK 64-Bit Server
VM v1.8.0_302-b08
22:52:58.542 INFO HaplotypeCaller - Start Date/Time: August 17, 2021
10:52:58 PM EDT
22:52:58.542 INFO HaplotypeCaller -
-----
22:52:58.542 INFO HaplotypeCaller -
-----
22:52:58.542 INFO HaplotypeCaller - HTSJDK Version: 2.24.0
22:52:58.542 INFO HaplotypeCaller - Picard Version: 2.25.0
22:52:58.542 INFO HaplotypeCaller - Built for Spark Version: 2.4.5
22:52:58.542 INFO HaplotypeCaller - HTSJDK Defaults.COMPRESSION_LEVEL : 2
22:52:58.543 INFO HaplotypeCaller - HTSJDK
Defaults.USE_ASYNC_IO_READ_FOR_SAMTOOLS : false
22:52:58.543 INFO HaplotypeCaller - HTSJDK
Defaults.USE_ASYNC_IO_WRITE_FOR_SAMTOOLS : true
22:52:58.543 INFO HaplotypeCaller - HTSJDK
Defaults.USE_ASYNC_IO_WRITE_FOR_TRIBBLE : false
22:52:58.543 INFO HaplotypeCaller - Deflater: IntelDeflater
22:52:58.543 INFO HaplotypeCaller - Inflater: IntelInflater
22:52:58.543 INFO HaplotypeCaller - GCS max retries/reopens: 20
22:52:58.543 INFO HaplotypeCaller - Requester pays: disabled
22:52:58.543 INFO HaplotypeCaller - Initializing engine
22:52:58.804 INFO HaplotypeCaller - Done initializing engine
22:52:58.809 INFO HaplotypeCallerEngine - Disabling physical phasing,
which is supported only for reference-model confidence output
22:52:58.820 INFO NativeLibraryLoader - Loading libgkl_utils.so from
jar:file:/mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-
local.jar!/com/intel/gkl/native/libgkl_utils.so
22:52:58.821 INFO NativeLibraryLoader - Loading libgkl_pairhmm_omp.so
from jar:file:/mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-
local.jar!/com/intel/gkl/native/libgkl_pairhmm_omp.so
22:52:58.854 INFO IntelPairHmm - Using CPU-supported AVX-512 instructions
22:52:58.854 INFO IntelPairHmm - Flush-to-zero (FTZ) is enabled when

```

```

running PairHMM
22:52:58.854 INFO   IntelPairHmm - Available threads: 16
22:52:58.854 INFO   IntelPairHmm - Requested threads: 4
22:52:58.854 INFO   PairHMM - Using the OpenMP multi-threaded AVX-
accelerated native PairHMM implementation
22:52:58.872 INFO   ProgressMeter - Starting traversal
22:52:58.873 INFO   ProgressMeter -           Current Locus   Elapsed Minutes
Regions Processed   Regions/Minute
22:53:00.733 WARN   InbreedingCoeff - InbreedingCoeff will not be
calculated at position 20:9999900 and possibly subsequent; at least 10
samples must have called genotypes
22:53:08.873 INFO   ProgressMeter -           20:17538652           0.2
58900              353400.0
22:53:17.681 INFO   HaplotypeCaller - 405 read(s) filtered by:
MappingQualityReadFilter
0 read(s) filtered by: MappingQualityAvailableReadFilter
0 read(s) filtered by: MappedReadFilter
0 read(s) filtered by: NotSecondaryAlignmentReadFilter
6628 read(s) filtered by: NotDuplicateReadFilter
0 read(s) filtered by: PassesVendorQualityCheckReadFilter
0 read(s) filtered by: NonZeroReferenceLengthAlignmentReadFilter
0 read(s) filtered by: GoodCigarReadFilter
0 read(s) filtered by: WellformedReadFilter
7033 total reads filtered
22:53:17.681 INFO   ProgressMeter -           20:63024652           0.3
210522             671592.9
22:53:17.681 INFO   ProgressMeter - Traversal complete. Processed 210522
total regions in 0.3 minutes.
22:53:17.687 INFO   VectorLoglessPairHMM - Time spent in setup for JNI call
: 0.010347438
22:53:17.687 INFO   PairHMM - Total compute time in PairHMM
computeLogLikelihoods() : 0.259172573
22:53:17.687 INFO   SmithWatermanAligner - Total compute time in java
Smith-Waterman : 1.27 sec
22:53:17.687 INFO   HaplotypeCaller - Shutting down engine
[August 17, 2021 10:53:17 PM EDT]
org.broadinstitute.hellbender.tools.walkers.haplotypecaller.HaplotypeCalle
r done. Elapsed time: 0.32 minutes.
Runtime.totalMemory()=5561122816
[root@genomics1 execution]#

```

Notice that the output file is located at the location specified after the execution.

Next: [Output for execution of GATK using the ./gatk script.](#)

Output for execution of GATK using the `./gatk` script

[Previous: Output for execution of GATK using the jar file.](#)

The execution of GATK using the `./gatk` script produced the following sample output.

```
[root@genomics1 gatk-4.2.0.0]# ./gatk --java-options "-Xmx4G" \
HaplotypeCaller \
-I /mnt/genomics/GATK/TEST\ DATA/bam/workshop_1906_2-
germline_bams_father.bam \
-R /mnt/genomics/GATK/TEST\ DATA/ref/workshop_1906_2-
germline_ref_ref.fasta \
-O /mnt/genomics/GATK/TEST\ DATA/variants.vcf
Using GATK jar /mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-
local.jar
Running:
    java -Dsamjdk.use_async_io_read_samtools=false
-Dsamjdk.use_async_io_write_samtools=true
-Dsamjdk.use_async_io_write_tribble=false -Dsamjdk.compression_level=2
-Xmx4G -jar /mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-local.jar
HaplotypeCaller -I /mnt/genomics/GATK/TEST DATA/bam/workshop_1906_2-
germline_bams_father.bam -R /mnt/genomics/GATK/TEST
DATA/ref/workshop_1906_2-germline_ref_ref.fasta -O /mnt/genomics/GATK/TEST
DATA/variants.vcf
23:29:45.553 INFO  NativeLibraryLoader - Loading libgkl_compression.so
from jar:file:/mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-
local.jar!/com/intel/gkl/native/libgkl_compression.so
Aug 17, 2021 11:29:45 PM
shaded.cloud_nio.com.google.auth.oauth2.ComputeEngineCredentials
runningOnComputeEngine
INFO: Failed to detect whether we are running on Google Compute Engine.
23:29:45.686 INFO  HaplotypeCaller -
-----
23:29:45.686 INFO  HaplotypeCaller - The Genome Analysis Toolkit (GATK)
v4.2.0.0
23:29:45.686 INFO  HaplotypeCaller - For support and documentation go to
https://software.broadinstitute.org/gatk/
23:29:45.687 INFO  HaplotypeCaller - Executing as
root@genomics1.healthyliife.fp on Linux v4.18.0-305.3.1.el8_4.x86_64 amd64
23:29:45.687 INFO  HaplotypeCaller - Java runtime: OpenJDK 64-Bit Server
VM v11.0.12+7-LTS
23:29:45.687 INFO  HaplotypeCaller - Start Date/Time: August 17, 2021 at
11:29:45 PM EDT
23:29:45.687 INFO  HaplotypeCaller -
-----
23:29:45.687 INFO  HaplotypeCaller -
```

```

-----
23:29:45.687 INFO HaplotypeCaller - HTSJDK Version: 2.24.0
23:29:45.687 INFO HaplotypeCaller - Picard Version: 2.25.0
23:29:45.687 INFO HaplotypeCaller - Built for Spark Version: 2.4.5
23:29:45.688 INFO HaplotypeCaller - HTSJDK Defaults.COMPRESSION_LEVEL : 2
23:29:45.688 INFO HaplotypeCaller - HTSJDK
Defaults.USE_ASYNC_IO_READ_FOR_SAMTOOLS : false
23:29:45.688 INFO HaplotypeCaller - HTSJDK
Defaults.USE_ASYNC_IO_WRITE_FOR_SAMTOOLS : true
23:29:45.688 INFO HaplotypeCaller - HTSJDK
Defaults.USE_ASYNC_IO_WRITE_FOR_TRIBBLE : false
23:29:45.688 INFO HaplotypeCaller - Deflater: IntelDeflater
23:29:45.688 INFO HaplotypeCaller - Inflater: IntelInflater
23:29:45.688 INFO HaplotypeCaller - GCS max retries/reopens: 20
23:29:45.688 INFO HaplotypeCaller - Requester pays: disabled
23:29:45.688 INFO HaplotypeCaller - Initializing engine
23:29:45.804 INFO HaplotypeCaller - Done initializing engine
23:29:45.809 INFO HaplotypeCallerEngine - Disabling physical phasing,
which is supported only for reference-model confidence output
23:29:45.818 INFO NativeLibraryLoader - Loading libgkl_utils.so from
jar:file:/mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-
local.jar!/com/intel/gkl/native/libgkl_utils.so
23:29:45.819 INFO NativeLibraryLoader - Loading libgkl_pairhmm_omp.so
from jar:file:/mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-
local.jar!/com/intel/gkl/native/libgkl_pairhmm_omp.so
23:29:45.852 INFO IntelPairHmm - Using CPU-supported AVX-512 instructions
23:29:45.852 INFO IntelPairHmm - Flush-to-zero (FTZ) is enabled when
running PairHMM
23:29:45.852 INFO IntelPairHmm - Available threads: 16
23:29:45.852 INFO IntelPairHmm - Requested threads: 4
23:29:45.852 INFO PairHMM - Using the OpenMP multi-threaded AVX-
accelerated native PairHMM implementation
23:29:45.868 INFO ProgressMeter - Starting traversal
23:29:45.868 INFO ProgressMeter -          Current Locus  Elapsed Minutes
Regions Processed  Regions/Minute
23:29:47.772 WARN InbreedingCoeff - InbreedingCoeff will not be
calculated at position 20:9999900 and possibly subsequent; at least 10
samples must have called genotypes
23:29:55.868 INFO ProgressMeter -          20:18885652          0.2
63390          380340.0
23:30:04.389 INFO HaplotypeCaller - 405 read(s) filtered by:
MappingQualityReadFilter
0 read(s) filtered by: MappingQualityAvailableReadFilter
0 read(s) filtered by: MappedReadFilter
0 read(s) filtered by: NotSecondaryAlignmentReadFilter
6628 read(s) filtered by: NotDuplicateReadFilter

```

```

0 read(s) filtered by: PassesVendorQualityCheckReadFilter
0 read(s) filtered by: NonZeroReferenceLengthAlignmentReadFilter
0 read(s) filtered by: GoodCigarReadFilter
0 read(s) filtered by: WellformedReadFilter
7033 total reads filtered
23:30:04.389 INFO ProgressMeter - 20:63024652 0.3
210522 681999.9
23:30:04.389 INFO ProgressMeter - Traversal complete. Processed 210522
total regions in 0.3 minutes.
23:30:04.395 INFO VectorLoglessPairHMM - Time spent in setup for JNI call
: 0.012129203000000002
23:30:04.395 INFO PairHMM - Total compute time in PairHMM
computeLogLikelihoods() : 0.267345217
23:30:04.395 INFO SmithWatermanAligner - Total compute time in java
Smith-Waterman : 1.23 sec
23:30:04.395 INFO HaplotypeCaller - Shutting down engine
[August 17, 2021 at 11:30:04 PM EDT]
org.broadinstitute.hellbender.tools.walkers.haplotypecaller.HaplotypeCalle
r done. Elapsed time: 0.31 minutes.
Runtime.totalMemory()=2111832064
[root@genomics1 gatk-4.2.0.0]#

```

Notice that the output file is located at the location specified after the execution.

[Next: Output for execution of GATK using the Cromwell engine.](#)

Output for execution of GATK using the Cromwell engine

[Previous: Output for execution of GATK using the ./gatk script.](#)

The execution of GATK using the Cromwell engine produced the following sample output.

```

[root@genomics1 genomics]# java -jar cromwell-65.jar run
/mnt/genomics/GATK/seq/ghplo.wdl --inputs
/mnt/genomics/GATK/seq/ghplo.json
[2021-08-18 17:10:50,78] [info] Running with database db.url =
jdbc:hsqldb:mem:856a1f0d-9a0d-42e5-9199-
5e6c1d0f72dd;shutdown=false;hsqldb.tx=mvcc
[2021-08-18 17:10:57,74] [info] Running migration
RenameWorkflowOptionsInMetadata with a read batch size of 100000 and a
write batch size of 100000
[2021-08-18 17:10:57,75] [info] [RenameWorkflowOptionsInMetadata] 100%
[2021-08-18 17:10:57,83] [info] Running with database db.url =
jdbc:hsqldb:mem:6afe0252-2dc9-4e57-8674-
ce63c67aa142;shutdown=false;hsqldb.tx=mvcc
[2021-08-18 17:10:58,17] [info] Slf4jLogger started

```

```

[2021-08-18 17:10:58,33] [info] Workflow heartbeat configuration:
{
  "cromwellId" : "cromid-41b7e30",
  "heartbeatInterval" : "2 minutes",
  "ttl" : "10 minutes",
  "failureShutdownDuration" : "5 minutes",
  "writeBatchSize" : 10000,
  "writeThreshold" : 10000
}
[2021-08-18 17:10:58,38] [info] Metadata summary refreshing every 1
second.
[2021-08-18 17:10:58,38] [info] No metadata archiver defined in config
[2021-08-18 17:10:58,38] [info] No metadata deleter defined in config
[2021-08-18 17:10:58,40] [info] KvWriteActor configured to flush with
batch size 200 and process rate 5 seconds.
[2021-08-18 17:10:58,40] [info] WriteMetadataActor configured to flush
with batch size 200 and process rate 5 seconds.
[2021-08-18 17:10:58,44] [info] CallCacheWriteActor configured to flush
with batch size 100 and process rate 3 seconds.
[2021-08-18 17:10:58,44] [warn] 'docker.hash-lookup.gcr-api-queries-per-
100-seconds' is being deprecated, use 'docker.hash-lookup.gcr.throttle'
instead (see reference.conf)
[2021-08-18 17:10:58,54] [info] JobExecutionTokenDispenser - Distribution
rate: 50 per 1 seconds.
[2021-08-18 17:10:58,58] [info] SingleWorkflowRunnerActor: Version 65
[2021-08-18 17:10:58,58] [info] SingleWorkflowRunnerActor: Submitting
workflow
[2021-08-18 17:10:58,64] [info] Unspecified type (Unspecified version)
workflow 3e246147-b1a9-41dc-8679-319f81b7701e submitted
[2021-08-18 17:10:58,66] [info] SingleWorkflowRunnerActor: Workflow
submitted 3e246147-b1a9-41dc-8679-319f81b7701e
[2021-08-18 17:10:58,66] [info] 1 new workflows fetched by cromid-41b7e30:
3e246147-b1a9-41dc-8679-319f81b7701e
[2021-08-18 17:10:58,67] [info] WorkflowManagerActor: Starting workflow
3e246147-b1a9-41dc-8679-319f81b7701e
[2021-08-18 17:10:58,68] [info] WorkflowManagerActor: Successfully started
WorkflowActor-3e246147-b1a9-41dc-8679-319f81b7701e
[2021-08-18 17:10:58,68] [info] Retrieved 1 workflows from the
WorkflowStoreActor
[2021-08-18 17:10:58,70] [info] WorkflowStoreHeartbeatWriteActor
configured to flush with batch size 10000 and process rate 2 minutes.
[2021-08-18 17:10:58,76] [info] MaterializeWorkflowDescriptorActor
[3e246147]: Parsing workflow as WDL draft-2
[2021-08-18 17:10:59,34] [info] MaterializeWorkflowDescriptorActor
[3e246147]: Call-to-Backend assignments:
helloHaplotypeCaller.haplotypeCaller -> Local

```



```

[2021-08-18 17:11:00,54] [info] WorkflowExecutionActor-3e246147-b1a9-41dc-8679-319f81b7701e [3e246147]: Starting
helloHaplotypeCaller.haplotypeCaller
[2021-08-18 17:11:01,56] [info] Assigned new job execution tokens to the
following groups: 3e246147: 1
[2021-08-18 17:11:01,70] [info] BackgroundConfigAsyncJobExecutionActor
[3e246147helloHaplotypeCaller.haplotypeCaller:NA:1]: java -jar
/mnt/genomics/cromwell-executions/helloHaplotypeCaller/3e246147-b1a9-41dc-8679-319f81b7701e/call-haplotypeCaller/inputs/-179397211/gatk-package-4.2.0.0-local.jar \
    HaplotypeCaller \
    -R /mnt/genomics/cromwell-executions/helloHaplotypeCaller/3e246147-b1a9-41dc-8679-319f81b7701e/call-haplotypeCaller/inputs/604632695/workshop_1906_2-germline_ref_ref.fasta \
    -I /mnt/genomics/cromwell-executions/helloHaplotypeCaller/3e246147-b1a9-41dc-8679-319f81b7701e/call-haplotypeCaller/inputs/604617202/workshop_1906_2-germline_bams_father.bam \
    -O fatherbam.raw.indels.snps.vcf
[2021-08-18 17:11:01,72] [info] BackgroundConfigAsyncJobExecutionActor
[3e246147helloHaplotypeCaller.haplotypeCaller:NA:1]: executing: /bin/bash
/mnt/genomics/cromwell-executions/helloHaplotypeCaller/3e246147-b1a9-41dc-8679-319f81b7701e/call-haplotypeCaller/execution/script
[2021-08-18 17:11:03,49] [info] BackgroundConfigAsyncJobExecutionActor
[3e246147helloHaplotypeCaller.haplotypeCaller:NA:1]: job id: 26867
[2021-08-18 17:11:03,53] [info] BackgroundConfigAsyncJobExecutionActor
[3e246147helloHaplotypeCaller.haplotypeCaller:NA:1]: Status change from -
to WaitingForReturnCode
[2021-08-18 17:11:03,54] [info] Not triggering log of token queue status.
Effective log interval = None
[2021-08-18 17:11:23,65] [info] BackgroundConfigAsyncJobExecutionActor
[3e246147helloHaplotypeCaller.haplotypeCaller:NA:1]: Status change from
WaitingForReturnCode to Done
[2021-08-18 17:11:25,04] [info] WorkflowExecutionActor-3e246147-b1a9-41dc-8679-319f81b7701e [3e246147]: Workflow helloHaplotypeCaller complete.
Final Outputs:
{
  "helloHaplotypeCaller.haplotypeCaller.rawVCF": "/mnt/genomics/cromwell-executions/helloHaplotypeCaller/3e246147-b1a9-41dc-8679-319f81b7701e/call-haplotypeCaller/execution/fatherbam.raw.indels.snps.vcf"
}
[2021-08-18 17:11:28,43] [info] WorkflowManagerActor: Workflow actor for
3e246147-b1a9-41dc-8679-319f81b7701e completed with status 'Succeeded'.
The workflow will be removed from the workflow store.
[2021-08-18 17:11:32,24] [info] SingleWorkflowRunnerActor workflow
finished with status 'Succeeded'.

```

```

{
  "outputs": {
    "helloHaplotypeCaller.haplotypeCaller.rawVCF":
"/mnt/genomics/cromwell-executions/helloHaplotypeCaller/3e246147-b1a9-
41dc-8679-319f81b7701e/call-
haplotypeCaller/execution/fatherbam.raw.indels.snps.vcf"
  },
  "id": "3e246147-b1a9-41dc-8679-319f81b7701e"
}
[2021-08-18 17:11:33,45] [info] Workflow polling stopped
[2021-08-18 17:11:33,46] [info] 0 workflows released by cromid-41b7e30
[2021-08-18 17:11:33,46] [info] Shutting down WorkflowStoreActor - Timeout
= 5 seconds
[2021-08-18 17:11:33,46] [info] Shutting down WorkflowLogCopyRouter -
Timeout = 5 seconds
[2021-08-18 17:11:33,46] [info] Shutting down JobExecutionTokenDispenser -
Timeout = 5 seconds
[2021-08-18 17:11:33,46] [info] Aborting all running workflows.
[2021-08-18 17:11:33,46] [info] JobExecutionTokenDispenser stopped
[2021-08-18 17:11:33,46] [info] WorkflowStoreActor stopped
[2021-08-18 17:11:33,47] [info] WorkflowLogCopyRouter stopped
[2021-08-18 17:11:33,47] [info] Shutting down WorkflowManagerActor -
Timeout = 3600 seconds
[2021-08-18 17:11:33,47] [info] WorkflowManagerActor: All workflows
finished
[2021-08-18 17:11:33,47] [info] WorkflowManagerActor stopped
[2021-08-18 17:11:33,64] [info] Connection pools shut down
[2021-08-18 17:11:33,64] [info] Shutting down SubWorkflowStoreActor -
Timeout = 1800 seconds
[2021-08-18 17:11:33,64] [info] Shutting down JobStoreActor - Timeout =
1800 seconds
[2021-08-18 17:11:33,64] [info] Shutting down CallCacheWriteActor -
Timeout = 1800 seconds
[2021-08-18 17:11:33,64] [info] SubWorkflowStoreActor stopped
[2021-08-18 17:11:33,64] [info] Shutting down ServiceRegistryActor -
Timeout = 1800 seconds
[2021-08-18 17:11:33,64] [info] Shutting down DockerHashActor - Timeout =
1800 seconds
[2021-08-18 17:11:33,64] [info] Shutting down IoProxy - Timeout = 1800
seconds
[2021-08-18 17:11:33,64] [info] CallCacheWriteActor Shutting down: 0
queued messages to process
[2021-08-18 17:11:33,64] [info] JobStoreActor stopped
[2021-08-18 17:11:33,64] [info] CallCacheWriteActor stopped
[2021-08-18 17:11:33,64] [info] KvWriteActor Shutting down: 0 queued
messages to process

```

```
[2021-08-18 17:11:33,64] [info] IoProxy stopped
[2021-08-18 17:11:33,64] [info] WriteMetadataActor Shutting down: 0 queued
messages to process
[2021-08-18 17:11:33,65] [info] ServiceRegistryActor stopped
[2021-08-18 17:11:33,65] [info] DockerHashActor stopped
[2021-08-18 17:11:33,67] [info] Database closed
[2021-08-18 17:11:33,67] [info] Stream materializer shut down
[2021-08-18 17:11:33,67] [info] WDL HTTP import resolver closed
[root@genomics1 genomics]#
```

[Next: GPU setup.](#)

GPU setup

[Previous: Output for execution of GATK using the Cromwell engine.](#)

At the time of publication, the GATK tool does not have native support for GPU-based execution on premises. The following setup and guidance is provided to enable the readers understand how simple it is to use FlexPod with a rear-mounted NVIDIA Tesla P6 GPU using a PCIe mezzanine card for GATK.

We used the following Cisco-Validated Design (CVD) as the reference architecture and best-practice guide to set up the FlexPod environment so that we can run applications that use GPUs.

- [FlexPod Datacenter for AI/ML with Cisco UCS 480 ML for Deep Learning](#)

Here is a set of key takeaways during this setup:

1. We used a PCIe NVIDIA Tesla P6 GPU in a mezzanine slot in the UCS B200 M5 servers.

Equipment / Chassis / Chassis 1 / Servers / Server 1				
<div> General Inventory Virtual Machines Installed Firmware CIMC Sessions SEL Logs VIF Paths Health </div>				
<div> Motherboard CIMC CPUs GPUs Memory Adapters HBAs NICs iSCSI vNICs Security </div>				
<div> Advanced Filter Export Print </div>				
Name	ID	Model	Serial	Mode
Graphics Card 2	2	UCSB-GPU-P6-R	FCH212373V7	Compute

Equipment / Chassis / Chassis 1 / Servers / Server 2				
<div> General Inventory Virtual Machines Installed Firmware CIMC Sessions SEL Logs VIF Paths Health </div>				
<div> Motherboard CIMC CPUs GPUs Memory Adapters HBAs NICs iSCSI vNICs Security </div>				
<div> Advanced Filter Export Print </div>				
Name	ID	Model	Serial	Mode
Graphics Card 2	2	UCSB-GPU-P6-R	FCH212373Y1	Compute

2. For this setup, we registered on the NVIDIA partner portal and obtained an evaluation license (also known as an entitlement) to be able to use the GPUs in compute mode.
3. We downloaded the NVIDIA vGPU software required from the NVIDIA partner website.
4. We downloaded the entitlement *.bin file from the NVIDIA partner website.
5. We installed an NVIDIA vGPU license server and added the entitlements to the license server using the *.bin file downloaded from the NVIDIA partner site.
6. Make sure to choose the correct NVIDIA vGPU software version for your deployment on the NVIDIA partner portal. For this setup we used driver version 460.73.02.
7. This command installs the [NVIDIA vGPU Manager](#) in ESXi.

```
[root@localhost:~] esxcli software vib install -v
/vmfs/volumes/infra_datastore_nfs/nvidia/vib/NVIDIA_bootbank_NVIDIA-
VMware_ESXi_7.0_Host_Driver_460.73.02-1OEM.700.0.0.15525992.vib
Installation Result
Message: Operation finished successfully.
Reboot Required: false
VIBs Installed: NVIDIA_bootbank_NVIDIA-
VMware_ESXi_7.0_Host_Driver_460.73.02-1OEM.700.0.0.15525992
VIBs Removed:
VIBs Skipped:
```


8. After rebooting the ESXi server, run the following command to validate the installation and check the health of the GPUs.

```

[root@localhost:~] nvidia-smi
Wed Aug 18 21:37:19 2021
+-----+
+-----+
| NVIDIA-SMI 460.73.02      Driver Version: 460.73.02      CUDA Version: N/A
|
|-----+-----+
+-----+
| GPU   Name               Persistence-M| Bus-Id        Disp.A | Volatile
Uncorr. ECC |
| Fan   Temp   Perf   Pwr:Usage/Cap|      Memory-Usage | GPU-Util
Compute M. |
|
|-----+-----+
MIG M. |
|=====+=====+=====+
=====|
|    0   Tesla P6             On      | 00000000:D8:00.0 Off |
0 |
| N/A    35C    P8      9W /  90W |  15208MiB /  15359MiB |      0%
Default |
|
|-----+-----+
N/A |
+-----+-----+
+-----+
+-----+
+-----+
+-----+
| Processes:
|
| GPU   GI    CI          PID    Type    Process name          GPU
Memory |
|      ID    ID                 |          Usage
|
|=====+=====+=====+
=====|
|    0   N/A   N/A     2812553      C+G     RHEL01
15168MiB |
+-----+-----+
+-----+
[root@localhost:~]

```

9. Using vCenter, [configure](#) the graphics device settings to “Shared Direct.”



10. Make sure that secure boot is disabled for the RedHat VM.
11. Make sure that the VM Boot Options firmware is set to EFI ([ref](#)).

Edit Settings
RHEL01

Virtual Hardware
VM Options

> General Options	VM Name: RHEL01
> VMware Remote Console Options	<input type="checkbox"/> Lock the guest operating system when the last remote user disconnects
> Encryption	Expand for encryption settings
> Power management	Expand for power management settings
> VMware Tools	Expand for VMware Tools settings
> Boot Options	
Firmware	EFI (recommended) ▼
Secure Boot	<input type="checkbox"/> Enabled
Boot Delay	When powering on or resetting, delay boot order by 0 milliseconds
Force EFI setup	<input type="checkbox"/> During the next boot, force entry into the EFI setup screen
Failed Boot Recovery	<input type="checkbox"/> If the VM fails to find boot device, automatically retry after 10 seconds
> Advanced	Expand for advanced settings
> Fibre Channel NPIV	Expand for Fibre Channel NPIV settings

CANCEL
OK

12. Make sure that the following PARAMS are added to the VM Options advanced Edit Configuration. The value of the `pciPassthru.64bitMMIOSizeGB` parameter depends on the GPU memory and number of GPUs assigned to the VM. For example:
 - a. If a VM is assigned 4 x 32GB V100 GPUs, then this value should be 128.
 - b. If a VM is assigned 4 x 16GB P6 GPUs, then this value should be 64.

Edit Settings
RHEL01

> Boot Options
Expand for boot options

Advanced

Settings

☐ Disable acceleration
☒ Enable logging

Debugging and statistics

Run normally

Swap file location

☒ Default

Use the settings of the cluster or host containing the virtual machine.

☐ Virtual machine directory

Store the swap files in the same directory as the virtual machine.

☐ Datastore specified by host

Store the swap files in the datastore specified by the host to be used for swap files. If not possible, store the swap files in the same directory as the virtual machine. Using a datastore that is not visible to both hosts during vMotion might affect the vMotion performance for the affected virtual machines.

Configuration Parameters

EDIT CONFIGURATION...

Latency Sensitivity

Normal

> Fibre Channel NPIV
Expand for Fibre Channel NPIV settings

Configuration Parameters


⚠

Modify or add configuration parameters as needed for experimental features or as instructed by technical support. Empty values will be removed (supported on ESXi 6.0 and later).

Name	Value
pciPassthru.64bitMMIOSizeGB	64
pciPassthru.use64bitMMIO	TRUE

- When adding vGPUs as a new PCI Device to the virtual machine in vCenter, make sure to select NVIDIA GRID vGPU as the PCI Device type.
- Choose the correct GPU profile that suites the GPU being used, the GPU memory, and the usage purpose: for example, graphics versus compute.

35



15. On the RedHat Linux VM, NVIDIA drivers can be installed by running the following command:

```
[root@genomics1 genomics]# sh NVIDIA-Linux-x86_64-460.73.01-grid.run
```

16. Verify that the correct vGPU profile is being reported by running the following command:

```
[root@genomics1 genomics]# nvidia-smi -query-gpu=gpu_name
-format=csv,noheader -id=0 | sed -e 's/ /-/g'
GRID-P6-16C
[root@genomics1 genomics]#
```

17. After reboot, verify that the correct NVIDIA vGPU are reported along with the driver versions.

```

[root@genomics1 genomics]# nvidia-smi
Wed Aug 18 20:30:56 2021
+-----+
+-----+
| NVIDIA-SMI 460.73.01      Driver Version: 460.73.01      CUDA Version:
11.2      |
|-----+-----+
+-----+
| GPU  Name                Persistence-M| Bus-Id        Disp.A | Volatile
Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|      Memory-Usage | GPU-Util
Compute M. |
|
MIG M. |
|=====+=====+=====
=====|
|   0  GRID P6-16C          On    | 00000000:02:02.0 Off |
N/A |
| N/A   N/A    P8      N/A /  N/A |   2205MiB / 16384MiB |      0%
Default |
|
N/A |
+-----+-----+
+-----+
+-----+
+-----+
+-----+
| Processes:
|
| GPU    GI    CI          PID    Type    Process name                  GPU
Memory |
|          ID    ID                                   Usage
|
|=====+=====+=====
=====|
|   0    N/A   N/A         8604      G    /usr/libexec/Xorg
13MiB |
+-----+-----+
+-----+
[root@genomics1 genomics]#

```

18. Make sure that the license server IP is configured on the VM in the vGPU grid configuration file.

a. Copy the template.


```
[root@genomics1 genomics]# cp /etc/nvidia/gridd.conf.template  
/etc/nvidia/gridd.conf
```

- b. Edit the file `/etc/nvidia/rid.conf`, add the license server IP address, and set the feature type to 1.

```
ServerAddress=192.168.169.10
```

```
FeatureType=1
```

19. After restarting the VM, you should see an entry under Licensed Clients in the license server as shown below.



20. Refer to the Solutions Setup section for more information on downloading the GATK and Cromwell software.
21. After GATK can use GPUs on premises, the workflow description language *.wdl has the runtime attributes as shown below.

```

task ValidateBAM {
  input {
    # Command parameters
    File input_bam
    String output_basename
    String? validation_mode
    String gatk_path
    # Runtime parameters
    String docker
    Int machine_mem_gb = 4
    Int additional_disk_space_gb = 50
  }
  Int disk_size = ceil(size(input_bam, "GB")) + additional_disk_space_gb
  String output_name = "${output_basename}_${validation_mode}.txt"
  command {
    ${gatk_path} \
      ValidateSamFile \
      --INPUT ${input_bam} \
      --OUTPUT ${output_name} \
      --MODE ${default="SUMMARY" validation_mode}
  }
  runtime {
    gpuCount: 1
    gpuType: "nvidia-tesla-p6"
    docker: docker
    memory: machine_mem_gb + " GB"
    disks: "local-disk " + disk_size + " HDD"
  }
  output {
    File validation_report = "${output_name}"
  }
}

```

[Next: Conclusion.](#)

Conclusion

[Previous: GPU setup.](#)

Many healthcare organizations around the world have standardized on FlexPod as a common platform. With FlexPod, you can deploy healthcare capabilities with confidence. FlexPod with NetApp ONTAP comes standard with the ability to implement an industry leading set of protocols out of the box. Irrespective of the origin of the request to run genomics of a given patient, interoperability, accessibility, availability, and scalability come standard with a FlexPod platform. When standardized on a FlexPod platform, the culture of innovation becomes contagious.

Where to find 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

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_480ml_aiml_deployment.pdf

- FlexPod Datacenter with VMware vSphere 7.0 and NetApp ONTAP 9.7

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/fp_vmware_vsphere_7_0_ontap_9_7.html

- ONTAP 9 Documentation Center

<http://docs.netapp.com>

- Epic honor roll

<https://www.netapp.com/blog/achieving-epic-honor-roll/>

- Agile and efficient—how FlexPod drives data center modernization

<https://www.flexpod.com/idc-white-paper/>

- AI in healthcare

<https://www.netapp.com/us/media/na-369.pdf>

- FlexPod for healthcare Ease Your Transformation

<https://flexpod.com/solutions/verticals/healthcare/>

- FlexPod from Cisco and NetApp

<https://flexpod.com/>

- AI and Analytics for healthcare (NetApp)

<https://www.netapp.com/us/artificial-intelligence/healthcare-ai-analytics/index.aspx>

- AI in healthcare Smart infrastructure Choices Increase Success

<https://www.netapp.com/pdf.html?item=/media/7410-wp-7314.pdf>

- FlexPod Datacenter with ONTAP 9.8, ONTAP Storage Connector for Cisco Intersight, and Cisco Intersight Managed Mode.

<https://www.netapp.com/pdf.html?item=/media/25001-tr-4883.pdf>

- FlexPod Datacenter with Red Hat Enterprise Linux OpenStack Platform

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_openstack_osp6.html

Version history

Version	Date	Document version history
Version 1.0	November 2021	Initial release.

FlexPod Datacenter for Epic Directional Sizing Guide

FlexPod for Epic Directional Sizing Guide

Brian O'Mahony, Ganesh Kamath, Atul Bhalodia, NetApp
Mike Brennan, Jon Ebmeier, Cisco

In partnership with:



Purpose

This technical report provides guidance for sizing FlexPod (NetApp storage and Cisco Unified Computing System) for an Epic Electronic Health Record (EHR) application software environment.

FlexPod systems that host Epic Hyperspace, InterSystems Caché database, Cogito Clarity analytics and reporting suite, and services servers hosting the Epic application layer provide an integrated platform for a dependable, high-performance infrastructure that can be deployed rapidly. The FlexPod integrated platform is deployed by skilled FlexPod channel partners and is supported by Cisco and NetApp technical assistance centers.

The sizing exercise described in this document covers users, global reference counts, availability, and disaster recovery (DR) requirements. The goal is to determine the optimal size of compute, network, and storage infrastructure components.

This document is outlined into the following main sections:

- [Reference Architecture](#), which describes the small, medium, and large compute storage architectures that can be used to host the Epic production database workload.
- [Technical Specifications](#), which details a sample bill of materials for the storage architectures. The configurations that are described are only for general guidance. Always size the systems according to your workload and tune the configurations as necessary.

Overall solution benefits

By running an Epic environment on the FlexPod architectural foundation, healthcare organizations can expect to see improved staff productivity and decreased capital and operating expenses. FlexPod, a prevalidated, rigorously tested converged infrastructure from the strategic partnership of Cisco and NetApp, is engineered and designed specifically to deliver predictable low-latency system performance and high availability. This approach results in high comfort levels and the best response time for users of the Epic EHR system.

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

- **Modular architecture.** FlexPod addresses the varied needs of the Epic modular architecture with purpose-configured FlexPod platforms for each specific workload. All components are connected through a clustered server and storage management fabric and a cohesive management toolset.
- **Accelerated application deployment.** The prevalidated architecture reduces implementation integration time and risk to expedite Epic project plans. NetApp OnCommand Workforce Automation (WFA) workflows for Epic automate Epic backup and refresh and remove the need for custom unsupported scripts. Whether the solution is used for an initial rollout of Epic, a hardware refresh, or expansion, more resources can be shifted to the business value of the project.
- **Simplified operations and lowered costs.** Eliminate the expense and complexity of legacy proprietary RISC and UNIX platforms by replacing them with a more efficient and scalable shared resource capable of supporting clinicians wherever they are. This solution delivers higher resource utilization for greater ROI.
- **Quicker deployment of infrastructure.** Whether it's in an existing data center or a remote location, the integrated and tested design of FlexPod Datacenter with Epic enables customers to have the new infrastructure up and running in less time with less effort.
- **Scale-out architecture.** Scale SAN and NAS from terabytes to tens of petabytes without reconfiguring running applications.
- **Nondisruptive operations.** Perform storage maintenance, hardware lifecycle operations, and software upgrades without interrupting the business.
- **Secure multitenancy.** Supports the increased needs of virtualized server and storage shared infrastructure, enabling secure multitenancy of facility-specific information, especially when hosting multiple instances of databases and software.
- **Pooled resource optimization.** Help reduce physical server and storage controller counts, load balance workload demands, and boost utilization while improving performance.
- **Quality of service (QoS).** FlexPod offers QoS on the entire stack. Industry-leading QoS 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.** The industry-leading workflow automation, orchestration, and management tools offered by FlexPod systems allow IT to be far more responsive to business requests. These requests can range from Epic backup and provisioning of additional test and training environments to analytics database replications for population health-management initiatives.
- **Productivity.** Quickly deploy and scale this solution for optimal clinician end-user experience.
- **Data Fabric.** The NetApp Data Fabric architecture weaves data together across sites, beyond physical boundaries, and across applications. The Data Fabric is built for data-driven enterprises in a data-centric world. Data is created and used in multiple locations, and it often needs to be leveraged and shared with other locations, applications, and infrastructures. Customers want a way to manage data that is consistent and integrated. The Data Fabric offers a way to manage data that puts IT in control and simplifies ever-increasing IT complexity.

Scope

This document covers environments that use Cisco Unified Computing System (Cisco UCS) and NetApp ONTAP based storage. It provides sample reference architectures for hosting Epic.

It does not cover:

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

- Sizing for nonproduction workloads

Audience

This document is for NetApp and partner systems engineers and professional services personnel. The reader is assumed to have a good understanding of compute and storage sizing concepts, as well as technical familiarity with Cisco UCS and NetApp storage systems.

Related documents

The following technical reports are relevant to this technical report. Together they make up a complete set of documents required for sizing, designing, and deploying Epic on FlexPod infrastructure:

- [TR-4693: FlexPod Datacenter for Epic EHR Deployment Guide](#)
- [TR-3930i: NetApp Sizing Guidelines for Epic](#) (requires Field Portal access)
- [TR-3928: NetApp Best Practices for Epic](#)

Reference architecture

NetApp storage reference architectures for Epic

An appropriate storage architecture can be determined by the overall database size and 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.

Given the storage requirements for Epic software environments, NetApp has three reference architectures based on the size of the environment. Epic requires the use of NetApp sizing methods to properly size a NetApp storage system for use in Epic environments. For quantitative performance requirements and sizing guidance, see NetApp [TR-3930i: NetApp Sizing Guidelines for Epic](#). NetApp Field Portal access is required to view this document.

The architectures listed here are a starting point for the design. The workloads must be validated in the SPM tool for the number of disks and controller utilization. Work with the NetApp Epic team to validate all designs.

All Epic production is deployed on all-flash arrays. In this report, the disk pools required for spinning disk have been consolidated to three disk pools for all-flash arrays. Before reading this section, you should review the Epic All-Flash Reference Architecture Strategy Handbook for the Epic storage layout requirements.

The three storage reference architectures are as follows:

- **Small.** Four-node architecture with two nodes in production and two nodes in DR (fewer than 5M global references)
- **Medium.** Six-node architecture with four nodes in production and two nodes in DR (more than 5M global references)
- **Large.** Twelve-or-more node architecture with six to ten nodes in production (5M-10M global references)



Global references = (Read IOPS + (Write Operations per 80-Second Write Burst / 45)) * 225.
These numbers are taken from the customer-specific Epic Hardware Configuration Guide.

Storage layout and LUN configuration

The first step in satisfying Epic's high-availability (HA_ and redundancy requirements is to design the storage

layout specifically for the Epic software environment. The design considerations should include isolating disk pool 1 from disk pool 2 on dedicated high-performance storage. See the Epic All-Flash Reference Architecture Strategy Handbook for information about what workloads are in each disk pool.

Placing each disk pool on a separate node creates the fault domains required for the isolation of Epic's production and nonproduction workloads. Using one aggregate per node maximizes disk utilization and aggregate affinity to provide better performance. This design also maximizes storage efficiency with aggregate-level deduplication.

Because Epic allows storage resources to be shared for nonproduction needs, a storage system can often service both the Clarity server and production services storage needs, such as virtual desktop infrastructure (VDI), CIFS, and other enterprise functions.

The Epic Database Storage Layout Recommendations document provides recommendations for the size and number of LUNs for each database. These recommendations might need to be adjusted according to your environment. It is important to review these recommendations with Epic support and finalize the number of LUNs and LUN sizes.



NetApp recommends starting with larger size LUNs because the size of the LUNs themselves have no cost to storage. For ease of operation, make sure that the number of LUNs and initial size can grow well beyond expected requirements after 3 years. Growing LUNs is much easier to manage than adding LUNs while scaling. With thin-provisioned LUNs and volumes, the storage-used space shows up in the aggregate.

Epic requires database, journal, and application or system storage to be presented to database servers as LUNs through FC.

Use one LUN per volume for Epic production and for Clarity. For larger deployments, NetApp recommends 24 to 32 LUNs for the Epic database. Factors that determine the number of LUNs to use are:

- Overall size of the Epic DB after 3 years. For larger DBs, determine the maximum size of the LUN for that operating system (OS) and make sure that you have enough LUNs to scale. For example, if you need a 60TB Epic database and the OS LUNs have a 4TB maximum, you need 24 to 32 LUNs to provide scale and head room.

Regardless of whether the architecture is small, medium, or large:

- ONTAP allows easy nondisruptive scale up and scale out. Disks and nodes can be upgraded, added, or removed by using ONTAP nondisruptive operations. Customers can start with four nodes and move to six nodes or upgrade to larger controllers nondisruptively.
- NetApp OnCommand Workflow Automation workflows can back up and refresh Epic full-copy test environments. This solution simplifies the architecture and saves on storage capacity with integrated efficiencies.
- The DR shadow database server is part of a customer's business continuity strategy (used to support SRO functionality and potentially configured to be an SRW instance). Therefore, the placement and sizing of the third storage system are usually the same as in the production database storage system.
- Database consistency requires some consideration. If NetApp SnapMirror backup copies are used in relation to business continuity, see the Epic document Business Continuity Technical Solutions Guide. For information about the use of SnapMirror technologies, see [TR-3446: SnapMirror Async Overview and Best Practices Guide](#).
- Isolation of production from potential bully workloads is a key design objective of Epic. A storage pool is a fault domain in which workload performance must be isolated and protected. Each node in an ONTAP cluster is a fault domain and can be considered as a pool of storage.

All platforms in the ONTAP family can run the full host of feature sets for Epic workloads.

Small configuration: four-node reference architecture for fewer than 5M global references (up to ~22K total IOPS)

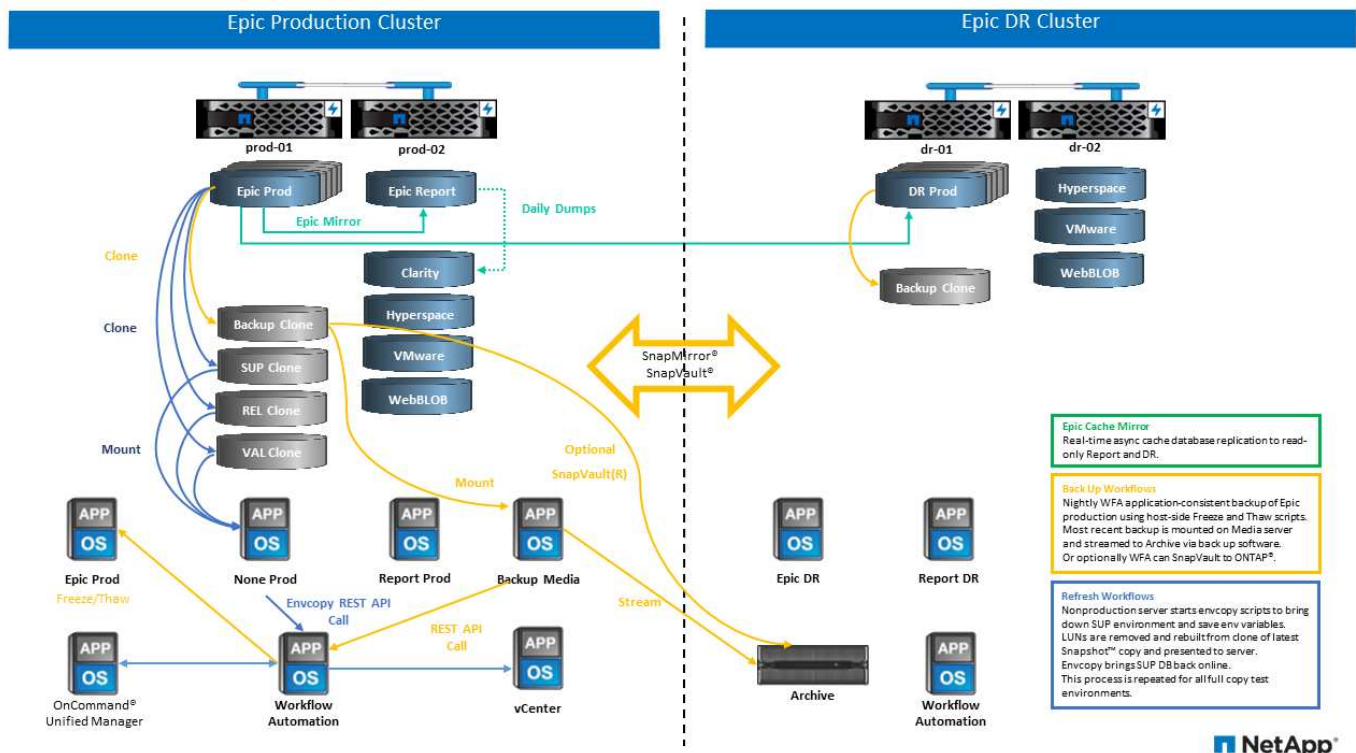
The small reference architecture is a four-node architecture with two nodes in production and two nodes in DR, with fewer than 5M global references. This architecture can be used by customers with fewer than 5M global references. At this size, the separation of Report and Clarity is not required.

With unique multiprotocol support from NetApp, QoS, and the ability to create fault domains in the same cluster, you can run all the production workload for disk pool1 and disk pool2 on a single HA pair and meet all of NetApp best practices and Epic's High Comfort rating requirements. All of disk pool1 would run on node1 and all of disk pool 2 would run on pool2.


With the ability of ONTAP to segregate workloads in the same cluster, and ONTAP multiprotocol support, all the production Epic workloads (Production, Report, Clarity, VMware, Citrix, CIFS, and Epic- related workloads) can be run on a single HA pair in a single cluster. This capability enables you to meet all of Epic's requirements (documented in the Epic All-Flash Reference Architecture Strategy Handbook) and all the NetApp best practices. Basically, pool1 runs on node prod-01 and pool2 runs on prod-02, as shown in the figure below. The NAS 1 workload can be placed on node 2 with NetApp multiprotocol NAS and SAN capabilities.

For disaster recovery, Epic DR pool 3 is split between the two nodes in the HA pair. Epic DR runs on node dr-01 and DR services run on dr-02.

NetApp SnapMirror or SnapVault replication can be set up as needed for workloads.



From a storage design and layout perspective, the following figure shows a high-level storage layout for the production database and the other constructs that comprise the Epic workload.



Medium configuration: six-node reference architecture for greater than 5M global references (22K-50K total IOPS)

The medium reference architecture is a six-node architecture with four nodes in production and two nodes in DR, with 5M-10M global references.

For this size, the All-Flash Reference Architecture Strategy Handbook states that you need to separate Epic Report workloads from Clarity, and that you need at least four nodes in production.

The six-node architecture is the most commonly deployed architecture in Epic environments. Customers with more than 5,000,000 global references are required to place Report and Clarity in separate fault domains. See the Epic All-Flash Reference Architecture Strategy Handbook.

Customers with fewer than 5,000,000 global references can opt to go with six nodes rather than four nodes for the following key advantages:

- Offload backup archive process from production
- Offload all test environments from production

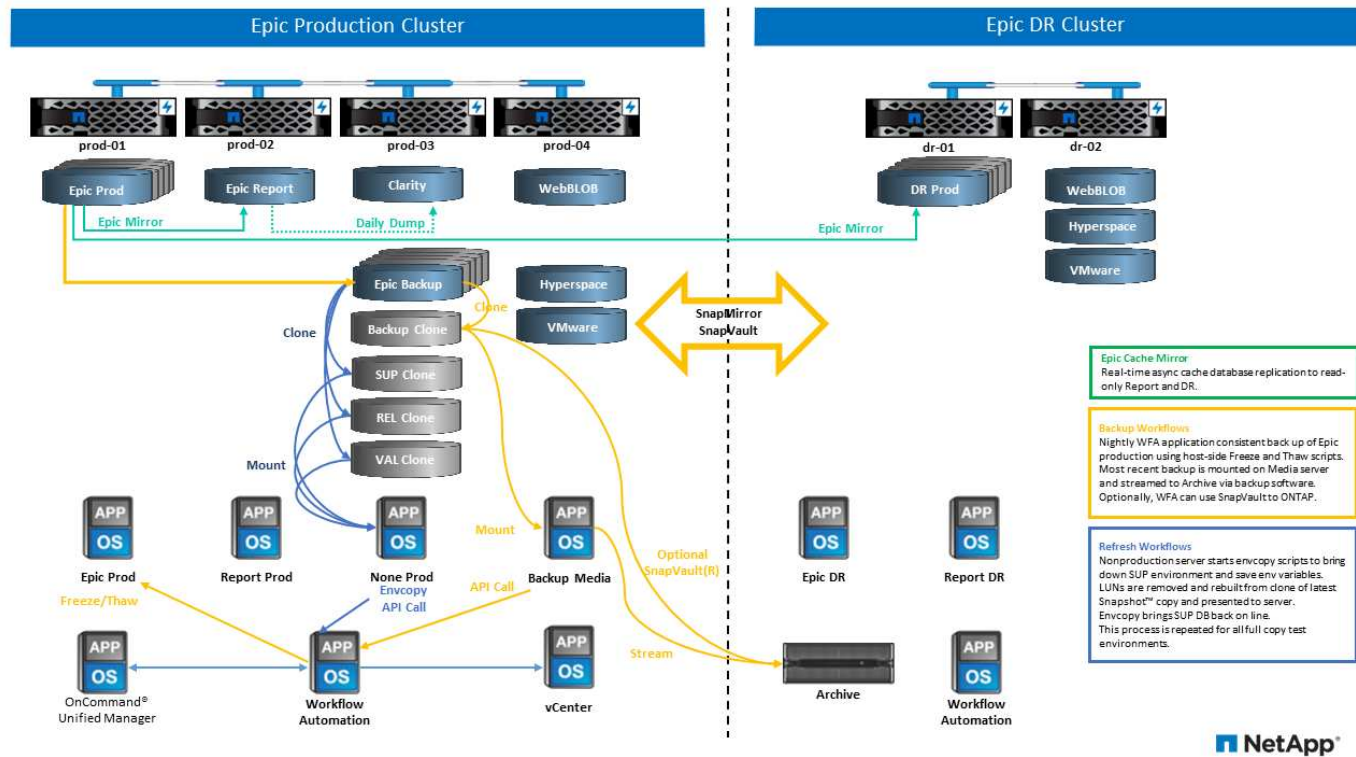
Production runs on node prod-01. Report runs on node prod-02, which is an up-to-the-minute Epic mirror copy of production. Test environments like support, release, and release validation can be cloned from either Epic production, Report, or DR. The figure below shows clones made from production for full-copy test environments.

The second HA pair is used for production services storage requirements. These workloads include storage for Clarity database servers (SQL or Oracle), VMware, Hyperspace, and CIFS. Customers might have non-Epic workloads that could be added to nodes 3 and node 4 in this architecture, or preferably added to a separate HA pair in the same cluster.

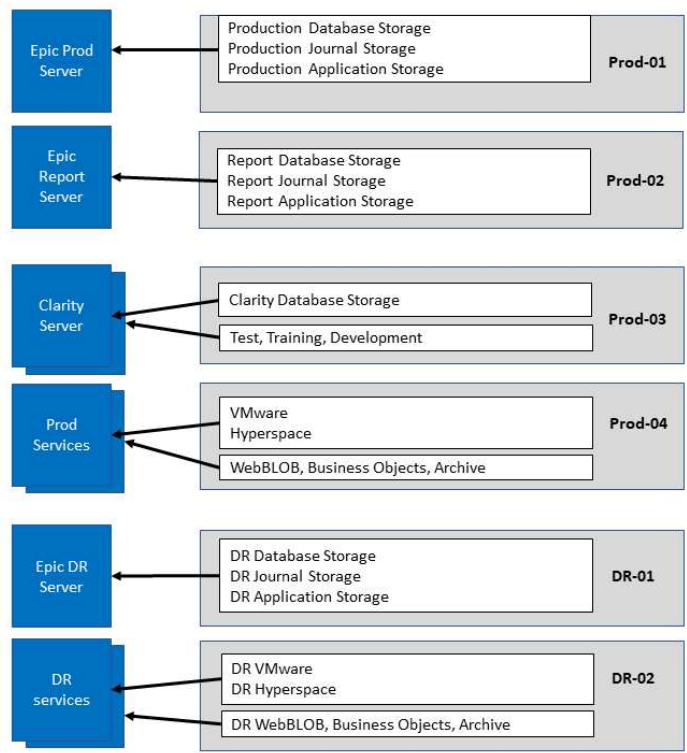
SnapMirror technology is used for storage-level replication of the production database to the second HA pair. SnapMirror backup copies can be used to create NetApp FlexClone volumes on the second storage system for nonproduction environments such as support, release, and release validation. Storage-level replicas of the

production database can also support customers' implementation of their DR strategy.

Optionally, to be more storage efficient, full-test clones can be created from the Report NetApp Snapshot copy backup and run directly on node 2. In this design, a SnapMirror destination copy is not required to be saved on disk.



The following figure shows the storage layout for a six-node architecture.



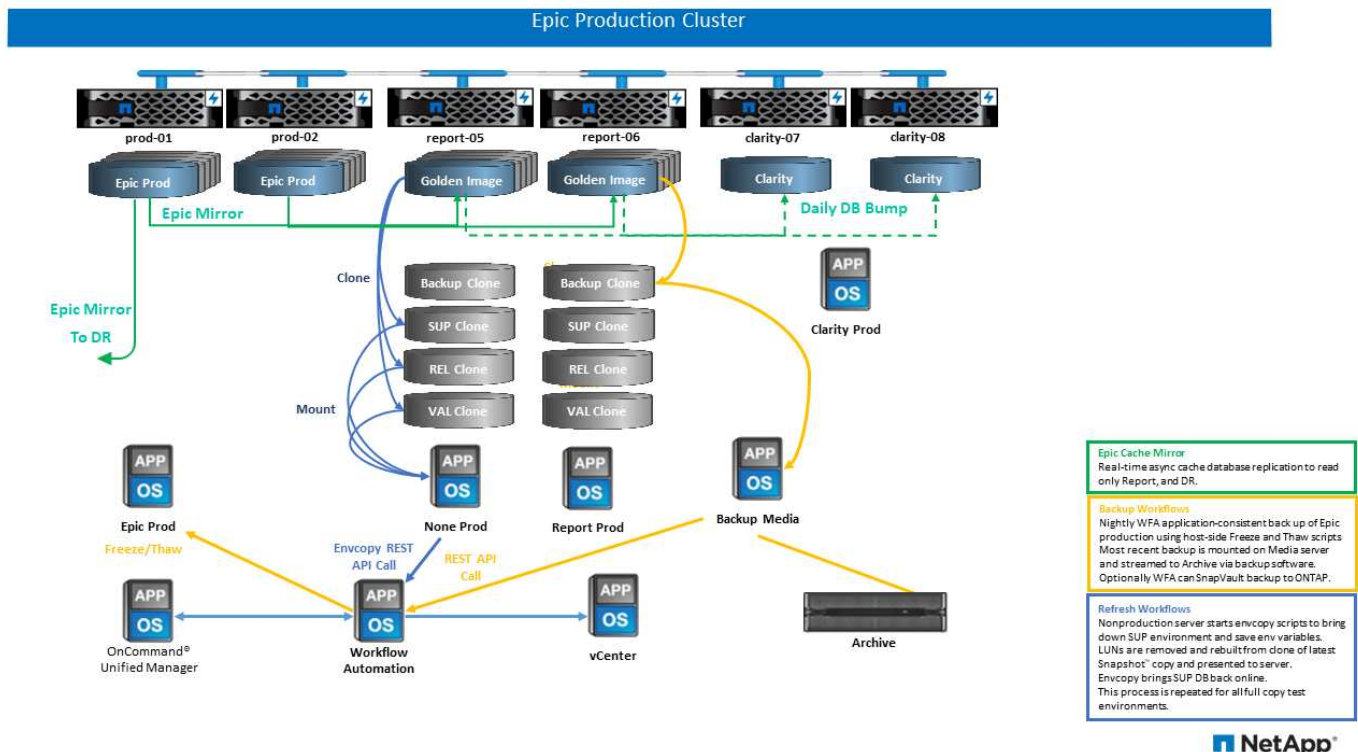
Large configuration: reference architecture for greater than 10M global references (more than 50K IOPS)

The large architecture is typically a twelve-or-more-node architecture with six to ten nodes in production, with more than 10M global references. For large Epic deployments, Epic Production, Epic Report, and Clarity can be placed on a dedicated HA pair with storage evenly balanced among the nodes, as shown in the figure below.

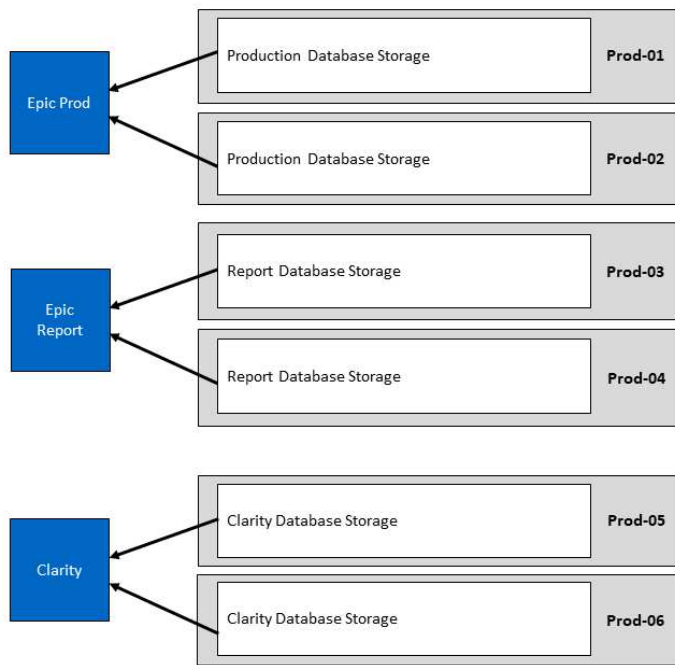
Larger customers have two options:

- Retain the six-node architecture and use AFF A700 controllers.
- Run Epic production, report, and DR on a dedicated AFF A300 HA pair.

You must use the SPM to compare controller utilization. Also, consider rack space and power when selecting controllers.



The following figure shows the storage layout for a large reference architecture.



Cisco UCS reference architecture for Epic

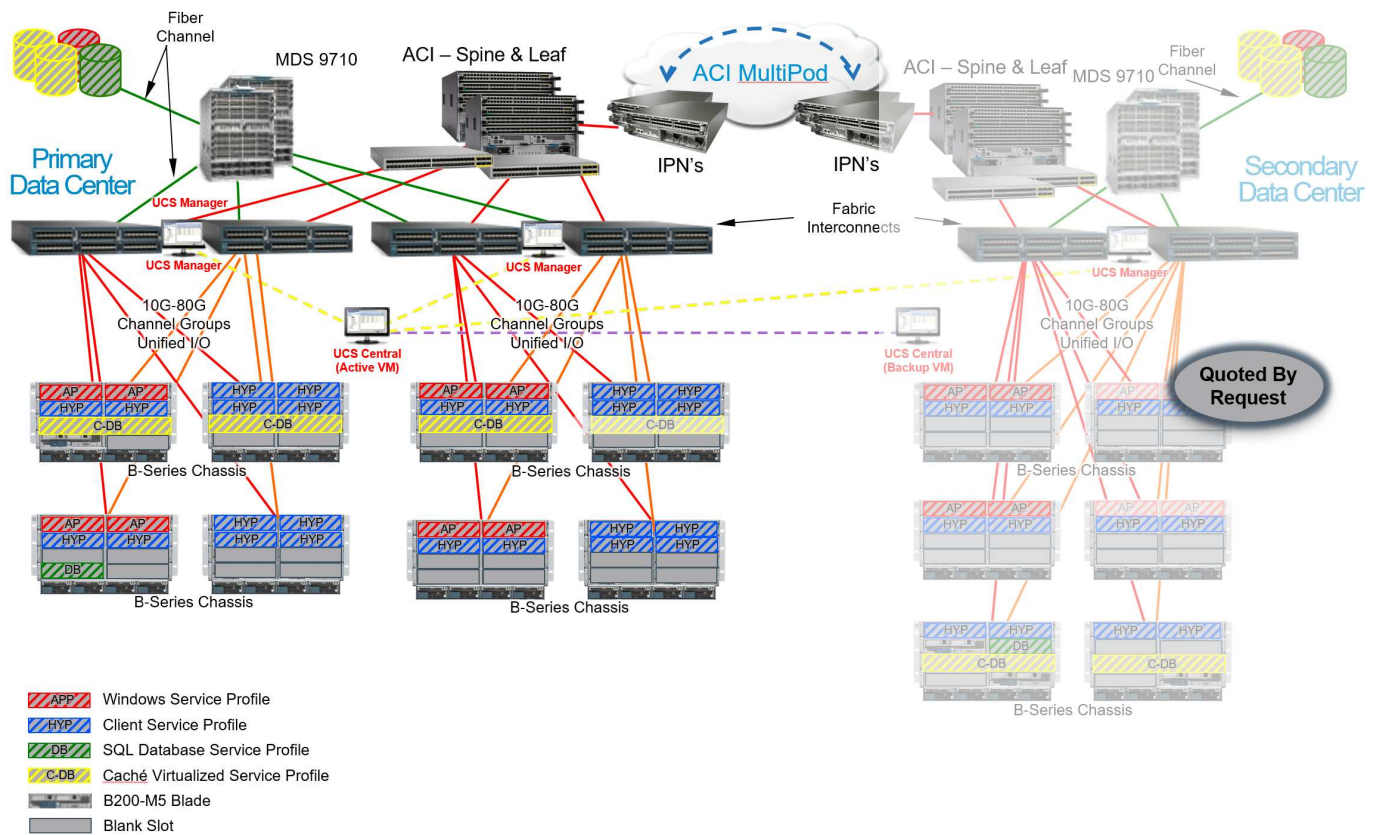
The architecture for Epic on FlexPod is based both on guidance from Epic, Cisco, and NetApp, and from partner experience in working with Epic customers of all sizes. The architecture is adaptable and applies best practices for Epic, depending on the customer's data center strategy—whether small or large,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

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

Basic design for smaller implementations

A basic design for Epic on Cisco UCS is less extensive than an expanded design. An example of a basic design use case might be a physician's practice with outpatient clinics. Such an organization might have few users of the Epic applications, or it might not need all components of Epic. For example, a physician's practice group might not require the Epic Willow Pharmacy application or Epic Monitor for in-patient monitoring. A basic design requires fewer virtual hosts and fewer physical servers. It is also likely to have fewer SAN requirements, and the WAN connections to the secondary data center might be handled with basic routing and TCP/IP.

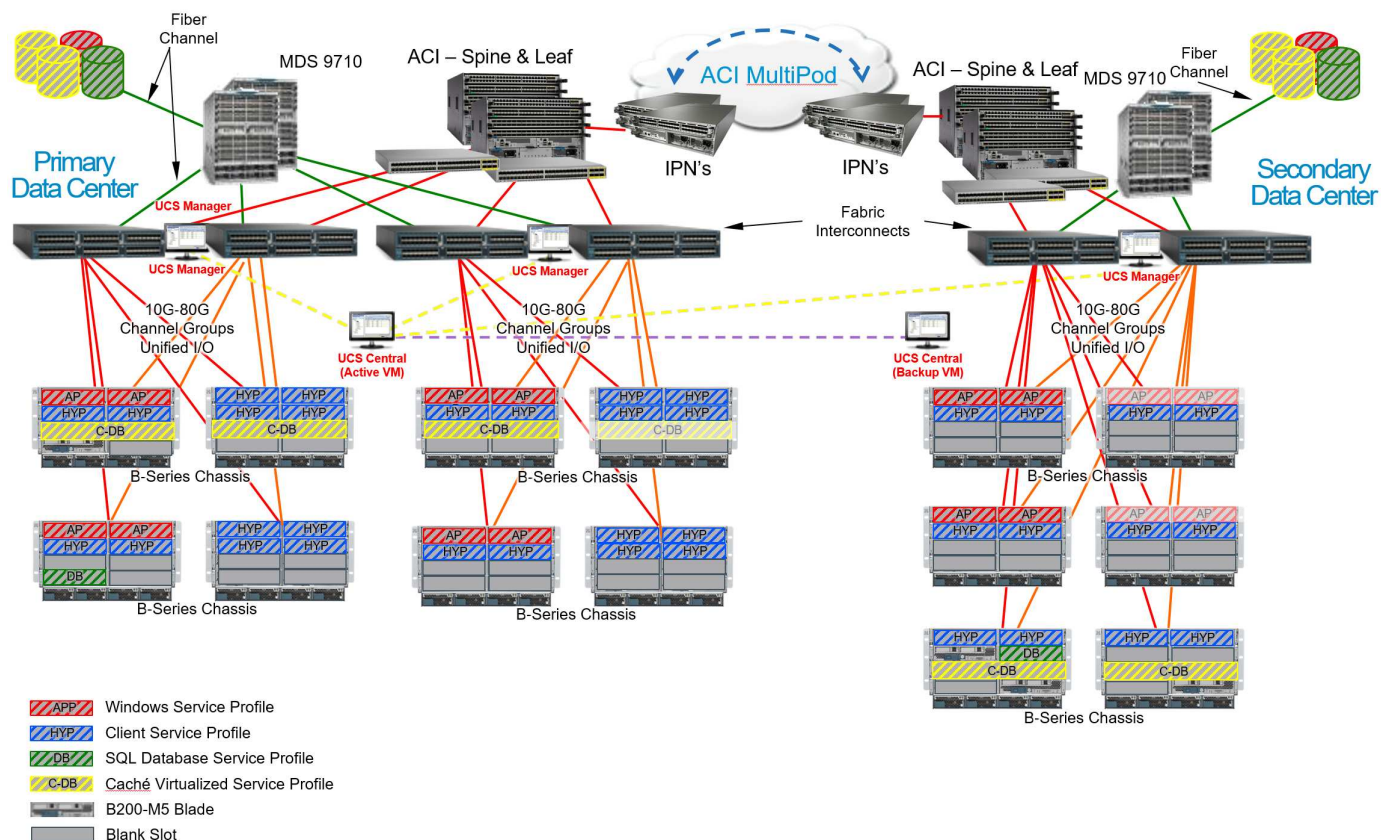
The following figure illustrates an example of a basic small Epic configuration.



Hyperspace active-active implementations

In the secondary data center, to avoid unused hardware resources and software costs, customers might use an active-active design for Epic Hyperspace. This design enables optimizing computing investment by delivering Hyperspace from both the primary data center and the secondary data center.

The Hyperspace active-active design, an example of which is shown in the following figure, takes the expanded design one step further and puts XenApp for Hyperspace or other Hyperspace virtual application servers into full operation in the secondary data center.



Technical specifications for small, medium, and large architectures

The FlexPod design enables a flexible infrastructure that encompasses many different components and software versions. Use [TR-4036: FlexPod Technical Specifications](#) as a guide for building or assembling a valid FlexPod configuration. The configurations that are detailed are only the minimum requirements for FlexPod, and they are just a sample. They can be expanded in the included product families as required for different environments and use cases.

The following table lists the capacity configurations for the Epic production database workload. The total capacity listed accommodates the need for all Epic components.

	Small	Medium	Large
Platform	One AFF A300 HA pair	One AFF A300 HA pair	One AFF A300 HA pair
Disk shelves	24 x 3.8TB	48 x 3.8TB	96 x 3.8TB
Epic database size	3 to 20TB	20TB-40TB	>40TB
Total IOPS	22,000	50,000	125,000
Raw	92.16TB	184.32TB	368.64TB
Usable capacity	65.02TiB	134.36TiB	269.51TiB
Effective capacity (2:1 storage efficiency)	130.04TiB	268.71TiB	539.03TiB

Epic production workloads can be easily satisfied with a single AFF A300 HA pair. An AFF A300 HA pair can push upward of 200k IOPs, which satisfies a large Epic deployment with room for more shared workloads.

Some customer environments might have multiple Epic production workloads running simultaneously, or they might simply have higher IOP requirements. In that case, work with the NetApp account team to size the storage systems according to the required IOPs and capacity and arrive at the right platform to serve the workloads. There are customers running multiple Epic environments on an AFF A700 HA pair.

The following table lists the standard software required for the small, medium, and large configurations.

Software	Product family	Version or release
Storage	Data ONTAP	ONTAP 9.3 GA
Network	Cisco UCS-FI	Cisco UCS Manager 3.2(2f)
	Cisco Ethernet switches	7.0(3)I7(2)
	Cisco FC: Cisco MDS 9132T	8.2(2)
Hypervisor	Hypervisor	VMware vSphere ESXi 6.5 U1
	VMs	RHEL 7.4
Management	Hypervisor management system	VMware vCenter Server 6.5 U1 (VCSA)
	NetApp Virtual Storage Console	VSC 7.0P1
	SnapCenter	SnapCenter 4.0
	Cisco UCS Manager	3.2(2f) or later

The following table lists small configuration infrastructure components.

Layer	Product family	Quantity and model	Details
Compute	Cisco UCS 5108 Chassis	Two	Based on the number of blades required to support the users
	Cisco UCS blade servers	4 x B200 M5	Each with 2 x 18 cores, 2.7GHz, and 384GB BIOS 3.2(2f)
	Cisco UCS VIC	4 x UCS 1340	VMware ESXi fNIC FC driver: 1.6.0.34 VMware ESXi eNIC Ethernet driver: 1.0.6.0 (see the matrix)
	2 x Cisco UCS FI	6332-16UP with Cisco UCS Manager 3.2 (2f)	
Network	Cisco Ethernet switches		2 x Cisco Nexus 93180YC-FX
Storage network	IP network N9k for BLOB storage		FI and UCS chassis
	FC: Cisco MDS 9132T		Two Cisco 9132T switches

Layer	Product family	Quantity and model	Details
Storage	NetApp AFF A300	1 HA pair	1 x 2-node cluster
	DS224C disk shelf	1 DS224C disk shelf (fully populated with 24 drives)	One fully populated disk shelf
	SSD	24 x 3.8TB	

A single disk shelf of 3.8TB SSD drives should suffice for most smaller Epic customer deployments. However, for shared workloads, more disk capacity might be required. You must size for your capacity accordingly.

The following table lists the medium configuration infrastructure components.

Layer	Product family	Quantity and model	Details
Compute	Cisco UCS 5108 Chassis	Four	Based on the number of blades required to support the users
	Cisco UCS blade servers	4 x B200 M5	Each with 2 x 18 cores, 2.7GHz/3.0Ghz, and 384GB 4 sockets for Cache DB BIOS 3.2(2f)
	Cisco UCS VIC	4 x UCS 1340	VMware ESXi fNIC FC driver: 1.6.0.34 VMware ESXi eNIC Ethernet driver: 1.0.6.0 (see the matrix)
	2 x Cisco UCS FI	6332-16UP with Cisco UCS Manager 3.2(2f)	
Network	Cisco Ethernet switches		2 x Cisco Nexus 93180YC-FX
Storage network	IP network: Cisco N9k for BLOB storage		FI and Cisco UCS chassis
	FC: Cisco MDS 9132T		Two Cisco 9132T switches
Storage	NetApp AFF A300	2 HA pairs	2 x 2-node cluster for all Epic workloads (Production, Report, Clarity, VMware, Citrix, CIFS, and so on)
	DS224C disk shelf	2 x DS224C disk shelves	2 fully populated disk shelves
	SSD	48 x 3.8TB	

Four disk shelves of 3.8TB SSD drives should suffice for almost all medium Epic customer deployments. However, assess your disk capacity requirements and size for required capacity accordingly.

The following table lists the large configuration infrastructure components.

Layer	Product family	Quantity and model	Details
Compute	Cisco UCS 5108 Chassis	8	
	Cisco UCS blade servers	4 x B200 M5	Each with 2 x 24 cores, 2.7GHz, and 576GB BIOS 3.2(2f)
	Cisco UCS VIC	4 x UCS 1340	VMware ESXi fNIC FC driver: 1.6.0.34 VMware ESXi eNIC Ethernet driver: 1.0.6.0 (see the matrix)
	2 x Cisco UCS FI	6332-16UP with Cisco UCS Manager 3.2(2f)	
Network	Cisco Ethernet switches		2 x Cisco Nexus 93180YC-FX
Storage network	IP network: Cisco N9k for BLOB storage		
	FC: Cisco MDS 9706		Two Cisco 9706 switches
Storage	NetApp AFF A300	3 HA pairs	3 x 2-node cluster for Epic workloads (Prod, Report, Clarity, VMware, Citrix, CIFS, and so on)
	DS224C disk shelf	4 x DS224C disk shelves	4 fully populated disk shelves
	SSD	96 x 3.8TB	

Some customer environments might have multiple Epic production workloads running simultaneously, or they might simply have higher IOPS requirements. In such cases, work with the NetApp account team to size the storage systems according to the required IOPS and capacity and determine the right platform to serve the workloads. There are customers running multiple Epic environments on an AFF A700 HA pair.

Additional information

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

- FlexPod Datacenter with FC Cisco Validated Design. Detailed deployment of FlexPod Datacenter environment.

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

- TR-3928: NetApp Best Practices for Epic. Overview of Epic software environments, reference architectures, and integration best practices guidance.

<https://fieldportal.netapp.com/?oparams=68646>

- TR-3930i: NetApp Sizing Guidelines for Epic (access to Field Portal is required to view this document)

<https://fieldportal.netapp.com/?oparams=68786>

- Epic_on_Cisco_UCS_tech_brief. Cisco Best practices with Epic on Cisco UCS.

https://www.cisco.com/c/dam/en_us/solutions/industries/healthcare/Epic_on_UCS_tech_brief_FNL.pdf

- NetApp FlexPod Design Zone

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

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

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

- TR-4693: FlexPod Datacenter for Epic EHR Deployment Guide

<https://www.netapp.com/us/media/tr-4693.pdf>

- NetApp Product Documentation

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

Acknowledgements

The following people contributed to the writing of this guide.

- Ganesh Kamath, Technical Marketing Engineer, NetApp
- Atul Bhalodia, Technical Marketing Engineer, NetApp
- Ketan Mota, Product Manager, NetApp
- Jon Ebmeier, Cisco Systems, Inc.
- Mike Brennan, Cisco Systems, Inc.

FlexPod Datacenter for Epic EHR Deployment Guide

TR-4693: FlexPod Datacenter for Epic EHR Deployment Guide

Brian O'Mahony, NetApp
Ganesh Kamath, NetApp
Mike Brennan, Cisco

In partnership with:



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

FlexPod systems deployed to host Epic HyperSpace, InterSystems Caché database, Cogito Clarity analytics

and reporting suite, and services servers hosting the Epic application layer provide an integrated platform for a dependable, high-performance infrastructure that can be deployed rapidly. The FlexPod integrated platform is deployed by skilled FlexPod channel partners and is supported by Cisco and NetApp technical assistance centers.

Overall solution benefits

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

- **Simplified operations and lowered costs.** Eliminate the expense and complexity of legacy proprietary RISC/UNIX platforms by replacing them with a more efficient and scalable shared resource capable of supporting clinicians wherever they are. This solution delivers higher resource utilization for greater ROI.
- **Quicker deployment of infrastructure.** Whether it's in an existing data center or a remote location, the integrated and tested design of FlexPod Datacenter with Epic enables customers to have the new infrastructure up and running in less time with less effort.
- **Scale-out architecture.** Scale SAN and NAS from terabytes to tens of petabytes without reconfiguring running applications.
- **Nondisruptive operations.** Perform storage maintenance, hardware lifecycle operations, and software upgrades without interrupting the business.
- **Secure multitenancy.** This benefit supports the increased needs of virtualized server and storage shared infrastructure, enabling secure multitenancy of facility-specific information, particularly if hosting multiple instances of databases and software.
- **Pooled resource optimization.** This benefit can help reduce physical server and storage controller counts, load balance workload demands, and boost utilization while improving performance.
- **Quality of service (QoS).** FlexPod offers QoS on the entire stack. Industry-leading QoS storage policies enable differentiated service levels in a shared environment. These policies enable optimal performance for workloads and help in isolating and controlling runaway applications.
- **Storage efficiency.** Reduce storage costs with the NetApp 7: 1 storage efficiency guarantee.
- **Agility.** The industry-leading workflow automation, orchestration, and management tools offered by FlexPod systems allow IT to be far more responsive to business requests. These business requests can range from Epic backup and provisioning of additional test and training environments to analytics database replications for population health management initiatives.
- **Productivity.** Quickly deploy and scale this solution for optimal clinician end- user experiences.
- **Data Fabric.** The NetApp Data Fabric architecture weaves data together across sites, beyond physical boundaries, and across applications. The NetApp Data Fabric is built for data-driven enterprises in a data-centric world. Data is created and used in multiple locations, and it often needs to be leveraged and shared with other locations, applications, and infrastructures. Customers want a way to manage data that is consistent and integrated. It provides a way to manage data that puts IT in control and simplifies ever-increasing IT complexity.

FlexPod

A New approach for infrastructure for Epic EHR

Healthcare provider organizations remain under pressure to maximize the benefits of their substantial investments in industry-leading Epic electronic health records (EHRs). For mission-critical applications, when customers design their data centers for Epic solutions, they often identify the following goals for their data center architecture:

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

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

Over the past decade, the Epic infrastructure customarily consisted of proprietary RISC processor- based servers running proprietary versions of UNIX and traditional SAN storage arrays. These server and storage platforms offer little by way of virtualization and can result in prohibitive capital and operating costs, given increasing IT budget constraints.

Epic now supports a production target platform consisting of a Cisco Unified Computing System (Cisco UCS) with Intel Xeon processors, virtualized with VMware ESXi, running Red Hat Enterprise Linux (RHEL). This platform coupled with Epic's High Comfort Level ranking for NetApp storage running ONTAP, a new era of Epic data center optimization has begun.

Value of prevalidated converged infrastructure

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

FlexPod, a prevalidated, rigorously tested converged infrastructure from the strategic partnership of Cisco and NetApp, is engineered and designed specifically for delivering predictable low-latency system performance and high availability. This approach results in Epic high comfort levels and ultimately the best response time for users of the Epic EHR system.

The FlexPod solution from Cisco and NetApp meets Epic system requirements with a high performing, modular, prevalidated, converged, virtualized, efficient, scalable, and cost-effective platform. It provides:

- **Modular architecture.** FlexPod addresses the varied needs of the Epic modular architecture with purpose-configured FlexPod platforms for each specific workload. All components are connected through a clustered server and storage management fabric and a cohesive management toolset.
- **Accelerated application deployment.** The prevalidated architecture reduces implementation integration time and risk to expedite Epic project plans. NetApp OnCommand Workforce Automation (OnCommand WFA) workflows for Epic automate Epic backup and refresh and remove the need for custom unsupported scripts. Whether the solution is used for an initial rollout of Epic, a hardware refresh, or expansion, more resources can be shifted to the business value of the project.
- **Industry-leading technology at each level of the converged stack.** Cisco, NetApp, VMware, and Red Hat are all ranked as number 1 or number 2 by industry analysts in their respective categories of servers, networking, storage, and open systems Linux.
- **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 some of Epic's largest customer organizations.


Proven FlexPod architecture and cooperative support

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

- **Performance to meet the Epic workload requirements.** Depending on the reference workload requirements (small, medium, large), different ONTAP platforms can be deployed to meet the required I/O profile.
- **Scalability to easily accommodate clinical data growth.** Dynamically scale virtual machines (VMs), servers, and storage capacity on demand, without traditional limits.
- **Enhanced efficiency.** Reduce both administration time and TCO with a converged virtualized infrastructure, which is easier to manage and stores data more efficiently while driving more performance from Epic software. NetApp OnCommand WFA automation simplifies the solution to reduce test environment refresh time from hours or days to minutes.
- **Reduced risk.** Minimize business disruption with a prevalidated platform built on a defined architecture that eliminates deployment guesswork and accommodates ongoing workload optimization.
- **FlexPod Cooperative Support.** NetApp and Cisco have established Cooperative Support, a strong, scalable, and flexible support model to address the unique support requirements of the FlexPod converged infrastructure. This model uses the combined experience, resources, and technical support expertise of NetApp and Cisco to provide a streamlined process for identifying and resolving a customer's FlexPod support issue, regardless of where the problem resides. The FlexPod Cooperative Support model helps to make sure that your FlexPod system operates efficiently and benefits from the most up-to-date technology, while providing an experienced team to help resolve integration issues.

FlexPod Cooperative Support is especially valuable to healthcare organizations running business-critical applications such as Epic 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 Epic solution delivers specific benefits for Epic EHR workflows.

Cisco Unified Computing System

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

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

The system is designed as a single virtual blade chassis that incorporates and scales across multiple blade chassis, rack servers, and racks. The system implements a radically simplified architecture that eliminates the multiple redundant devices that populate traditional blade server chassis and result in layers of complexity: Ethernet and FC switches and chassis management modules. Cisco UCS consists of a redundant pair of Cisco 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 include critical server information about the server identity such as LAN and SAN addressing, I/O configurations, firmware versions, boot order, network VLAN, physical port, and QoS

policies. Service profiles can be dynamically created and associated with any physical server in the system in minutes rather than hours or days. The association of service profiles with physical servers is performed as a simple, single operation and enables migration of identities between servers in the environment without requiring any physical configuration changes. It facilitates rapid bare-metal provisioning of replacements for failed servers.

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

Cisco UCS has been extensively tested with Epic over a multi- year period to demonstrate that it meets the server configuration requirements. Cisco UCS is a supported server platform, as listed in customers' "Epic Hardware Configuration Guide."

Cisco Nexus

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

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

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

- **Infrastructure scalability.** Virtualization, efficient power and cooling, cloud scale with automation, high density, and performance all support efficient data center growth.
- **Operational continuity.** The design integrates hardware, NX-OS software features, and management to support zero-downtime environments.
- **Transport flexibility.** Incrementally adopt new networking technologies with a cost-effective solution.

Together, Cisco UCS with Cisco Nexus switches and MDS multilayer directors provide a compute, networking, and SAN connectivity solution for Epic.

NetApp ONTAP

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

The scale-out architecture of the ONTAP software can flexibly adapt to various I/O workloads. To deliver the necessary throughput and low latency required for clinical applications while providing a modular scale-out architecture, all-flash configurations are typically used in ONTAP architectures. All- flash arrays will be required by Epic by year 2020 and are required by Epic today for customers with more than 5 million global references.

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. Customers can clone, replicate, and back up the Epic environment (from expensive SSD storage) to more economical HDD storage on other nodes, meeting or exceeding Epic guidelines for SAN-based cloning and backup of production disk pools. With NetApp cloud-enabled storage and Data Fabric, you can back up to object storage on the premises or in the cloud.

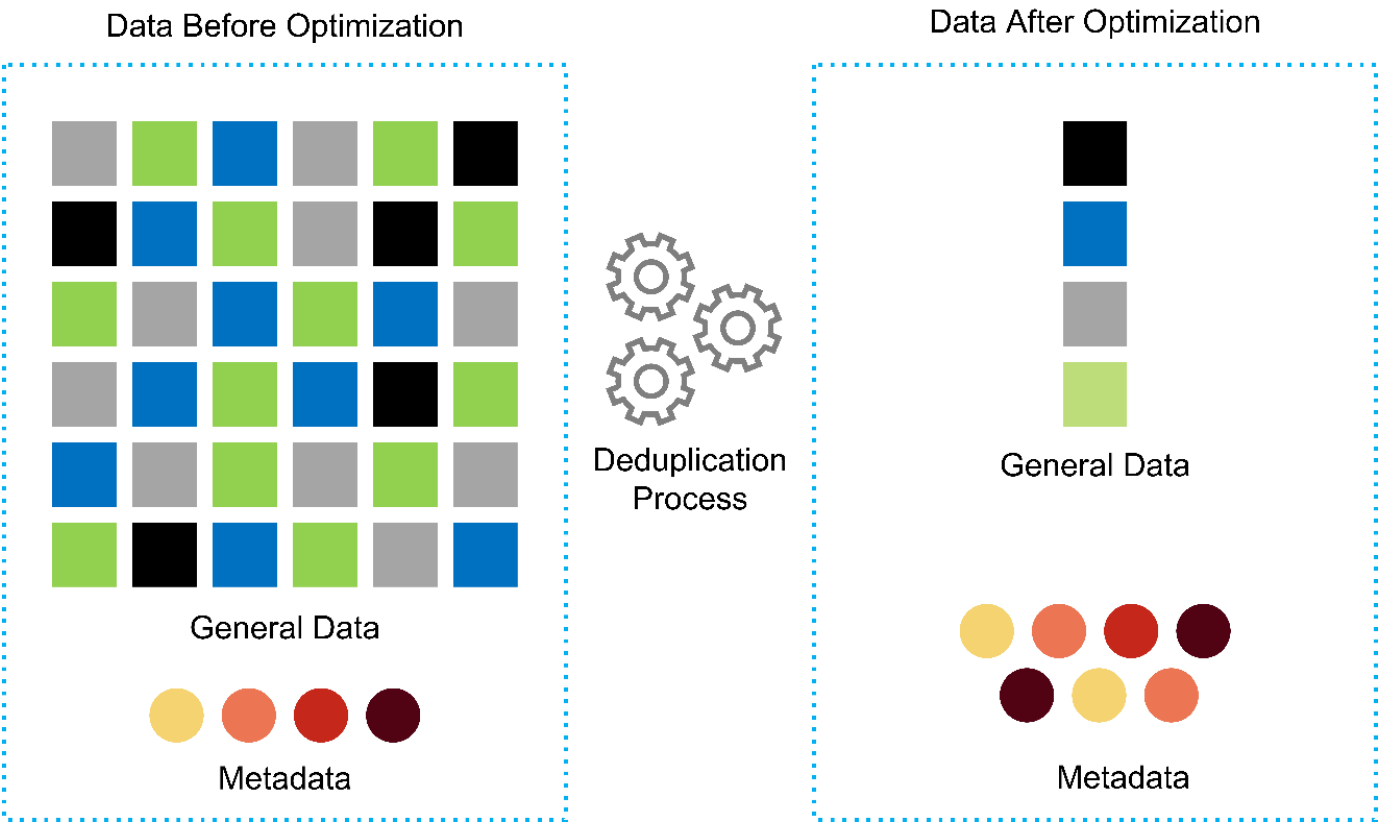
ONTAP offers features that are extremely useful in Epic environments, simplifying management, increasing availability and automation, and reducing the total amount of storage needed:

- **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 FAS product families. This innovative combination of all-flash media with ONTAP delivers the consistent low latency and high IOPS of all-flash storage with the industry-leading ONTAP software.
- **Storage efficiency.** Reduce total capacity requirements with deduplication, NetApp FlexClone, inline compression, inline compaction, thin replication, thin provisioning, and aggregate deduplication.

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

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

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

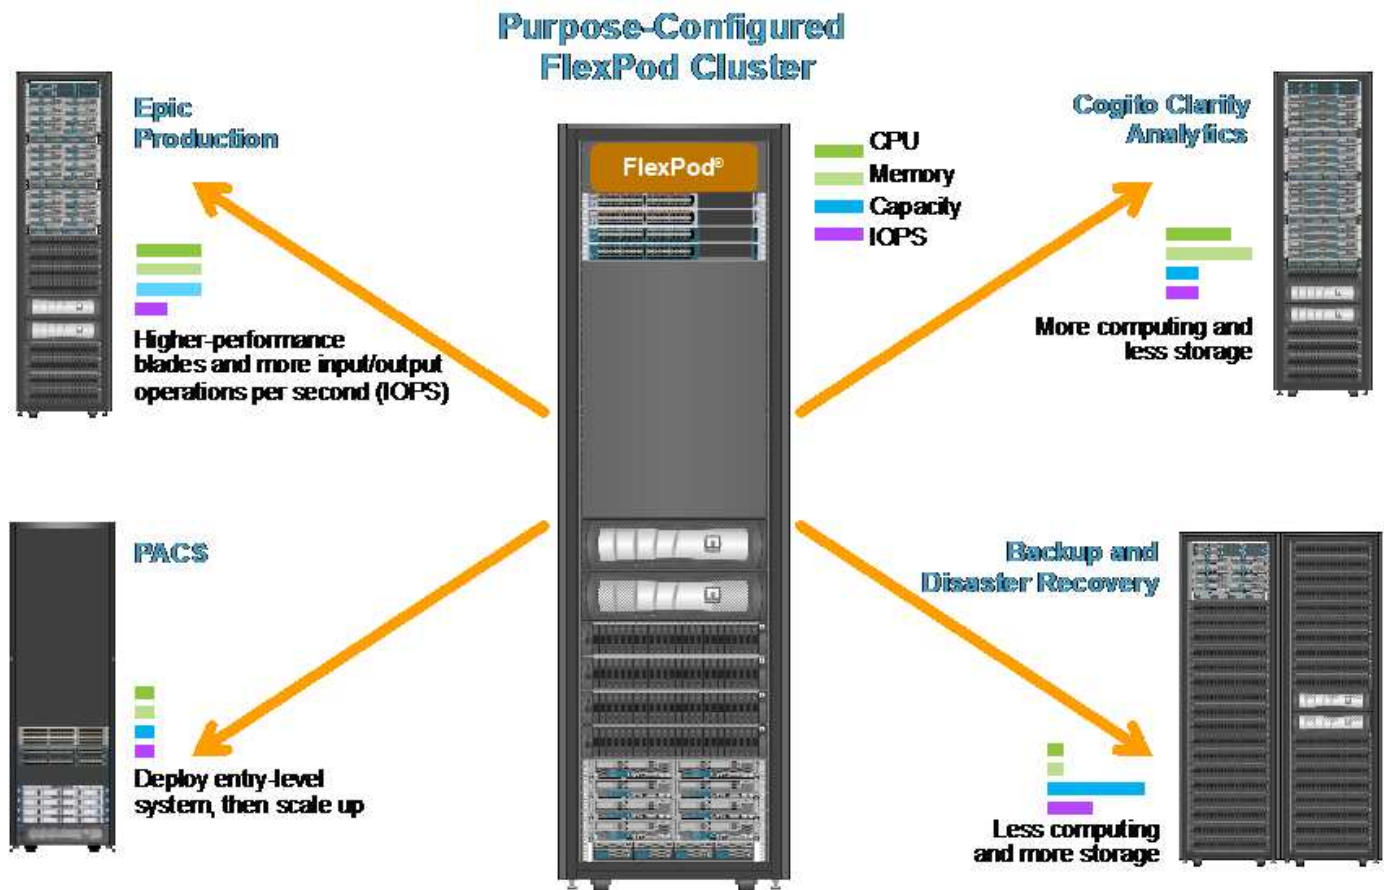


- **Space-efficient cloning.** The FlexClone capability allows you to almost instantly create clones to support

backup and test environment refresh. These clones consume additional storage only as changes are made.

- **Integrated data protection.** Full data protection and disaster recovery features help customers protect critical data assets and provide disaster recovery.
- **Nondisruptive operations.** Upgrading and maintenance can be performed without taking data offline.
- **Epic workflow automation.** NetApp has designed OnCommand WFA workflows to automate and simplify the Epic backup solution and refresh of test environments such as SUP, REL, and REL VAL. This approach eliminates the need for any custom unsupported scripts, reducing deployment time, operations hours, and disk capacity required for NetApp and Epic best practices.
- **QoS.** Storage QoS allows you to limit potential bully workloads. More importantly, QoS can guarantee minimum performance for critical workloads such as Epic production. NetApp QoS can reduce performance-related issues by limiting contention.
- **OnCommand Insight Epic dashboard.** The Epic Pulse tool can identify an application issue and its effect on the end user. The OnCommand Insight Epic dashboard can help identify the root cause of the issue and gives full visibility into the complete infrastructure stack.
- **Data Fabric.** NetApp Data Fabric simplifies and integrates data management across cloud and on-premises to accelerate digital transformation. It delivers consistent and integrated data management services and applications for data visibility and insights, data access and control, and data protection and security. NetApp is integrated with AWS, Azure, Google Public Cloud, and IBM Cloud clouds, giving customers a wide breadth of choice.

The following figure illustrates FlexPod for Epic workloads.



Epic overview

Overview

Epic is a software company headquartered in Verona, Wisconsin. The following excerpt from the company's website describes the span of functions supported by Epic software:

"Epic makes software for midsize and large medical groups, hospitals, and integrated healthcare organizations—working with customers that include community hospitals, academic facilities, children's organizations, safety net providers, and multi-hospital systems. Our integrated software spans clinical, access, and revenue functions and extends into the home. "


It is beyond the scope of this document to cover the wide span of functions supported by Epic software. From the storage system point of view, however, for each deployment, all Epic software shares a single patient-centric database. Epic uses the InterSystems Caché database, which is available for various operating systems, including IBM AIX and Linux.

The primary focus of this document is to enable the FlexPod stack (servers and storage) to satisfy performance-driven requirements for the InterSystems Caché database used in an Epic software environment. Generally, dedicated storage resources are provided for the production database, whereas shadow database instances share secondary storage resources with other Epic software-related components, such as Clarity reporting tools. Other software environment storage, such as that used for application and system files, is also provided by the secondary storage resources.

Purpose-built for specific Epic workloads

Though Epic does not resell server, network, or storage hardware, hypervisors, or operating systems, the company has specific requirements for each component of the infrastructure stack. Therefore, Cisco and NetApp worked together to test and enable FlexPod Datacenter to be successfully configured, deployed, and supported to meet customers' Epic production environment requirements. This testing, technical documentation, and growing number of successful mutual customers have resulted in Epic expressing an increasingly high level of comfort in FlexPod Datacenter's ability to meet Epic customers' needs. See the "Epic Storage Products and Technology Status" document and the "Epic Hardware Configuration Guide. "

The end-to-end Epic reference architecture is not monolithic, but modular. The figure below outlines five distinct modules, each with unique workload characteristics.



These interconnected but distinct modules have often resulted in Epic customers having to purchase and manage specialty silos of storage and servers. These might include a vendor's platform for traditional tier 1 SAN; a different platform for NAS file services; platforms specific to protocol requirements of FC, FCoE, iSCSI, NFS, and SMB/CIFS; separate platforms for flash storage; and appliances and tools to attempt to manage these silos as virtual storage pools.

With FlexPod connected through ONTAP, you can implement purpose-built nodes optimized for each targeted workload, achieving the economies of scale and streamlined operational management of a consistent compute, network, and storage data center.

Caché production database

Caché, manufactured by InterSystems, is the database system on which Epic is built. All patient data in Epic is stored in a Caché database.

In an InterSystems Caché database, the data server is the access point for persistently stored data. The application server services database queries and makes data requests to the data server. For most Epic software environments, the use of the symmetric multiprocessor architecture in a single database server suffices to service the Epic applications' database requests. In large deployments, using InterSystems' Enterprise Caché Protocol can support a distributed database model.

By using failover-enabled clustered hardware, a standby data server can access the same disks (that is, storage) as the primary data server and take over the processing responsibilities in the event of a hardware failure.

InterSystems also provides technologies to satisfy shadow, disaster recovery, and high-availability (HA) requirements. InterSystems' shadow technology can be used to asynchronously replicate a Caché database from a primary data server to one or more secondary data servers.

Cogito Clarity

Cogito Clarity is Epic's integrated analytics and reporting suite. Starting as a copy of the production Caché database, Cogito Clarity delivers information that can help improve patient care, analyze clinical performance, manage revenue, and measure compliance. As an OLAP environment, Cogito Clarity utilizes either Microsoft SQL Server or Oracle RDBMS. Because this environment is distinct from the Caché production database environment, it is important to architect a FlexPod platform that supports the Cogito Clarity requirements following Cisco and NetApp published validated design guides for SQL Server and Oracle environments.

Epic Hyperspace Desktop Services

Hyperspace is the presentation component of the Epic suite. It reads and writes data from the Caché database and presents it to the user. Most hospital and clinic staff members interact with Epic using the Hyperspace application.

Although Hyperspace can be installed directly on client workstations, many healthcare organizations use application virtualization through a Citrix XenApp farm or a virtual desktop infrastructure (VDI) to deliver applications to users. Virtualizing XenApp server farms using ESXi is supported. See the validated designs for FlexPod for ESXi in the "References" section for configuration and implementation guidelines.

For customers interested in deploying full VDI Citrix XenDesktop or VMware Horizon View systems, careful attention must be paid for an optimal clinical workflow experience. A foundational step for obtaining precise configurations is to clearly understand and document the scope of the project, including detailed mapping of user profiles. Many user profiles include access to applications beyond Epic. Variables in profiles include:

- Authentication, especially Imprivata or similar tap- and-go single sign-on (SSO), for nomadic clinician users
- PACS Image Viewer
- Dictation software and devices such as Dragon NaturallySpeaking
- Document management such as Hyland OnBase or Perceptive Software integration
- Departmental applications such as health information management coding from 3M Health Care or OptumHealth
- Pre-Epic legacy EMR or revenue cycle apps, which the customer might still use
- Video conferencing capabilities that could require use of video acceleration cards in the servers

Your certified FlexPod reseller, with specific certifications in VMware Horizon View or Citrix XenDesktop, will work with your Cisco and NetApp Epic solutions architect and professional services provider to scope and architect the solution for your specific VDI requirements.

Disaster recovery and shadow copies

Evolving to active-active dual data centers

In Epic software environments, a single patient-centric database is deployed. Epic's hardware requirements refer to the physical server hosting the primary Caché data server as the production database server. This server requires dedicated, high-performance storage for files belonging to the primary database instance. For HA, Epic supports the use of a failover database server that has access to the same files.

A reporting shadow database server is typically deployed to provide read-only access to production data. It hosts a Caché data server configured as a backup shadow of the production Caché data server. This database server has the same storage capacity requirements as the production database server. This storage is sized differently from a performance perspective because reporting workload characteristics are different.

A shadow database server can also be deployed to support Epic's read-only (SRO) functionality, in which access is provided to a copy of production in read-only mode. This type of database server can be switched to read-write mode for business continuity reasons.

To meet business continuity and disaster recovery (DR) objectives, a DR shadow database server is commonly deployed at a site geographically separate from the production and/or reporting shadow database servers. A DR shadow database server also hosts a Caché data server configured as a backup shadow of the production Caché data server. It can be configured to act as a shadow read-write instance if the production site is unavailable for an extended time. Like the reporting shadow database server, the storage for its database files has the same capacity requirements as the production database server. In contrast, this storage is sized the same as production from a performance perspective, for business continuity reasons.

For healthcare organizations that need continuous uptime for Epic and have multiple data centers, FlexPod can be used to build an active-active design for Epic deployment. In an active-active scenario, FlexPod hardware is installed into a second data center and is used to provide continuous availability and quick failover or disaster recovery solutions for Epic. The "Epic Hardware Configuration Guide" provided to customers should be shared with Cisco and NetApp to facilitate the design of an active-active architecture that meets Epic's guidelines.

Licensing Caché

NetApp and Cisco are experienced in migrating legacy Epic installations to FlexPod systems following Epic's best practices for platform migration. They can work through any details if a platform migration is required.

One consideration for new customers moving to Epic or existing customers evaluating a hardware and software refresh is the licensing of the Caché database. InterSystems Caché can be purchased with either a platform-specific license (limited to a single hardware OS architecture) or a platform-independent license. A platform-independent license allows the Caché database to be migrated from one architecture to another, but it costs more than a platform-specific license.



Customers with platform-specific licensing might need to budget for additional licensing costs to switch platforms.

Epic storage considerations

RAID performance and protection

Epic recognizes the value of NetApp RAID DP, RAID-TEC, and WAFL technologies in achieving levels of data protection and performance that meet Epic-defined requirements. Furthermore, with NetApp efficiency technologies, NetApp storage systems can deliver the overall read performance required for Epic environments while using fewer disk drives.

Epic requires the use of NetApp sizing methods to properly size a NetApp storage system for use in Epic environments. For more information, see TR-3930i: NetApp Sizing Guidelines for Epic. NetApp Field Portal access is required to view this document.

Isolation of production disk groups

See the Epic All-Flash Reference Architecture Strategy Handbook for details about the storage layout on an all-flash array. In summary, disk pool 1 (production) must be stored on a separate storage fault domain from disk pool 2. An ONTAP node in the same cluster is a fault domain.

Epic recommends the use of flash for all full-size operational databases, not just the production operational databases. At present this approach is only a recommendation; however, by calendar year 2020 it will be a requirement for all customers.

For very large sites, where the production OLTP database is expected to exceed 5 million global references per second, the Cogito workloads should be placed on a third array to minimize the impact to the performance of the production OLTP database. The test bed configuration used in this document is an all-flash array.

High availability and redundancy

Epic recommends the use of HA storage systems to mitigate hardware component failure. This recommendation extends from basic hardware, such as redundant power supplies, to networking, such as multipath networking.

At the storage node level, Epic highlights the use of redundancy to enable nondisruptive upgrades and nondisruptive storage expansion.

Pool 1 storage must reside on separate disks from the pool 2 storage for the performance isolation reasons previously stated, both of which NetApp storage arrays provide by default out of the box. This separation also provides data-level redundancy for disk-level failures.

Storage monitoring

Epic recommends the use of effective monitoring tools to identify or predict any storage system bottlenecks.

NetApp OnCommand Unified Manager, bundled with ONTAP, can be used to monitor capacity, performance, and headroom. For customers with OnCommand Insight, an Insight dashboard has been developed for Epic that gives complete visibility into storage, network, and compute beyond what the Epic Pulse monitoring tool provides. Although Pulse can detect an issue, Insight can identify the issue early, before it has an impact.

Snapshot technology

Epic recognizes that storage node-based NetApp Snapshot technology can minimize performance impacts on production workloads compared to traditional file-based backups. When Snapshot backups are intended for use as a recovery source for the production database, the backup method must be implemented with database consistency in mind.

Storage expansion

Epic cautions against expanding storage without considering storage hotspots. For example, if storage is frequently added in small increments, storage hotspots can develop where data is not evenly spread across

disks.

Comprehensive management tools and automation capabilities

Cisco Unified Computing System with Cisco UCS Manager

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

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

Cisco UCS Manager (UCSM) 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 faster, more successful Epic infrastructure deployment. Cisco and its technology partners such as VMware and Citrix and storage partners IBM, NetApp, and Pure Storage have developed integrated, validated architectures, including predefined converged architecture infrastructure packages such as FlexPod. Cisco virtualization solutions have been tested with VMware vSphere, Linux, Citrix XenDesktop, and XenApp.

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

Virtualization, therefore, significantly increases the need for virtual machine-level awareness of policy and security, especially given the dynamic and fluid nature of virtual machine 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, virtual machine-aware policies and administration, and network security across the LAN and WAN infrastructure.

- **Scalable.** Growth of virtualization solutions is all but inevitable, so a solution must be able to scale, and scale predictably, with that growth. The Cisco virtualization solutions support high virtual machine density (VMs per server), and additional 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 just as easy to deploy dozens of hosts as it is to deploy hundreds.

Cisco UCS servers provide near-linear performance and scale. Cisco UCS implements the patented Cisco Extended Memory Technology to offer large memory footprints with fewer sockets (with scalability to up to

1TB of memory with 2- and 4-socket servers). Using unified fabric technology as a building block, Cisco UCS server aggregate bandwidth can scale to up to 80Gbps per server, and the northbound Cisco UCS fabric interconnect can output 2Tbps at line rate, helping prevent desktop virtualization I/O and memory bottlenecks. Cisco UCS, with its high-performance, low-latency unified fabric-based networking architecture, supports high volumes of virtual desktop traffic, including high-resolution video and communications traffic. In addition, Cisco storage partner NetApp helps to maintain data availability and optimal performance during boot and login storms as part of the Cisco virtualization solutions.

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

VMware vCenter Server

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

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

Virtual Storage Console for VMware vSphere

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

The virtual appliance for VSC, VASA Provider, and SRA integrates smoothly with the VMware vSphere Web Client and enables you to use SSO services. In an environment with multiple 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:

- **Using VSC to deploy and manage storage and configure the ESXi host.** You can use VSC to add credentials, remove credentials, assign credentials, and set up permissions for storage controllers in your VMware environment. In addition, you can manage ESXi servers that are connected to NetApp storage systems. You can set recommended best practice values for host timeouts, NAS, and multipathing for all the hosts with a couple of clicks. You can also view storage details and collect diagnostic information.
- **Using VASA Provider to create storage capability profiles and set alarms.** VASA Provider for ONTAP is registered with VSC as soon as 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 virtual machine disks (VMDKs) and the VMs that are created on virtual datastores.
- **Using 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 Epic service delivery chain. This

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

- **Optimize existing resources.** Identify misused, underused, or orphaned assets using established best practices to avoid problems and meet service levels.
- **Make better decisions.** Real-time data helps resolve capacity problems more quickly to accurately plan future purchases, avoid overspending, and defer capital expenditures.
- **Accelerate IT initiatives.** Better understand virtual environments to manage risks, minimize downtime, and speed cloud deployment.
- **OnCommand Insight dashboard.** This Epic dashboard was developed by NetApp for Epic and provides a comprehensive view of the complete infrastructure stack and goes beyond Pulse monitoring. OnCommand Insight can proactively identify contention issues in compute, network, and storage.

NetApp OnCommand workflow automation

OnCommand WFA is a free software solution that helps to automate storage management tasks, such as provisioning, migration, decommissioning, data protection configurations, and cloning storage. You can use OnCommand WFA to build workflows to complete tasks that are specified by your processes.

A workflow is a repetitive and procedural task that consists of steps, including the following types of tasks:

- Provisioning, migrating, or decommissioning storage for databases or file systems
- Setting up a new virtualization environment, including storage switches and datastores
- Setting up storage for an application as part of an end-to-end orchestration process

Workflows can be built to quickly set up and configure NetApp storage as per recommended best practices for Epic workloads. OnCommand WFA workflows for Epic replace all customer unsupported scripting required for Epic workflows to automate backup and test environment refresh.

NetApp SnapCenter

SnapCenter is a unified, scalable platform for data protection. SnapCenter provides centralized control and oversight, allowing users to manage application-consistent, database-consistent Snapshots copies. SnapCenter enables the backup, restore, clone, and backup, verification of virtual machine (VMs) from both primary and secondary destinations (SnapMirror and SnapVault). With SnapCenter, database, storage, and virtualization administrators have a single tool to manage backup, restore, and clone operations for various applications, databases, and VMs.

SnapCenter enables centralized application resource management and easy data protection job execution by using resource groups and policy management (including scheduling and retention settings). SnapCenter provides unified reporting by using a dashboard, multiple reporting options, job monitoring, and log and event viewers.

SnapCenter can back up VMware, RHEL, SQL, Oracle, and CIFS. Combined with Epic WFA backup workflow integration, NetApp provides a backup solution for any Epic environment.

Design

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

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

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

- **Easy nondisruptive scale up and scale out.** Disks and nodes can be upgraded, added, or removed by using ONTAP nondisruptive operations. Customers can start with four nodes and move to six nodes or upgrade to larger controllers nondisruptively.
- **Storage efficiencies.** Reduce total capacity requirements with deduplication, FlexClone, inline compression, inline compaction, thin replication, thin provisioning, and aggregate deduplication. The FlexClone capability allows you to almost instantly create clones to support backup and test environment refreshes. These clones consume additional storage only as changes are made.
- **Ability of OnCommand WFA workflows to back up and refresh Epic full-copy test environments.** This solution simplifies the architecture and saves on storage capacity with integrated efficiencies. These architectures factor in the backup solution for Epic and leverage storage integration to integrate with any backup solution.
- **DR shadow database server.** The DR shadow database server is part of a customer's business continuity strategy (used to support storage read-only [SRO] functionality and potentially configured to be a storage read-write [SRW] instance). Therefore, the placement and sizing of the third storage system are in most cases the same as in the production database storage system.
- **Database consistency (requires some consideration).** If SnapMirror backup copies are used in relation to business continuity, see the document "Epic Business Continuity Technical Solutions Guide." For information about the use of SnapMirror technologies, see [TR-3446: SnapMirror Async Overview and Best Practices Guide](#).
- **Isolation of production from potential bully workloads is a key design objective of Epic.** A storage pool is a fault domain in which workload performance must be isolated and protected. Each node in an ONTAP cluster is a fault domain and can be considered as a pool of storage.

All platforms in the ONTAP family can run the full host of feature sets for Epic workloads.

Storage architecture

The figure below depicts a 6-node architecture, which is a commonly deployed architecture in Epic environments. There is also a 4- node or 12- node deployment, but these architectures are simply a reference or starting point for the design. The workloads must be validated in the [SPM](#) sizing tool for the number of disks and controller utilization. All Epic production is deployed on AFF arrays. See the Epic All-Flash Reference Architecture Strategy Handbook for Epic storage layout requirements.




Work with the NetApp Epic team to validate all designs. Epic requires the use of NetApp sizing methods to properly size a NetApp storage system for use in Epic environments. For more information, see [TR-3930i: NetApp Sizing Guidelines for Epic](#). NetApp Field Portal access is required to view this document.

about what workloads are in each disk pool.

Placing each disk pool on a separate node creates the fault domains required for Epic isolation of production and nonproduction workloads. Using one aggregate per node maximizes disk utilization and aggregate affinity to provide better performance. This design also maximizes storage efficiency with aggregate-level deduplication.

Because Epic allows storage resources to be shared for nonproduction needs, a storage system can often service both the Clarity server and production services storage needs, such as VDI, CIFS, and other enterprise functions.

The figure below shows the storage layout for the 6-node architecture. Each storage system is a single node in a fully redundant HA pair. This layout ensures maximum utilization on each controller and storage efficiency.



Storage node configuration

High availability

Storage systems configured with nodes in an HA pair mitigate the effect of node failure and enable nondisruptive upgrades of the storage system. Disk shelves connected to nodes with multiple paths increase storage resiliency by protecting against a single-path failure while providing improved performance consistency during a node failover.

Hardware- assisted failover

Hardware-assisted failover minimizes storage node failover time by enabling the remote LAN module or service processor module of one node to notify its partner of a node failure faster than a heartbeat-timeout trigger, reducing the time elapsed before failover. When storage is virtualized, failover times improve because controller identity does not need to move during failover. Only software disk ownership changes.

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 storage systems to call home if a hardware failure or system misconfiguration occurs. For mission-critical environments, NetApp also recommends the SupportEdge Premium package, which provides access to operational expertise, extended support hours, and fast response times on parts replacement.

All-flash optimized personality on AFF A300 and AFF A700 controllers

For the AFF solution to function properly, the environment variable `bootarg.init.flash_optimized` must be set to `true` on both nodes in an HA pair of all-flash-optimized FAS80x0 systems. Platforms with the all-flash-optimized personality support only SSDs.

Volume configuration

Snapshot Copies

A nightly volume-level Snapshot schedule should be set for volumes that provide storage for the production database. Volume-level Snapshot copies can also be used as the source for cloning the production database for use in nonproduction environments such as development, test, and staging. NetApp has developed OnCommand WFA workflows for Epic that automate the backup of production databases and the refresh of test environments. These workflows freeze and thaw the database for application-consistent Snapshot copies. The backup copies of production are automatically presented to test servers for support, release, and release validation. These workflows can also be used for backup streaming and integrity checks.

Snapshot copies can be used to support the restore operations of Epic's production database.

You can use SnapMirror to maintain Snapshot copies on storage systems separate from production.

For SAN volumes, disable the default Snapshot policy on each volume. These Snapshot copies are typically managed by a backup application or by OnCommand WFA workflows. NetApp recommends turning on all efficiency settings to maximize disk utilization.

Volume affinity

To support concurrent processing, ONTAP assesses its available hardware on startup and divides its aggregates and volumes into separate classes, called affinities. In general terms, volumes that belong to one affinity can be serviced in parallel with volumes that are in other affinities. In contrast, two volumes that are in the same affinity often have to take turns waiting for scheduling time (serial processing) on the node's CPU.

The AFF A300 and AFF A700 have a single aggregate affinity and four volume affinities per node. For best node utilization and use of volume affinity, the storage layout should be one aggregate per node and at least four volumes per node. Typically, eight volumes or LUNs are used for an Epic database.

LUN configuration

The document "Epic Database Storage Layout Recommendations" details the size and number of LUNs for each database. It is important for the customer to review that with Epic support and finalize the number of LUNs and LUN sizes; they might need to be adjusted slightly.

Starting with larger size LUNs is recommended because the size of the LUNs themselves has no cost to storage. For ease of operation, make sure that the number of LUNs and initial size can grow well beyond expected requirements after three years. Growing LUNs is much easier to manage than adding LUNs when scaling. With thin provisioning on the LUN and volume, only storage used shows in the aggregate.

Use one LUN per volume for Epic production and for Clarity. For larger deployments, NetApp recommends 24 to 32 LUNs for Epic databases.

Factors that determine the number of LUNs to use are:

- Overall size of the Epic DB after three years. For larger DBs, determine the maximum size of the LUN for that OS and make sure that you have enough LUNs to scale. For example, if you need a 60TB Epic database and the OS LUNs have a 4TB maximum, you would need 24 to 32 LUNs to provide scale and headroom.



Epic requires database, journal, and application or system storage to be presented to database servers as LUNs through FC.

Deployment and configuration

Overview

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

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

This document does not cover:

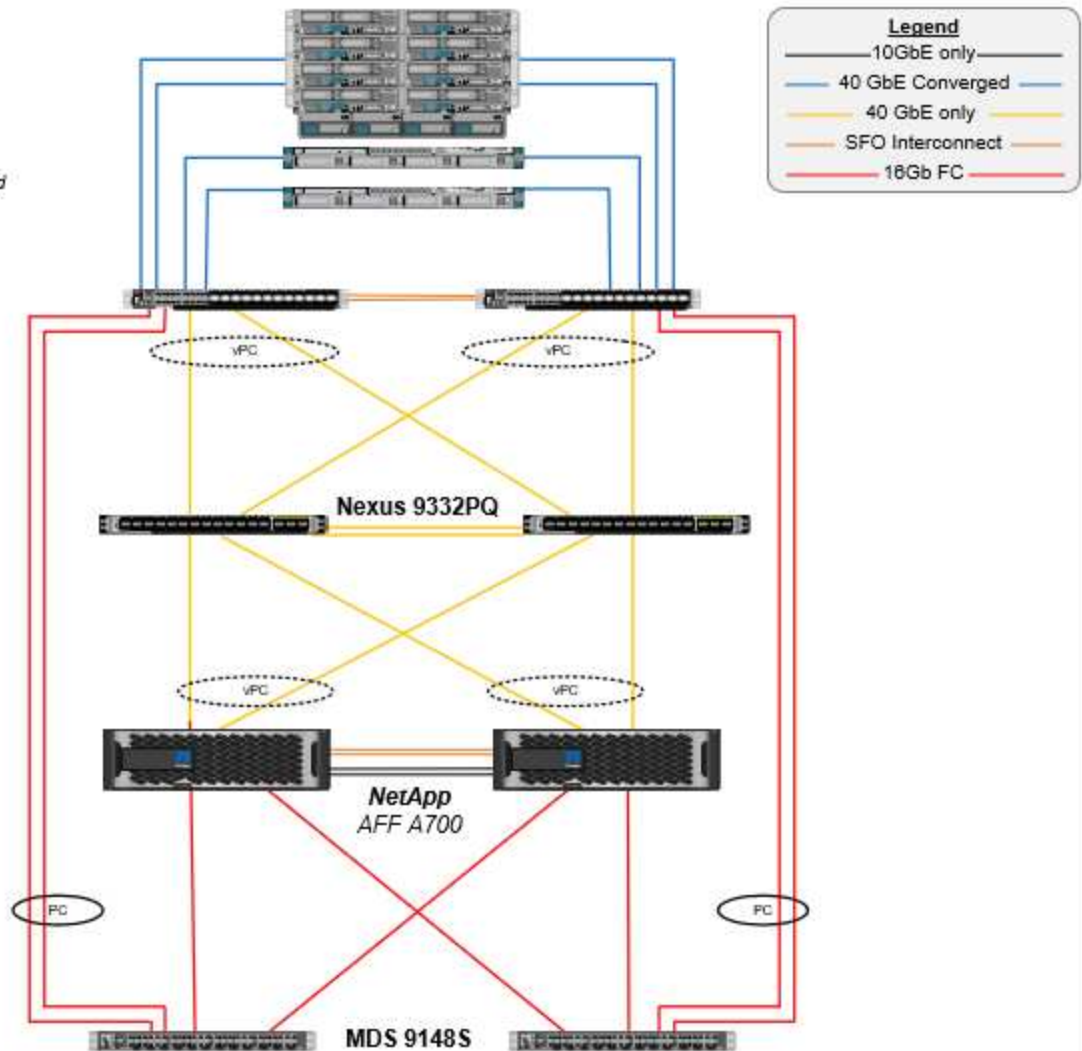
- Detailed deployment of FlexPod Datacenter environment. See [FlexPod Datacenter with FC Cisco Validated Design](#).
- Overview of Epic software environments, reference architectures, and integration best practices guidance. See NetApp [TR-3928: NetApp Best Practices for Epic](#).
- Quantitative performance requirements and sizing guidance. See NetApp [TR-3930: NetApp Sizing Guidelines for Epic](#).
- Use of NetApp SnapMirror technologies to meet backup and disaster recovery requirements.
- Epic database backup and recovery, including SnapCenter.
- Generic NetApp storage deployment guidance.
- Deployment guidance for Clarity reporting environments. See NetApp [TR-4590: Best Practice Guide for Microsoft SQL Server with ONTAP](#).

This section describes the lab environment setup with infrastructure deployment best practices. The GenIO tool is used to simulate the Epic EHR application workload. This section lists the various infrastructure hardware and software components and the versions used.

Cabling diagram

The following figure illustrates the 16Gb FC/40GbE topology diagram for an Epic deployment.

Unified Computing System
 UCS 6332-16UP Fabric Interconnects & UCS B and C Series Servers



Next: Infrastructure Hardware and Software Components.

Infrastructure hardware and software components

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

Layer	Product family	Version or release	Details
Compute	Cisco UCS 5108		One chassis
	Cisco UCS blade servers	4 x B200 M5	Each with 18 CPU cores and 768GB RAM BIOS 2.2(8))
	Cisco UCS VIC	4 x UCS 1340	VMware ESXi fNIC FC driver: 1.6.0.34 VMware ESXi eNIC Ethernet driver: 1.0.6.0

Layer	Product family	Version or release	Details
	2 x Cisco UCS FI	6332-16UP with Cisco UCSM 3.2 (2f)	
Network	Cisco Ethernet switches	7.0(3)I7(2)	2 x Cisco Nexus 9372PX-E
Storage network	iSCSI: IP solution using N9k		FI and UCS chassis
	FC: Cisco MDS 9148S	8.2(2)	Two Cisco 9148S switches
Storage	2 x NetApp AFF A700s	ONTAP 9.3 GA	1 x 2-node cluster
	2 x DS224C disk shelf		
	SSD	48 x 960GB	
Software	Hypervisor	VMware vSphere ESXi 6.5 U1	
	VMs	RHEL 7.4	
	Hypervisor management system	VMware vCenter Server 6.5 U1 (VCSA)	vCenter Server Appliance
	NetApp Virtual Storage Console	VSC 7.0P1	
	SnapCenter	SnapCenter 4.0	
	Cisco UCS Manager	3.2 (2f) *	

[Next: Base Infrastructure Configuration.](#)

Base infrastructure configuration

Network connectivity

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

- Link aggregation using port channels and virtual port channels is used throughout, enabling the design for higher bandwidth and HA.
 - Virtual port channel is used between the Cisco FI and Cisco Nexus switches.
 - Each server has vNICs with redundant connectivity to the unified fabric. NIC failover is used between FI for redundancy.
 - Each server has vHBAs with redundant connectivity to the unified fabric.
- The Cisco UCS FI are configured in end-host mode as recommended, providing dynamic pinning of vNICs to uplink switches.

Storage connectivity

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

- Storage ports ifgroups (vPC)

- 10G link to switch N9k-A
- 10G link to switch N9k-B
- In- band management (active-passive bond):
 - 1G link to management switch N9k-A
 - 1G link to management switch N9k-B
- 16G FC end-to-end connectivity through Cisco MDS switches. Single initiator zoning configured.
- FC SAN boot to fully achieve stateless computing. Servers are booted from LUNs in the boot volume hosted on the AFF storage cluster.
- All Epic workloads are hosted on FC LUNs, which are spread across the storage controller nodes.

Host software

The following software must be installed:

- ESXi is installed on the Cisco UCS blades.
- vCenter is installed and configured, with all the hosts registered in vCenter.
- VSC is installed and registered in vCenter.
- A NetApp cluster is configured.

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

Cisco UCS blade server and switch configuration

The FlexPod for Epic software is designed with fault tolerance at every level. There is no single point of failure in the system. We recommend the use of hot spare blade servers for optimal performance.

This document is intended to provide high-level guidance on the basic configuration of a FlexPod environment for Epic software. In this section, we present high-level steps with some examples to prepare the Cisco UCS compute platform element of the FlexPod configuration. A prerequisite for this guidance is that the FlexPod configuration is racked, powered, and cabled per the instructions in the FlexPod Datacenter with FC Storage.

Cisco Nexus switch configuration

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

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

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

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

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

3. Create the NTP distribution interface, port channels, port channel parameters, and port descriptions for troubleshooting according to the [FlexPod Datacenter with FC Cisco Validated Design](#).

Next: [ESXi Configuration Best Practices](#)

Cisco MDS 9148S configuration

The Cisco MDS 9100 Series FC switches provide redundant 16Gb FC connectivity between the NetApp AFF A700 controllers and the Cisco UCS compute fabric. The cables should be connected as described in the section [“Cabling Diagram.”](#)

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

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

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

configuration mode:

- a. For the fabric A MDS switch:

```
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 using the reference document.

- b. For the fabric B MDS switch:

```
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 using the reference document.

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

- a. For the fabric A MDS switch:

```

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:

```

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

```

Cisco UCS configuration guidance


Cisco UCS allows Epic customers to use their subject matter experts in network, storage, and compute to create policies and templates that tailor the environment to their specific needs. After being created, these

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

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

- Cisco UCS Manager HTML 5 GUI
- Cisco UCS CLI
- Cisco UCS Central for multidomain environments

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 additional fault tolerance at the major Epic functional component level.

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

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

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

and default maintenance.

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

For the detailed steps to configure each key element of the Cisco UCS service profiles for FlexPod, see the [FlexPod Datacenter with FC Cisco Validated Design](#) document.

For Epic deployments, Cisco recommends a range of service profile types, based on the Epic elements being deployed. By using server pools and server pool qualification, customers can identify and automate the deployment of service profiles to particular host roles. A sample list of service profiles are as follows:

- For the Epic Chronicle Caché database hosts:
 - Production host service profile
 - Reporting service host profile
 - Disaster recovery host service profile
 - Hot spare host service profile
- For Epic Hyperspace hosts:
 - VDI host service profile
 - Citrix XenApp host service profile
 - Disaster recovery host service profile
 - Hot spare host service profile
- For the Epic Cogito and Clarity database hosts:
 - Database host service profile (Clarity RDBMS and business objects)
- For the Epic Services hosts:
 - Application host profile (print format and relay, communications, web BLOB, and so on)

ESXi configuration best practices

For the ESXi host-side configuration, see the [InterSystems Best practices for VMware](#). Configure the VMware hosts as you would to run any enterprise database workload:

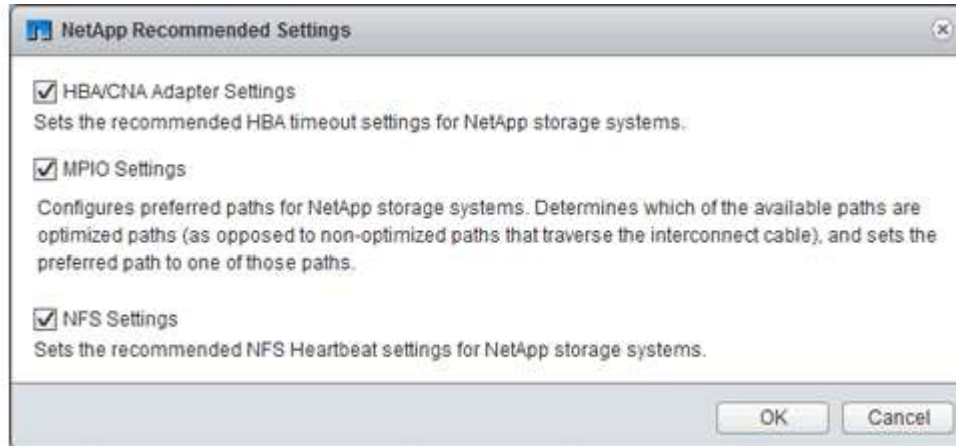
- Virtual Storage Console (VSC) for VMware vSphere checks and sets the ESXi host multipathing settings and HBA timeout settings that work best with NetApp storage systems. The values that VSC sets are based on rigorous internal testing by NetApp.
- For the best storage performance, customers should consider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

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

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

1. From the VMware vSphere Web Client home page, click 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 deployed for Epic software environments uses storage controllers in a high-availability (HA) pair configuration. Storage is required to be presented from both controllers to Epic database servers over the FC Protocol (FCP). The configuration presents storage from both controllers to evenly balance the application load during normal operation.

Epic requirements for separating production workloads into fault domains call pools is detailed in the Epic All-Flash Reference Architecture Strategy Handbook. Read this document in detail before continuing. Note that an ONTAP node can be considered a separate pool of storage.

ONTAP configuration

This section describes a sample deployment and provisioning procedures using the relevant ONTAP commands. The emphasis is to show how storage is provisioned to implement the storage layout recommended by NetApp, which uses an HA controller pair. One of the major advantages with ONTAP is the ability to scale out without disturbing the existing HA pairs.

Epic provides detailed storage performance requirements and layout guidance, including the storage presentation and host-side storage layout, to each customer. Epic will provide these custom documents:

- The Epic Hardware Configuration Guide used for sizing during presales.
- The Epic Database Storage Layout Recommendations used for LUN and volume layout during deployment.

A customer-specific storage system layout and configuration that meet these requirements must be developed by referring to the Epic Database Storage Layout Recommendations.

The following example describes the deployment of an AFF A700 storage system supporting a 10TB database. The provisioning parameters of the storage used to support the production database in the example deployment are shown in the table below.

Parameter	Controller 1	Controller 2
Controller host name	Prod1-01	Prod1-02
Aggregates ONTAP	aggr0_prod1-01 (ADP 11-partitions)	aggr0_prod1-02 (ADP 11-partitions)
Aggregates data	Prod1-01_aggr1 (22-partitions)	Prod1-02_aggr1 (22-partitions)
Volumes (size)	epic_prod_db1 (2TB) epic_prod_db2 (2TB) epic_prod_db3 (2TB) epic_prod_db4 (2TB) epic_prod_db5 (2TB) epic_prod_db6 (2TB) epic_prod_db7 (2TB) epic_prod_db8 (2TB) epic_prod_inst (1TB) epic_prod_jrn1 (1200GB) epic_prod_jrn2 (1200GB)	epic_report_db1 (2TB) epic_report_db2 (2TB) epic_report_db3 (2TB) epic_report_db4 (2TB) epic_report_db5 (2TB) epic_report_db6 (2TB) epic_report_db7 (2TB) epic_report_db8 (2TB) epic_report_inst (1TB) epic_report_jrn1 (1200GB) epic_report_jrn2 (1200GB)
LUN paths (size)	/epic_prod_db1/epic_prod_db1 (1.4TB) /epic_prod_db2/epic_prod_db2 (1.4TB) /epic_prod_db3/epic_prod_db3 (1.4TB) /epic_prod_db4/epic_prod_db4 (1.4TB) /epic_prod_db5/epic_prod_db5 (1.4TB) /epic_prod_db6/epic_prod_db6 (1.4TB) /epic_prod_db7/epic_prod_db7 (1.4TB) /epic_prod_db8/epic_prod_db8 (1.4TB) /epic_prod_inst/epic_prod_inst (700GB) /epic_prod_jrn1/epic_prod_jrn1 (800GB) /epic_prod_jrn2/epic_prod_jrn2 (800GB)	/epic_prod_db1/epic_report_db1 (1.4TB) /epic_prod_db2/epic_report_db2 (1.4TB) /epic_prod_db3/epic_report_db3 (1.4TB) /epic_prod_db4/epic_report_db4 (1.4TB) /epic_prod_db5/epic_report_db5 (1.4TB) /epic_prod_db6/epic_report_db6 (1.4TB) /epic_prod_db7/epic_report_db7 (1.4TB) /epic_prod_db8/epic_report_db8 (1.4TB) /epic_report_inst/epic_report_inst (700GB) /epic_report_jrn1/epic_report_jrn1 (800GB) /epic_report_jrn2/epic_report_jrn2 (800GB)
VMs	RHEL	RHEL
LUN type	Linux (mounted as RDMs directly by the RHEL VMs using FC)	Linux (mounted as RDMs directly by the RHEL VMs using FC)
FCP initiator group (igroup) name	ig_epic_prod (Linux)	ig_epic_report (Linux)
Host operating system	VMware	VMware
Epic database server host name	epic_prod	epic_report

Parameter	Controller 1	Controller 2
SVM	svm_prod	svm_ps (production services) svm_cifs

ONTAP licenses

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

To apply the licenses, open NetApp System Manager and go to Configuration-Licenses and add appropriate licenses. Alternatively, run the following command to add licenses using the CLI:

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

AutoSupport configuration

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

1. Run the following ONTAP command. Set the partner address option to the IP address of the management port for prod1-01.

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

2. Run the following ONTAP command. Set the partner address option to the IP address of the management port for cluster1-02.

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

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

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

ONTAP storage provisioning

The storage provisioning workflow is as follows:

1. Create the aggregates.
2. Create a storage virtual machine (SVM).

After aggregate creation, the next step is to create an SVM. In ONTAP the storage is virtualized in the form of an SVM. Hosts and clients no longer access the physical storage hardware. Create an SVM using the System Manager GUI or the CLI.

3. Create FC LIFs.

Ports and storage are provisioned on the SVM and presented to hosts and clients through virtual ports called logical interfaces (LIFs).

You can run all the workloads in one SVM with all the protocols. For Epic, NetApp recommends having an SVM for production FC and one SVM for CIFS.

- a. Enable and start FC from SVM settings in the System Manager GUI.
 - b. Add FC LIFs to the SVM. Configure multiple FC LIFs on each storage node, depending on the number of paths architected per LUN.
4. Create initiator groups (igroups).

Igroups are tables of FC- protocol host WWPNs or iSCSI host node names that define which LUNs are available to the hosts. For example, if you have a host cluster, you can use igroups to ensure that specific LUNs are visible to only one host in the cluster or to all the hosts in the cluster. You can define multiple igroups and map them to LUNs to control which initiators have access to LUNs.

Create FC igroups of type VMware using the System Manager GUI or the CLI.

5. Create zones on the FC switch.

An FC or FCoE zone is a logical grouping of one or more ports in a fabric. For devices to be able to see each other, connect, create sessions with one another, and communicate, both ports need to have a common zone membership. Single initiator zoning is recommended.

- a. Create zones on the switch and add the NetApp target and the Cisco UCS blade initiators in the zone.

NetApp best practice is single initiator zoning. Each zone contains only one initiator and the target WWPN on the controller. The zones use the port name and not the node name.

6. Create volumes and LUNs.

- a. Create volumes to host the LUNs using the System Manager GUI (or the CLI). All the storage efficiency settings and data protection are set by default on the volume. You can optionally turn on volume

encryption and QoS policies on the volume using the `vol modify` command. Note that the volumes need to be large enough to contain the LUNs and Snapshot copies. To protect the volume from capacity issues, enable the `autosize` and `autodelete` options. After the volumes are created, create the LUNs that will house the Epic workload.

- b. Create FC LUNs of type `VMware` that will host the Epic workload using the System Manager GUI (or the CLI). NetApp has simplified LUN creation in a very easy to follow wizard in System Manager.

You can also use VSC to provision volumes and LUNs. See the [FC Configuration for ESX Express Guide](#).

See the [SAN Administration and the SAN Configuration Guide](#) if you are not using VSC.

7. Map the LUNs to the igroups.

After the LUNs and igroups are created, map the LUNs to the relevant igroups that give the desired hosts access to the LUNs.

The LUNs are now ready to be discovered and mapped to the ESXi servers. Refresh the storage on the ESXi hosts and add the newly discovered LUNs.

Next: [GenIO Tool](#).

GenIO tool

GenIO is the storage-performance testing tool used by Epic. It simulates the workload generated by an InterSystems Caché database used in an Epic production environment, including the write-cycle patterns. It is available as a command-line application on various host operating systems on which Caché is deployed. Always test with the latest copy of the GenIO tool from Epic.

A performance test run involves executing the GenIO application on the production Epic database host with a set of I/O parameters. These parameters simulate the I/O patterns for the customer's Epic environment, including the write cycles.

Epic pushes the controller past the 100% full load detailed in the hardware configuration guide to determine how much headroom is on the controller. Epic also runs a full load test and simulates backup operations.

Epic server support representatives use it to verify storage performance from the host perspective. NetApp has also used GenIO to validate the performance of NetApp storage systems in the lab.

Where to find additional information

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

FlexPod design zone

- FlexPod design zone

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

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

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

- Cisco Best Practices with Epic on Cisco UCS

https://www.cisco.com/c/dam/en_us/solutions/industries/healthcare/Epic_on_UCS_tech_brief_FNL.pdf

NetApp technical reports

- TR-3929: Reallocate Best Practices Guide

<https://fieldportal.netapp.com/content/192896>

- TR-3987: Snap Creator Framework Plug-In for InterSystems Caché

<https://fieldportal.netapp.com/content/248308>

- TR-3928: NetApp Best Practices for Epic

<https://fieldportal.netapp.com/?oparams=68646>

- TR-4017: FC SAN Best Practices

<http://media.netapp.com/documents/tr-4017.pdf>

- TR-3446: SnapMirror Async Overview and Best Practices Guide

<http://media.netapp.com/documents/tr-3446.pdf>

ONTAP documentation

- NetApp Product Documentation

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

- Virtual Storage Console (VSC) for vSphere documentation

<https://mysupport.netapp.com/documentation/productlibrary/index.html?productID=30048>

- ONTAP 9 Documentation Center

<http://docs.netapp.com/ontap-9/index.jsp>

- FC Express Guide for ESXi

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.exp-fc-esx-cpg/home.html>

- All ONTAP 9.3 Documentation

<https://mysupport.netapp.com/documentation/docweb/index.html?productID=62579>

- Software Setup Guide

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.dot-cm-ssg/home.html?lang=dot-cm-ssg>

- Disks and Aggregates Power Guide

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.dot-cm-psmg/home.html?lang=dot-cm-psmg>

- SAN Administration Guide

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.dot-cm-sanag/home.html?lang=dot-cm-sanag>

- SAN Configuration Guide

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.dot-cm-sanconf/home.html?lang=dot-cm-sanconf>

- FC Configuration for Red Hat Enterprise Linux Express Guide

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.exp-fc-rhel-cg/home.html?lang=exp-fc-rhel-cg>

- FC Configuration for Windows Express Guide

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.exp-fc-cpg/home.html?lang=exp-fc-cpg>

- FC SAN Optimized AFF Setup Guide

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.cdott-fcsan-optaff-sg/home.html?lang=cdott-fcsan-optaff-sg>

- High-Availability Configuration Guide

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.dot-cm-hacg/home.html?lang=dot-cm-hacg>

- Logical Storage Management Guide

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.dot-cm-vsmg/home.html?lang=dot-cm-vsmg>

- Performance Management Power Guide

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.pow-perf-mon/home.html?lang=pow-perf-mon>

- SMB/CIFS Configuration Power Guide

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.pow-cifs-cg/home.html?lang=pow-cifs-cg>

- SMB/CIFS Reference

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.cdott-famg-cifs/home.html?lang=cdott-famg-cifs>

- Data Protection Power Guide

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.pow-dap/home.html?lang=pow-dap>

- Data Protection Tape Backup and Recovery Guide

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.dot-cm-ptbrg/home.html?lang=dot-cm-ptbrg>

- NetApp Encryption Power Guide

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.pow-nve/home.html?lang=pow-nve>

- Network Management Guide

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.dot-cm-nmg/home.html?lang=dot-cm-nmg>

- Commands: Manual Page Reference for ONTAP 9.3

<http://docs.netapp.com/ontap-9/topic/com.netapp.doc.dot-cm-cmpr-930/home.html?lang=dot-cm-cmpr-930>

Cisco Nexus, MDS, Cisco UCS, and Cisco UCS Manager guides

- Cisco UCS Servers Overview

<https://www.cisco.com/c/en/us/products/servers-unified-computing/index.html>

- Cisco UCS Blade Servers Overview

<https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/index.html>

- Cisco UCS B200 M5 Datasheet

<https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/index.html>

- Cisco UCS Manager Overview

<https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-manager/index.html>

- Cisco UCS Manager 3.2(3a) Infrastructure Bundle (requires Cisco.com authorization)

<https://software.cisco.com/download/home/283612660/type/283655658/release/3.2%25283a%2529>

- Cisco Nexus 9300 Platform Switches

<https://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/datasheet-c78-736967.html>

- Cisco MDS 9148S FC Switch

<https://www.cisco.com/c/en/us/products/storage-networking/mds-9148s-16g-multilayer-fabric-switch/index.html>

Acknowledgements

The following people contributed to the writing of this guide.

- Ganesh Kamath, Technical Marketing Engineer, NetApp
- Atul Bhalodia, Technical Marketing Engineer, NetApp
- Ketan Mota, Product Manager, NetApp
- Jon Ebmeier, Cisco
- Mike Brennan, Cisco

FlexPod for Epic Performance Testing

TR-4784: FlexPod for Epic Performance Testing

Brian O'Mahony,
Ganesh Kamath
Atul Bhalodia
Brandon Agee

In partnership with:



Objective

The objective of this report is to highlight the performance of FlexPod with NetApp All Flash A300 and A700 storage systems with Epic Healthcare workloads.

Epic Hardware configuration guide

For acceptable end-user performance, Epic production and disaster recovery operational database (ODB) target- read and target- write time requirements are as follows:

- For randomly placed reads to database files measured at the system call level:
 - Average read latencies must be 2ms or less
 - 99% of read latencies must be below 60ms
 - 99.9% of read latencies must be below 200ms
 - 99.99% of read latencies must be below 600ms
- For randomly placed writes to database files measured at the system call level:
 - Average write latencies must be 1ms or less depending on size



These requirements change with time. Epic prepares a customer- specific Epic Hardware Configuration Guide (HCG). Refer to your HCG for details on requirements.

Overall solution benefits

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

- **Simplified operations and lowered costs.** Eliminate the expense and complexity of legacy proprietary RISC/UNIX platforms by replacing them with a more efficient and scalable shared resource capable of supporting clinicians wherever they are. This solution delivers higher resource utilization for greater ROI.
- **Quicker deployment of infrastructure.** Whether it's in an existing data center or in a remote location, the integrated and tested design of FlexPod Datacenter with Epic enables customers to have new infrastructure up and running in less time with less effort.
- **Scale-out architecture.** Scale SAN and NAS from terabytes to tens of petabytes without reconfiguring

running applications.

- **Nondisruptive operations.** Perform storage maintenance, hardware lifecycle operations, and software upgrades without interrupting business operations.
- **Secure multitenancy.** FlexPod supports the needs of shared virtualized server and storage infrastructure, enabling secure multitenancy of facility-specific information, particularly if you are hosting multiple instances of databases and software.
- **Pooled resource optimization.** FlexPod can help reduce physical server and storage controller counts and load- balance workload demands. It can also boost utilization while improving performance.
- **Quality of service (QoS).** FlexPod offers QoS on the entire stack. Industry-leading QoS storage policies enable differentiated service levels in a shared environment. These policies enable optimal performance for workloads and help isolate and control runaway applications.
- **Storage efficiency.** Reduce storage costs with the NetApp 7:1 storage efficiency guarantee.
- **Agility.** The industry-leading workflow automation, orchestration, and management tools offered by FlexPod systems allow IT to be far more responsive to business requests. These business requests can range from Epic backup and provisioning of additional test and training environments to analytics database replications for population health management initiatives.
- **Productivity.** Quickly deploy and scale this solution for optimal clinician end-user experiences.
- **Data Fabric.** The NetApp Data Fabric architecture weaves data together across sites, beyond physical boundaries, and across applications. The NetApp Data Fabric is built for data-driven enterprises in a data-centric world. Data is created and used in multiple locations, and it often must be leveraged and shared with other locations, applications, and infrastructures. Customers want a way to manage data that is consistent and integrated. It provides a way to manage data that puts IT in control and simplifies ever-increasing IT complexity.

Cisco Unified Computing System, Cisco Nexus and MDS Switching, and ONTAP all-flash storage

The FlexPod for Epic Healthcare delivers the performance, efficiency, manageability, scalability, and data protection that IT organizations need to meet for the most stringent Epic requirements. By accelerating Epic production database performance and by reducing application deployment time from months to weeks, FlexPod helps organizations maximize the potential of their Epic investment.

Cisco Unified Computing System

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

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

The system is designed as a single virtual blade chassis that incorporates and scales across multiple blade chassis. The system implements a radically simplified architecture that eliminates the multiple redundant devices that populate traditional blade server chassis and result in layers of complexity. Examples include Ethernet switches, Fibre Channel switches, and chassis management modules. The Cisco UCS contains a redundant pair of Cisco fabric interconnects that provide a single point of management and a single point of

control for all I/O traffic.

The Cisco UCS uses service profiles to help ensure that virtual servers in the UCS infrastructure are configured correctly. Service profiles include critical server information about the server identity such as LAN and SAN addressing, I/O configurations, firmware versions, boot order, network VLAN, physical port, and quality-of-service (QoS) policies. Service profiles can be dynamically created and associated with any physical server in the system within minutes rather than within 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 any physical configuration changes. It facilitates rapid bare-metal provisioning of replacements for failed servers.

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

Cisco UCS has been extensively tested with Epic over a multiyear period to demonstrate that it meets server configuration requirements. Cisco UCS is a supported server platform, as listed in customers' "Epic Hardware Configuration Guide."

Cisco Nexus and Cisco MDS Ethernet and Fibre Channel switching

Cisco Nexus switches and MDS multilayer directors provide enterprise-class connectivity and SAN consolidation. Cisco multiprotocol storage networking reduces business risk by providing flexibility and options. Supported protocols include Fibre Channel (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 the data center and the campus core. They also offer a full feature set for data-center aggregation, end-of-row deployments, and data center interconnect deployments in a highly resilient, modular platform.

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

In summary, the Cisco UCS provides the following important advantages for Epic deployments:

- **Infrastructure scalability.** Virtualization, efficient power and cooling, cloud scale with automation, high density, and performance all support efficient data-center growth.
- **Operational continuity.** The design integrates hardware, NX-OS software features, and management to support zero-downtime environments.
- **Transport flexibility.** Incrementally adopt new networking technologies with a cost-effective solution.

Together, Cisco UCS with Cisco Nexus switches and MDS multilayer directors provide a compelling computer, networking, and SAN connectivity solution for Epic.

NetApp all-flash storage systems

NetApp AFF systems address enterprise storage requirements with high performance, superior flexibility, and best-in-class data management. Built on ONTAP data management software, AFF systems speed up your business without compromising the efficiency, reliability, or flexibility of your IT operations. With enterprise-grade all-flash arrays, AFF systems accelerate, manage, and protect your business-critical data and enable an

easy and risk-free transition to flash media for your data center.

Designed specifically for flash, AFF A-series all-flash systems deliver industry-leading performance, capacity, density, scalability, security, and network connectivity in a dense form factor. With the addition of a new entry-level system, the new AFF A-series family extends enterprise-grade flash to midsize businesses. At up to seven million IOPS per cluster with sub-millisecond latency, the AFF A series is the fastest family of all-flash arrays, built on a true unified scale-out architecture.

With the AFF A series, you can complete twice the work at half the latency relative to the previous generation of AFF systems. The members of the AFF A series are the industry's first all-flash arrays that provide both 40Gb Ethernet (40GbE) and 32Gb Fibre Channel (FC) connectivity. Therefore, they eliminate the bandwidth bottlenecks that are increasingly moving from storage to the network as flash becomes faster and faster.

NetApp has taken the lead for all-flash storage innovations with the latest solid-state-drive (SSD) technologies. As the first all-flash array to support 15TB SSDs, AFF systems, with the introduction of the A series, also become the first to use multistream write SSDs. Multistream write capability significantly increases the usable capacity of SSDs.

NetApp ONTAP Flash Essentials is the power behind the performance of All Flash FAS. ONTAP is industry-leading data management software. However, it is not widely known that ONTAP, with its NetApp WAFL (Write Anywhere File Layout) file system, is natively optimized for flash media.

ONTAP Flash Essentials optimizes SSD performance and endurance with the following features, among others:

- NetApp data-reduction technologies, including inline compression, inline deduplication, and inline data compaction, can provide significant space savings. Savings can be further increased by using NetApp Snapshot and NetApp FlexClone technologies. Studies based on customer deployments have shown that these data-reduction technologies have enabled space savings of up to 933 times.
- Coalesced writes to free blocks maximize performance and flash media longevity.
- Flash-specific read-path optimizations provide consistent low latency.
- Parallelized processing handles more requests at once.
- Software-defined access to flash maximizes deployment flexibility.
- Advanced Disk Partitioning (ADP) increases storage efficiency and further increases usable capacity by almost 20%.
- The Data Fabric enables live workload migration between flash and hard-disk-drive tiers on the premises or to the cloud.

QoS capability guarantees minimum service-level objectives in multiworkload and multitenant environments.

The key differentiators with adaptive QOS are as follows:

- Simple self-managing IOPS/TB or throughput MB/TB. Performance grows as data capacity grows.
- Simplified consumption of storage based on service-level performance policies.
- Consolidation of mixed workloads onto a single cluster with guaranteed performance service levels. No more silos are required for critical applications.
- Major cost saving by consolidating nodes and disk.

Executive Summary

To showcase the storage efficiency and performance of NetApp’s All Flash FAS platform, NetApp performed a study to measure Epic EHR performance on AFF A300 and AFF A700 systems. NetApp measured the data throughput, peak IOPS, and average latency of an AFF A300 running ONTAP 9.5 and an AFF A700 storage controller running ONTAP 9.4, each running an Epic EHR workload. In a manner similar to SPC-3 testing, all inline storage efficiency features were enabled.

We ran the Epic GenIO workload generator on an AFF A300 cluster that contained a total of twenty-four 3.8TB SSDs and on an AFF A700 cluster that contained a total of forty-eight 3.8TB SSDs. We tested our cluster at a range of load points that drove the storage to peak CPU utilization. At each load point, we collected information about the storage IOPS and latency.

NetApp has consistently with each software upgrade improved performance in the range of 40-50%. Innovation with performance enhancements has varied based on workload and protocol.

The Epic performance test demonstrated that the AFF A300 cluster IOPS increased from 75,000 IOPS at <1ms to a peak performance of 188,929 IOPS at <1ms. For all load points at or below 200,000 IOPS, we were able to maintain consistent storage latencies of no greater than 1ms. Additionally, the Epic performance test demonstrated that the AFF A700 cluster IOPS increased from 75,000 IOPS at <1ms to a peak performance of 319,000 IOPS at <1ms. For all load points at or below 320,000 IOPS, we were able to maintain consistent storage latencies of no greater than 1ms.

Test methodology

Test plan

The GenerationIO tool (GenIO) is used by Epic to validate that storage is production ready. This test focuses on performance by pushing storage to its limits and determining the headroom on storage controllers by ramping up until requirements fail.

The tests performed here are focused on determining headroom as well as using Adaptive Quality of Service (AQOS) to protect critical Epic workloads. For AFF A300 testing, two servers are used with GenIO loaded on both to drive I/O on the storage controllers. Three servers are used with GenIO loaded on all three to drive I/O on the AFF A700 storage controllers. Three servers are used because of server performance limits, and three servers are required for an AFF A700.

Test environment

Hardware and software

For this study, we configured three Red Hat Linux virtual machines (VMs) on VMware ESXi 6.5 running on Cisco UCS B200-M5s. We connected the ESXi hosts to the AFF storage controller nodes with Cisco MDS-series switches by using 16Gb FC on the server side and 16Gb FC on the storage side. The AFF A700 nodes were connected to one DS2446 disk shelf with 3.8TB SSDs by following NetApp cabling best practices.

The three tables below list the hardware and software components that we used for the Epic performance test configuration.

The following table lists Epic Test hardware and software components.

Hardware and software components	Details
Operating system for VM	RHEL 7.4 VMs

Hardware and software components	Details
Operating system on server blades	VMware ESXi 6.5
Physical server	Cisco UCS B200 M5 x 3
Processors per server	Two 20-core Intel Xeon Gold 6148 2.4Ghz
Physical memory per server	768GB
FC network	16Gb FC with multipathing
FC HBA	FC vHBA on Cisco UCS VIC 1340
Dedicated public 1GbE ports for cluster management	Two Intel 1350GbE ports
16Gb FC switch	Cisco MDS 9148s
40GbE switch	Cisco Nexus 9332 switch

The following table lists NetApp AFF A700 and AFF A300 storage system hardware and software.

Hardware and software components	AFF A700 details	AFF A300 details
Storage system	AFF A700 controller configured as a high-availability (HA) active-active pair	AFF A300 controller configured as a high-availability (HA) active-active pair
ONTAP version	9.4	9.5
Total number of drives	36	24
Drive size	3.8TB	3.8TB
Drive type	SSD	SSD
FC target ports	Eight 16Gb ports (four per node)	Eight 16Gb ports (four per node)
Ethernet ports	Four 10Gb ports (two per node)	Four 10Gb ports (two per node)
Storage virtual machines (SVMs)	One SVM across both node aggregates	One SVM across both node aggregates
Ethernet logical interfaces (LIFs)	Four 1Gb management LIFs (two per node connected to separate private VLANs)	Four 1Gb management LIFs (two per node connected to separate private VLANs)
FC LIFs	Four 16Gb data LIFs	Four 16Gb data LIFs

The following table lists NetApp AFF A700 and AFF A300 storage system layout.

Storage layout	AFF A700 details	AFF A300 details
SVM	Single SVM for Epic application databases	Single SVM for Epic application databases
Aggregates	Two 20TB each	Two 30TB each
Volumes for production	Sixteen 342GB volumes per RHEL VM	Sixteen 512GB volumes per RHEL VM

Storage layout	AFF A700 details	AFF A300 details
LUNs for production	Sixteen 307GB LUNs, one per volume	Sixteen 460GB LUNs, one per volume
Volumes for journal	Two 95Gb volumes per RHEL VM	Two 240Gb volumes per RHEL VM
LUNs for journal	Two 75Gb LUNs, one per volume	Two 190Gb LUNs, one per volume

Workload testing

AFF A300 procedure

The AFF A300 HA pair can comfortably run the largest Epic instance in existence. If you have two or more very large Epic instances, you might need to use an AFF A700, based on the outcome of the NetApp SPM tool.

Data generation

Data inside the LUNs were generated with Epic's Dgen.pl script. The script is designed to create data similar to what would be found inside an Epic database.

The following Dgen command was run from both RHEL VMs, epic-rhel1 and epic-rhel2:

```
./dgen.pl --directory "/epic" --jobs 2 --quiet --pctfull 20
```

-pctfull is optional and defines the percentage of the LUN to fill with data. The default is 95%. The size does not affect performance, but it does affect the time to write the data to the LUNs.

After the dgen process is complete, you can run the GenIO tests for each server.

Run GenIO

Two servers were tested. A ramp run from 75,000 to 110,000 IOPS was executed, which represents a very large Epic environment. Both tests were run at the same time.

Run the following GenIO command from the server epic-rhel1:

```
./RampRun.pl -miniops 75000 --maxiops 110000 --background --disable-warmup
--runtime 30 --wijfile /epic/epicjrn/GENIO.WIJ --numruns 10 --system epic-
rhell --comment Ramp 75-110k
```

GenIO result on the AFF A300

The following table lists GenIO results on the AFF A300

Read IOPs	Write IOPs	Total IOPs	Longest write cycle (sec)	Effective write latency (ms)	Randread average (ms)
142505	46442	188929	44.68	0.115	0.66

AFF A700 procedure

For larger Epic environments, typically greater than ten million global references, customers can choose the AFF A700.

Data generation

Data inside the LUNs was generated with Epic's Dgen.pl script. The script is designed to create data similar to what would be found inside an Epic database.

Run the following Dgen command on all three RHEL VMs.

```
./dgen.pl --directory "/epic" --jobs 2 --quiet --pctfull 20
```

-pctfull is optional and defines the percentage of the LUN to fill with data. The default is 95%. The size does not affect performance, but it does affect the time to write the data to the LUNs.

After the dgen process is complete you are ready to run the GenIO tests for each server.

Run GenIO

Three servers were tested. On two servers, a ramp run from 75,000 to 100,000 IOPs was executed, which represents a very large Epic environment. The third server was set up as a bully to ramp run from 75,000 IOPS to 170,000 IOPS. All three tests were run at the same time.

Run the following GenIO command from the server epic-rhel1:

```
./RampRun.pl -miniops 75000 --maxiops 100000 --background --disable-warmup  
--runtime 30 --wijfile /epic/epicjrn/GENIO.WIJ --numruns 10 --system epic-  
rhel1 --comment Ramp 75-100k
```

GenIO results on the AFF A700

The following table presents the GenIO results a test of the AFF A700.

Read IOPs	Write IOPs	Total IOPs	Longest write cycle (sec)	Effective write latency (ms)	Randread average (ms)
241,180	78,654	319,837	43.24	0.09	1.05

Performance SLA with AQOS

NetApp can set floor and ceiling performance values for workloads using AQOS policies. The floor setting guarantees minimum performance. IOPS/TB can be applied to a group of volumes for an application like Epic. The Epic workload assigned to a QoS policy is protected from other workloads on the same cluster. The minimum requirements are guaranteed while allowing the workload to peak and use available resources on the controller.

In this test, server 1 and server 2 were protected with AQOS, and the third server acted as a bully workload to cause performance degradation within the cluster. AQOS allowed servers 1 and 2 to perform at the specified SLA, while the bully workload showed signs of degradation with longer write cycles.

Adaptive quality of service defaults

ONTAP comes configured with three default AQOS policies: value, performance, and extreme. The values for each policy can be view with the `qos` command. Use `-instant` at the end of the command to view all AQOS settings.

```
::> qos adaptive-policy-group show
Name          Vserver Wklds  Expected IOPS      Peak IOPS
extreme       fp-g9a  0      6144IOPS/TB 12288IOPS/TB
performance   fp-g9a  0      2048IOPS/TB 4096IOPS/TB
value         fp-g9a  0      128IOPS/TB  512IOPS/TB
```

Here is the syntax to create an AQOS policy:

```
::> qos adaptive-policy-group modify -policy-group aqos- epic-prod1
-expected-iops 5000 -peak-iops 10000 -absolute-min-iops 4000 -peak-iops
-allocation used-space
```

There are a few important settings in an AQOS policy:

- **Expected IOPS.** This adaptive setting is the minimum IOPS/TB value for the policy. Workloads are guaranteed to get at least this level of IOPS/TB. This is the most important setting in this testing. In our example test, the performance AQOS policy was set to 2048IOPS/TB.
- **Peak IOPS.** This adaptive setting is the maximum IOPS/TB value for the policy. In our example test, the performance AQOS policy was set to 4096IOPS/TB.
- **Peak IOPS allocation.** Options are allocated space or used space. Set this parameter to used space, because this value changes as the database grows in the LUNs.
- **Absolute minimum IOPS.** This setting is static and not adaptive. This parameter sets the minimum IOPS regardless of size. This value is only used when size is less than 1TB and has no effect on this testing.

Typically, Epic workloads in production run at about ~1000 IOPS/TB of storage and capacity, and IOPS grows linearly. The default AQOS performance profile is more than adequate for an Epic workload.

For this testing the lab did not reflect a production size database with a smaller size of 5TB. The goal was to run each test at 75,000 IOPS. The setting for the EpicProd AQOS policy is shown below.


- Expected IOPS/TB = Total IOPS/used space
- 15,000 IOPS/TB = 75,000 IOPS/5TB

The following table presents the settings that were used for the EpicProd AQOS policy.

Setting	Value
Volume size	5TB
Required IOPS	75,000
peak-iops-allocation	Used space
Absolute minimum IOPS	7,500

Setting	Value
Expected IOPS/TB	15,000
Peak IOPS/TB	30,000

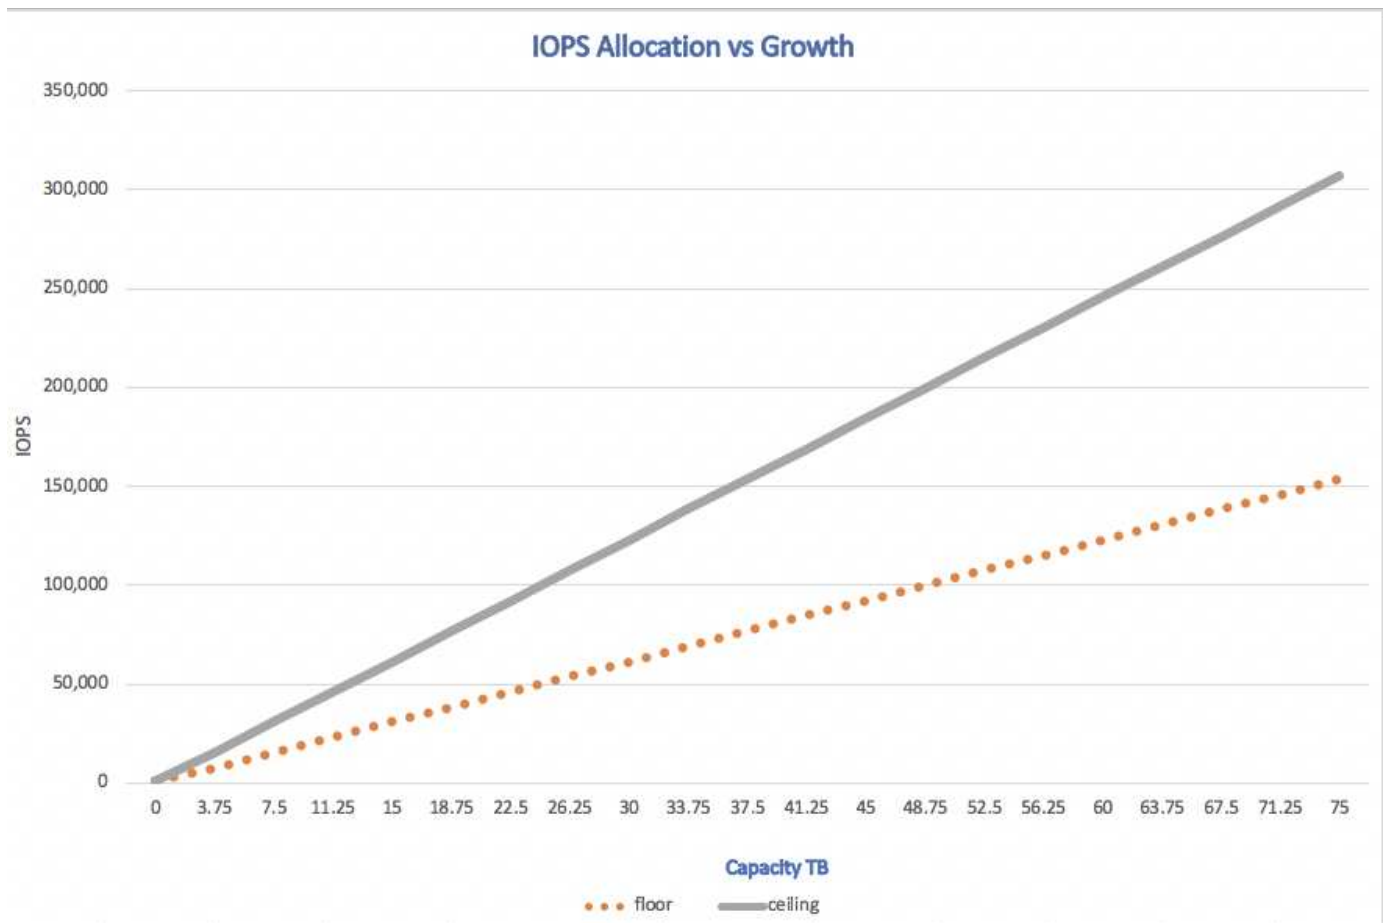
The following figure shows how floor IOPS and ceiling IOPS are calculated as the used space grows over time.



For a production-sized database, you can either create a custom AQOS profile like the one used in the last example, or you can use the default performance AQOS policy. The settings for the performance AQOS policy are show in the table below.

Setting	Value
Volume size	75TB
Required IOPS	75,000
peak-iops-allocation	Used space
Absolute minimum IOPS	500
Expected IOPS/TB	1,000
Peak IOPS/TB	2,000

The following figure shows how floor and ceiling IOPS are calculated as the used space grows over time for the default performance AQOS policy.



Parameters

- The following parameter specifies the name of the adaptive policy group:

```
-policy-group <text> - Name
```

Adaptive policy group names must be unique and are restricted to 127 alphanumeric characters including underscores "_" and hyphens "-". Adaptive policy group names must start with an alphanumeric character. Use the `qos adaptive-policy-group rename` command to change the adaptive policy group name.

- The following parameter specifies the data SVM (called vserver in the command line) to which this adaptive policy group belongs.

```
-vserver <vserver name> - Vserver
```

You can apply this adaptive policy group to only the storage objects contained in the specified SVM. If the system has only one SVM, then the command uses that SVM by default.

- The following parameter specifies the minimum expected IOPS/TB or IOPS/GB allocated based on the storage object allocated size.

```
-expected-iops {<integer>[IOPS[/{GB|TB}]] (default: TB)} - Expected IOPS
```

- The following parameter specifies the maximum possible IOPS/TB or IOPS/GB allocated based on the storage object allocated size or the storage object used size.

```
-peak-iops {<integer>[IOPS[/{GB|TB}]] (default: TB)} - Peak IOPS
```

- The following parameter specifies the absolute minimum IOPS that is used as an override when the expected IOPS is less than this value.

```
[-absolute-min-iops <qos_tput>] - Absolute Minimum IOPS
```

The default value is computed as follows:

```
qos adaptive-policy-group modify -policy-group aqos- epic-prod1  
-expected-iops 5000 -peak-iops 10000 -absolute-min-iops 4000 -peak-iops  
-allocation used-space
```

```
qos adaptive-policy-group modify -policy-group aqos- epic-prod2  
-expected-iops 6000 -peak-iops 20000 -absolute-min-iops 5000 -peak-iops  
-allocation used-space
```

```
qos adaptive-policy-group modify -policy-group aqos- epic-bully  
-expected-iops 3000 -peak-iops 2000 -absolute-min-iops 2000 -peak-iops  
-allocation used-space
```

Data generation

Data inside the LUNs was generated with the Epic `Dgen.pl` script. The script is designed to create data similar to what would be found inside an Epic database.

The following `Dgen` command was run on all three RHEL VMs:

```
./dgen.pl --directory "/epic" --jobs 2 --quiet --pctfull 20
```

Run GenIO

Three servers were tested. Two ran at a constant 75,000 IOPS, which represents a very large Epic environment. The third server was setup as a bully to ramp run from 75,000 IOPS to 150,000 IOPS. All three

tests were run at the same time.

Server epic_rhel1 GenIO test

The following command was run to assign EpicProd AQOS settings to each volume:

```
::> vol modify -vserver epic -volume epic_rhel1_* -qos-adaptive-policy  
-group AqosEpicProd
```

The following GenIO command was run from the server epic-rhel1:

```
./RampRun.pl --miniops 75000 --maxiops 75000 --background --disable-warmup  
--runtime 30 --wijfile /epic/GENIO.WIJ --numruns 10 --system epic-rhel1  
--comment Ramp constant 75k
```

Server epic_rhel2 GenIO test

The following command was run to assign EpicProd AQOS settings to each volume:

```
::> vol modify -vserver epic -volume epic_rhel2_* -qos-adaptive-policy  
-group AqosEpicProd
```

The following GenIO command was run from the server epic-rhel2:

```
./RampRun.pl --miniops 75000 --maxiops 75000 --background --disable-warmup  
--runtime 30 --wijfile /epic/GENIO.WIJ --numruns 10 --system epic-rhel2  
--comment Ramp constant 75k
```

Server epic_rhel3 GenIO test (bully)

The following command assigns no AQOS policy to each volume:

```
::> vol modify -vserver epic -volume epic_rhel3_* -qos-adaptive-policy  
-group non
```

The following GenIO command was run from the server epic-rhel3:

```
./RampRun.pl --miniops 75000 --maxiops 150000 --background --disable-  
warmup --runtime 30 --wijfile /epic/GENIO.WIJ --numruns 10 --system epic-  
rhel3 --comment Ramp 75-150k
```

AQOS test results

The tables in the following sections contain the output from the summary.csv files from each concurrent GenIO test. To pass the test, the longest write cycle must have been below 45 seconds. The effective write latency must have been below 1 millisecond.

Server epic_rhel1 GenIO results

The following table illustrates GenIO results for AQOS server epic_rhel1.

Run	Read IOPS	Write IOPS	Total IOPS	Longest write cycle (sec)	Effective write latency (ms)
10	55655	18176	73832	32.66	0.12
11	55653	18114	73768	34.66	0.1
12	55623	18099	73722	35.17	0.1
13	55646	18093	73740	35.16	0.1
14	55643	18082	73726	35.66	0.1
15	55634	18156	73791	32.54	0.1
16	55629	18138	73767	34.74	0.11
17	55646	18131	73777	35.81	0.11
18	55639	18136	73775	35.48	0.11
19	55597	18141	73739	35.42	0.11

Server epic_rhel2 GenIO results

The following table illustrates GenIO results for AQOS server epic_rhel2.

Run	Read IOPS	Write IOPS	Total IOPS	Longest write cycle (sec)	Effective write latency (ms)
10	55629	18081	73711	33.96	0.1
11	55635	18152	73788	28.59	0.09
12	55606	18154	73761	30.44	0.09
13	55639	18148	73787	30.37	0.09
14	55629	18145	73774	30.13	0.09
15	55619	18125	73745	30.03	0.09
16	55640	18156	73796	33.48	0.09
17	55613	18177	73790	33.32	0.09
18	55605	18173	73779	32.11	0.09
19	55606	18178	73785	33.19	0.09

Server epic_rhel3 GenIO results (bully)

The following table illustrates GenIO results for AQOS server epic_rhel3.

Run	Write IOPS	Total IOPS	Longest WIJ Time (sec)	Longest Write Cycle (sec)	Effective Write Latency (ms)
10	19980	81207	21.48	40.05	0.1
11	21835	88610	17.57	46.32	0.12
12	23657	95955	19.77	53.03	0.12
13	25493	103387	21.93	57.53	0.12
14	27331	110766	23.17	60.57	0.12
15	28893	117906	26.93	56.56	0.1
16	30704	125233	28.05	60.5	0.12
17	32521	132585	28.43	64.38	0.12
18	34335	139881	30	70.38	0.12
19	36361	147633	22.78	73.66	0.13

AQOS test results analysis

The results from the previous section demonstrate that the performance of the servers epic_rhel1 and epic_rhel2 are not affected by the bully workload on epic_rhel3. epic_rhel3 ramps up to 150,000 IOPS and starts to fail the GenIO test as it hits the limits of the controllers. The write cycle and latency on epic_rhel1 and epic_rhel2 stay constant while the bully server spirals out of control.

This illustrates how an AQOS minimum policy can effectively isolate workloads from bullies and guarantee a minimum level of performance.

AQOS has a number of benefits:

- It allows for a more flexible and simplified architecture. Critical workloads no longer need to be siloed and can coexist with noncritical workloads. All capacity and performance can be managed and allocated with software rather than by using physical separation.
- It saves on the amount of disk and controllers required for Epic running on an ONTAP cluster.
- It simplifies the provisioning of workloads to performance policies that guarantee consistent performance.
- Optionally, you can also implement of NetApp Service Level Manager to perform the following tasks:
 - Create a catalog of services to simplify provisioning of storage.
 - Deliver predictable service levels so that you can consistently meet utilization goals.
 - Define service-level objectives.

Conclusion

By 2020, all Epic customers must be on flash storage. NetApp ONTAP was the first all- flash array to get a high-comfort rating from Epic, and it is listed under Enterprise Storage Arrays. All NetApp platforms that run a GA version of ONTAP are high comfort.

Epic requires that critical workloads like Production, Report, and Clarity are physically separated on storage

allocations called pools. NetApp provides multiple pools of storage in a single cluster with each node and offers a simplified single cluster and single OS for the entire Epic solution. ONTAP supports all protocols for NAS and SAN, with mixed tiers of storage for SSD, HDD, and cloud.

The introduction of Adaptive QoS in ONTAP 9.3, with significant enhancements in ONTAP 9.4, allows for the creation of storage pools with software without the need for physical separation. This capability greatly simplifies architecture development, permits the consolidation of nodes and disks, and improves performance for critical workloads like production by spreading across nodes. It also eliminates storage performance issues caused by bullies and guarantees consistent performance for the life of the workload.

Where to find additional information

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

FlexPod Design Zone

- NetApp FlexPod Design Zone

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

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

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

- Cisco Best Practices with Epic on Cisco UCS

https://www.cisco.com/c/dam/en_us/solutions/industries/healthcare/Epic_on_UCS_tech_brief_FNL.pdf

NetApp technical reports

- TR-4693: FlexPod Datacenter for Epic EHR Deployment Guide

<https://www.netapp.com/us/media/tr-4693.pdf>

- TR-4707: FlexPod for Epic Directional Sizing Guide

<https://www.netapp.com/us/media/tr-4707.pdf>

- TR-3929: Reallocate Best Practices Guide

<https://www.netapp.com/us/media/tr-3929.pdf>

- TR-3987: Snap Creator Framework Plug-In for InterSystems Caché

<https://www.netapp.com/us/media/tr-3987.pdf>

- TR-3928: NetApp Best Practices for Epic

<https://www.netapp.com/us/media/tr-3928.pdf>

- TR-4017: FC SAN Best Practices

<https://www.netapp.com/us/media/tr-4017.pdf>

- TR-3446: SnapMirror Async Overview and Best Practices Guide

<https://www.netapp.com/us/media/tr-3446.pdf>

ONTAP documentation

- NetApp product documentation

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

- Virtual Storage Console (VSC) for vSphere documentation

<https://mysupport.netapp.com/documentation/productlibrary/index.html?productID=30048>

- ONTAP 9 Documentation Center

<http://docs.netapp.com/ontap-9/index.jsp>

Cisco Nexus, MDS, Cisco UCS, and Cisco UCS Manager guides

- Cisco UCS Servers Overview

<https://www.cisco.com/c/en/us/products/servers-unified-computing/index.html>

- Cisco UCS Blade Servers Overview

<https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/index.html>

- Cisco UCS B200 M5 Datasheet

<https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/index.html>

- Cisco UCS Manager Overview

<https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-manager/index.html>

- Cisco UCS Manager 3.2(3a) Infrastructure Bundle (requires Cisco.com authorization)

<https://software.cisco.com/download/home/283612660/type/283655658/release/3.2%25283a%2529>

- Cisco Nexus 9300 Platform Switches

<https://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/datasheet-c78-736967.html>

- Cisco MDS 9148S FC Switch

<https://www.cisco.com/c/en/us/products/storage-networking/mds-9148s-16g-multilayer-fabric-switch/index.html>

Acknowledgements

The following people contributed to the creation of this guide:

- Ganesh Kamath, Technical Marketing Engineer, NetApp
- Atul Bhalodia, Technical Marketing Engineer, NetApp
- Brandon Agee, Technical Marketing Engineer, NetApp
- Brian O'Mahony, Solution Architect – Healthcare, NetApp
- Ketan Mota, Product Manager, NetApp
- Jon Ebmeier, Technical Solutions Architect, Cisco Systems, Inc
- Mike Brennan, Product Manager, Cisco Systems, Inc

FlexPod for MEDITECH Directional Sizing Guide

TR-4774: FlexPod for MEDITECH Directional Sizing

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

In partnership with:



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

Purpose

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

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

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

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

Overall solution benefits

Running a MEDITECH environment on the FlexPod architectural foundation can help healthcare organizations improve productivity and decrease capital and operating expenses. FlexPod provides a prevalidated, rigorously

tested, converged infrastructure from the strategic partnership of Cisco and NetApp. It is engineered and designed specifically for delivering predictable low-latency system performance and high availability. This approach results in faster response time for users of the MEDITECH EHR system.

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

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

Scope

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

It does not cover:

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

- Sizing for nonproduction workloads.

Audience

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

Related Documents

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

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



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

MEDITECH Workload Overview

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

MEDITECH and backup workloads

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

MEDITECH Host

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

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

MEDITECH workload

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

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

You can scale the MEDITECH system to provide increased capacity and performance by adding hosts. Each

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

MEDITECH Categories

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

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

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

MEDITECH Platforms

MEDITECH has four platforms:

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

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

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

MEDITECH Category	MEDITECH Platform	Average Random Read %	Average Random Write %	Average Sustained IOPS per MEDITECH Host
1	EXPANSE, 6.x	20	80	750

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

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

MEDITECH requires the average random write latency to stay below 1ms for each host. However, temporary increases of write latency up to 2ms during backup and reallocation jobs are considered acceptable. MEDITECH also requires the average random read latency to stay below 7ms for category 1 hosts and below 5ms for category 2 hosts. These latency requirements apply to every host regardless of which MEDITECH platform is being used.

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

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



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

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

Backup Workload Description

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

The workload that is generated by the backup software comes from the sequential reading of the LUNs that reside in the FlexClone volumes. The workload is defined as a 100% sequential read workload with a request size of 64,000. For the MEDITECH production workload, the performance criterion is to maintain the required IOPS and the associated read/write latency levels. For the backup workload, however, the attention is shifted to the overall data throughput (MBps) that is generated during the backup operation. MEDITECH LUN backups are required to be completed in an eight-hour backup window, but NetApp recommends that the backup of all MEDITECH LUNs be completed in six hours or less. Aiming to complete the backup in less than six hours mitigates for events such as an unplanned increase in the MEDITECH workload, NetApp ONTAP background operations, or data growth over time. Any of these events might incur extra backup time. Regardless of the amount of application data stored, the backup software performs a full block-level backup of the entire LUN for each MEDITECH host.

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

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

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

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

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

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

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

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

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

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

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

Cisco UCS Reference Architecture for MEDITECH

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

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

Technical specifications for small, medium and large architectures

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

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

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

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

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

	Small	Medium	Large
Effective capacity (2:1 storage efficiency)	55.6TiB	83.89TiB	101.47TiB



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

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

Software	Product family	Version or release	Details
Storage	ONTAP	ONTAP 9.4 general availability (GA)	
Network	Cisco UCS fabric interconnects	Cisco UCSM 4.x	Current recommended release
	Cisco Nexus Ethernet switches	7.0(3)I7(6)	Current recommended release
	Cisco FC: Cisco MDS 9132T	8.3(2)	Current recommended release
Hypervisor	Hypervisor	VMware vSphere ESXi 6.7	
	Virtual machines (VMs)	Windows 2016	
Management	Hypervisor management system	VMware vCenter Server 6.7 U1 (VCSA)	
	NetApp Virtual Storage Console (VSC)	VSC 7.0P1	
	NetApp SnapCenter	SnapCenter 4.0	
	Cisco UCS Manager	4.x	

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

Layer	Product family	Quantity and model	Details
Compute	Cisco UCS 5108 Chassis	1	Supports up to eight half-width or four full-width blades. Add chassis as server requirement grows.
	Cisco Chassis I/O Modules	2 x 2208	8GB x 10GB uplink ports
	Cisco UCS blade servers	4 x B200 M5	Each with 2 x 14 cores, 2.6GHz or higher clock speed, and 384GB BIOS 3.2(3#)
	Cisco UCS Virtual Interface Cards	4 x UCS 1440	VMware ESXi fNIC FC driver: 1.6.0.47 VMware ESXi eNIC Ethernet driver: 1.0.27.0 (See interoperability matrix: https://ucshcltool.cloudapps.cisco.com/public/)
	2 x Cisco UCS Fabric Interconnects (FI)	2 x UCS 6454 FI	4th-generation fabric interconnects supporting 10/25/100GB Ethernet and 32GB FC
Network	Cisco Ethernet switches	2 x Nexus 9336c-FX2	1GB, 10GB, 25GB, 40GB, 100GB
Storage network	IP Network Nexus 9k for BLOB storage		FI and UCS chassis
	FC: Cisco MDS 9132T		Two Cisco 9132T switches
Storage	NetApp AFF A300 all-flash storage system	1 HA Pair	2-node cluster for all MEDITECH workloads (File Server, Image Server, SQL Server, VMware, and so on)
	DS224C disk shelf	1 DS224C disk shelf	
	Solid-state drive (SSD)	9 x 3.8TB	

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

Layer	Product family	Quantity and model	Details
Compute	Cisco UCS 5108 chassis	1	Supports up to eight half-width or four full-width blades. Add chassis as server requirement grows.
	Cisco chassis I/O modules	2 x 2208	8GB x 10GB uplink ports
	Cisco UCS blade servers	6 x B200 M5	Each with 2 x 16 cores, 2.5GHz/or higher clock speed, and 384GB or more memory BIOS 3.2(3#)
	Cisco UCS virtual interface card (VIC)	6 x UCS 1440 VICs	VMware ESXi fNIC FC driver: 1.6.0.47 VMware ESXi eNIC Ethernet driver: 1.0.27.0 (See interoperability matrix:)
	2 x Cisco UCS Fabric Interconnects (FI)	2 x UCS 6454 FI	4th-generation fabric interconnects supporting 10GB/25GB/100GB Ethernet and 32GB FC
Network	Cisco Ethernet switches	2 x Nexus 9336c-FX2	1GB, 10GB, 25GB, 40GB, 100GB
Storage network	IP Network Nexus 9k for BLOB storage		
	FC: Cisco MDS 9132T		Two Cisco 9132T switches
Storage	NetApp AFF A220 all-flash storage system	2 HA Pair	2-node cluster for all MEDITECH workloads (File Server, Image Server, SQL Server, VMware, and so on)
	DS224C disk shelf	1 x DS224C disk shelf	
	SSD	13 x 3.8TB	

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

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



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

Additional Information

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

- FlexPod Datacenter with FC Cisco Validated Design.

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

- NetApp Deployment Guidelines for MEDITECH Environments.
<https://fieldportal.netapp.com/content/248456> (NetApp login required)
- NetApp Sizing Guidelines for MEDITECH Environments.
www.netapp.com/us/media/tr-4190.pdf
- FlexPod Datacenter for Epic EHR Deployment
www.netapp.com/us/media/tr-4693.pdf
- FlexPod Design Zone
<https://www.cisco.com/c/en/us/solutions/design-zone/data-center-design-guides/flexpod-design-guides.html>
- FlexPod DC with FC Storage (MDS Switches) Using NetApp AFF, vSphere 6.5U1, and Cisco UCS Manager
https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_esxi65u1_n9fc.html
- Cisco Healthcare
<https://www.cisco.com/c/en/us/solutions/industries/healthcare.html?dtid=osscdc000283>

Acknowledgments

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

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

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.

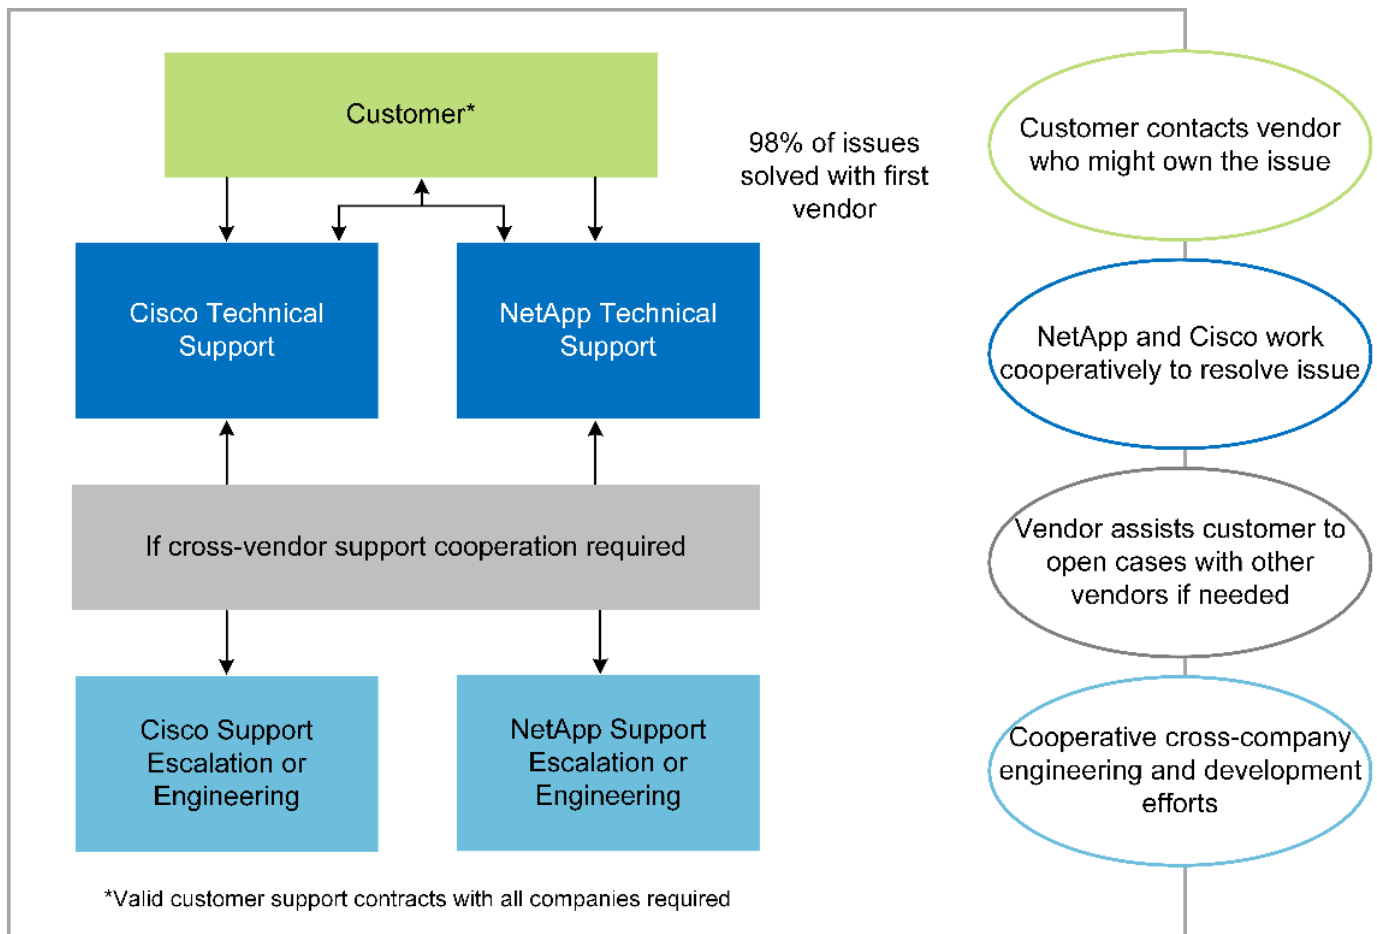
Proven FlexPod architecture and cooperative support

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:

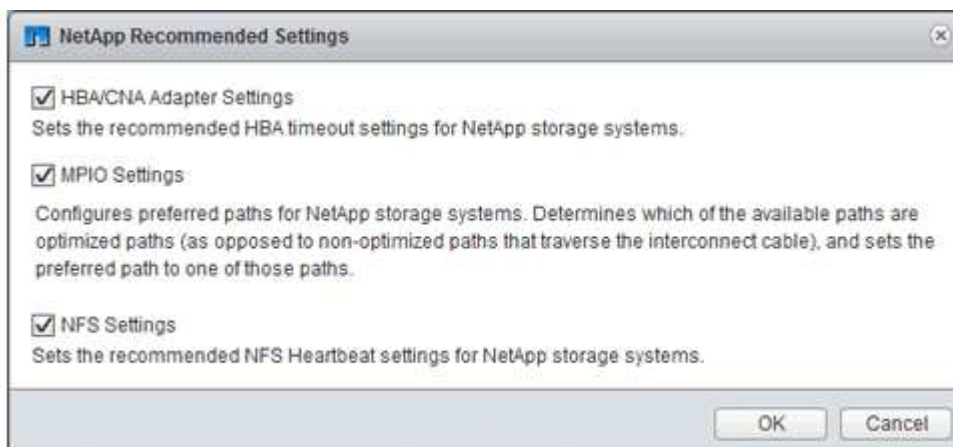
- VSC for VMware vSphere checks and sets the ESXi host multipathing settings and HBA timeout settings that work best with NetApp storage systems. The values that VSC sets are based on rigorous internal testing by NetApp.
- For optimal storage performance, consider using storage hardware that supports 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)

Function	Type
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)
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](#)

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 use 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

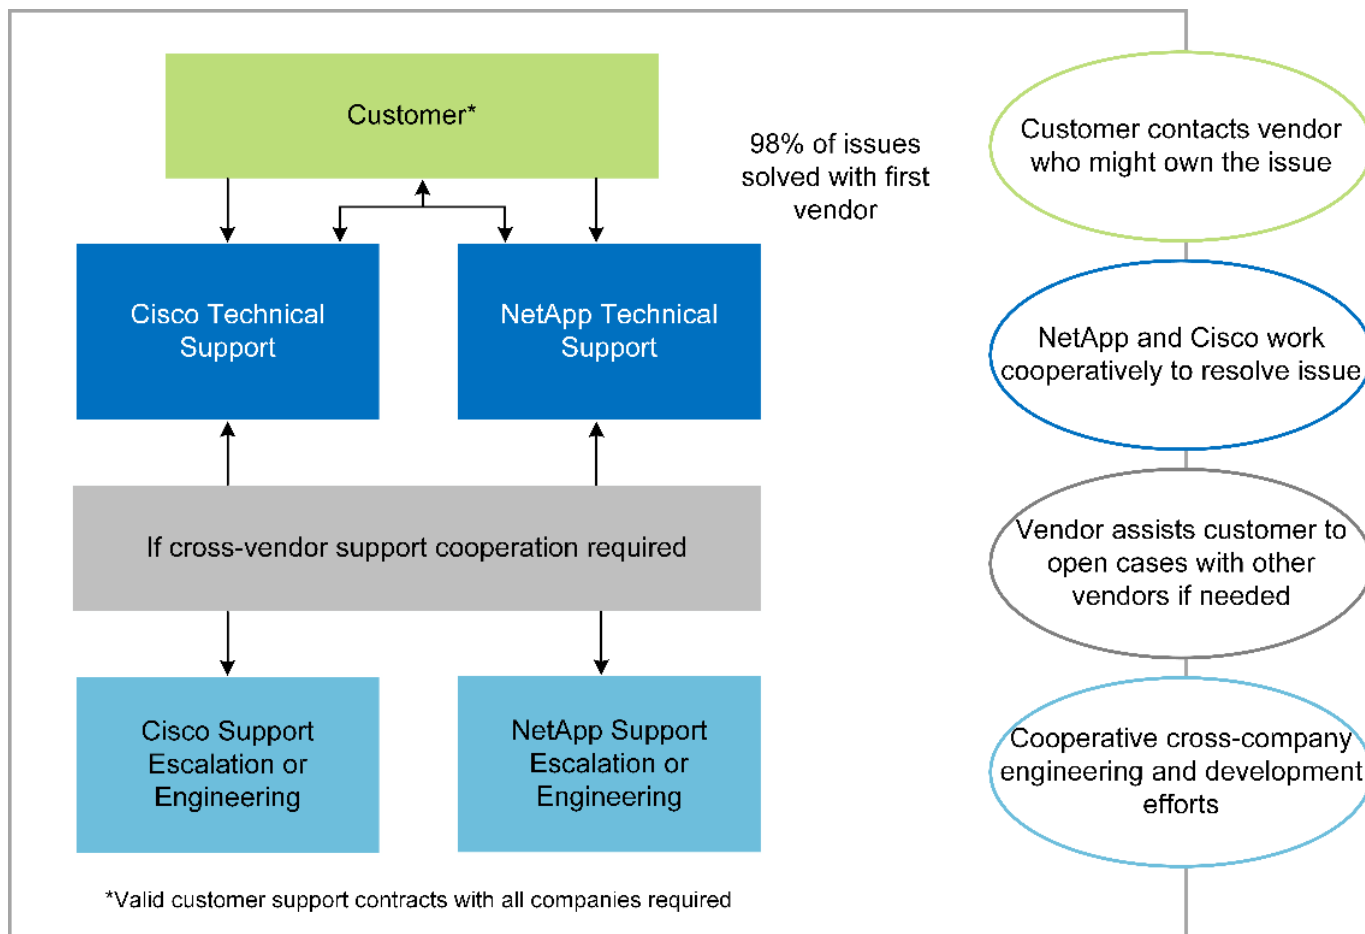
proprietary platforms by replacing them with a more efficient and scalable shared resource that can meet the dynamic needs of your workload. This solution delivers higher infrastructure resource utilization for greater return on investment (ROI).

- **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 as 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:

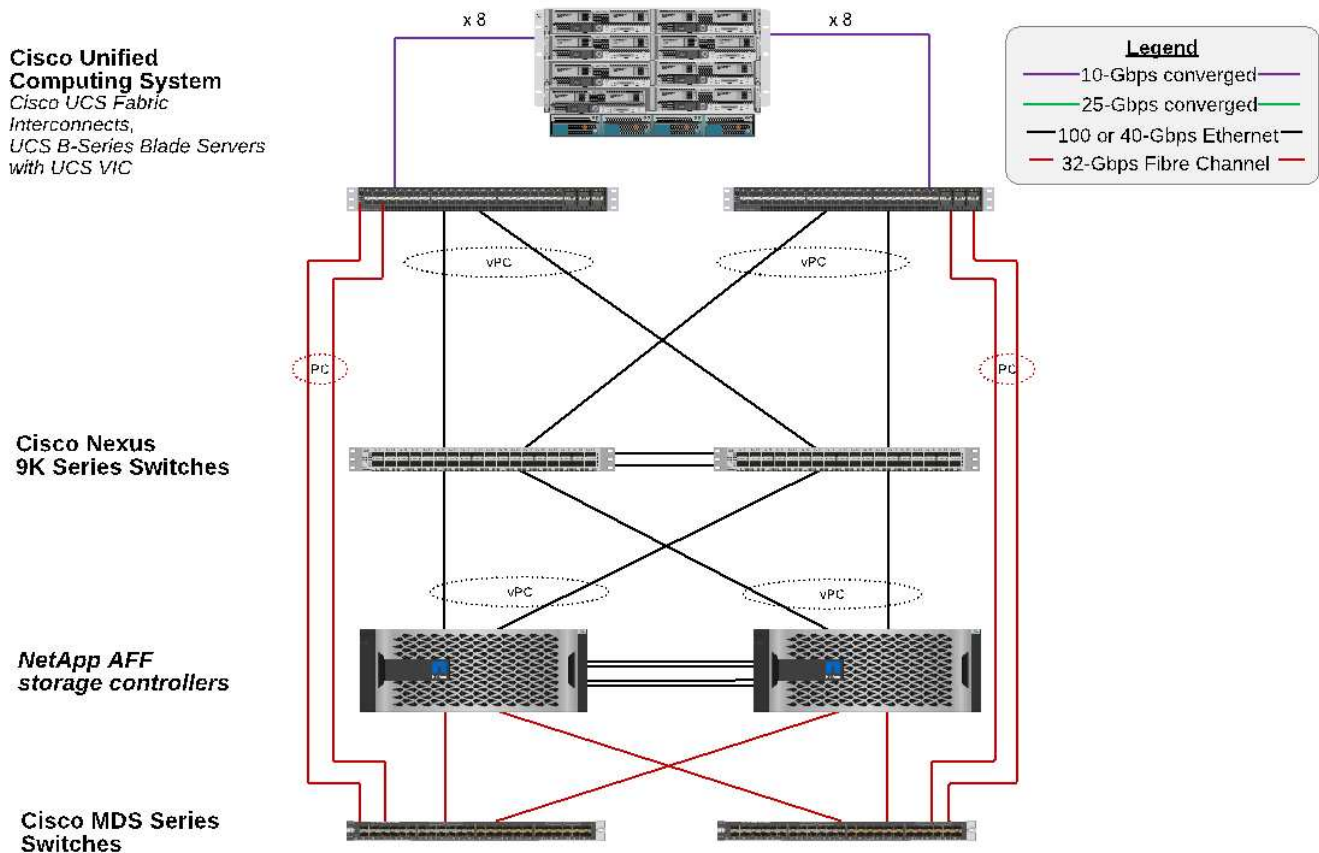
- **Simple deployment.** Quickly and easily deploy vCenter Server by using a virtual appliance.
- **Centralized control and visibility.** Administer the entire vSphere infrastructure from a single location.
- **Proactive optimization.** Allocate, optimize, and migrate resources for maximum efficiency.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

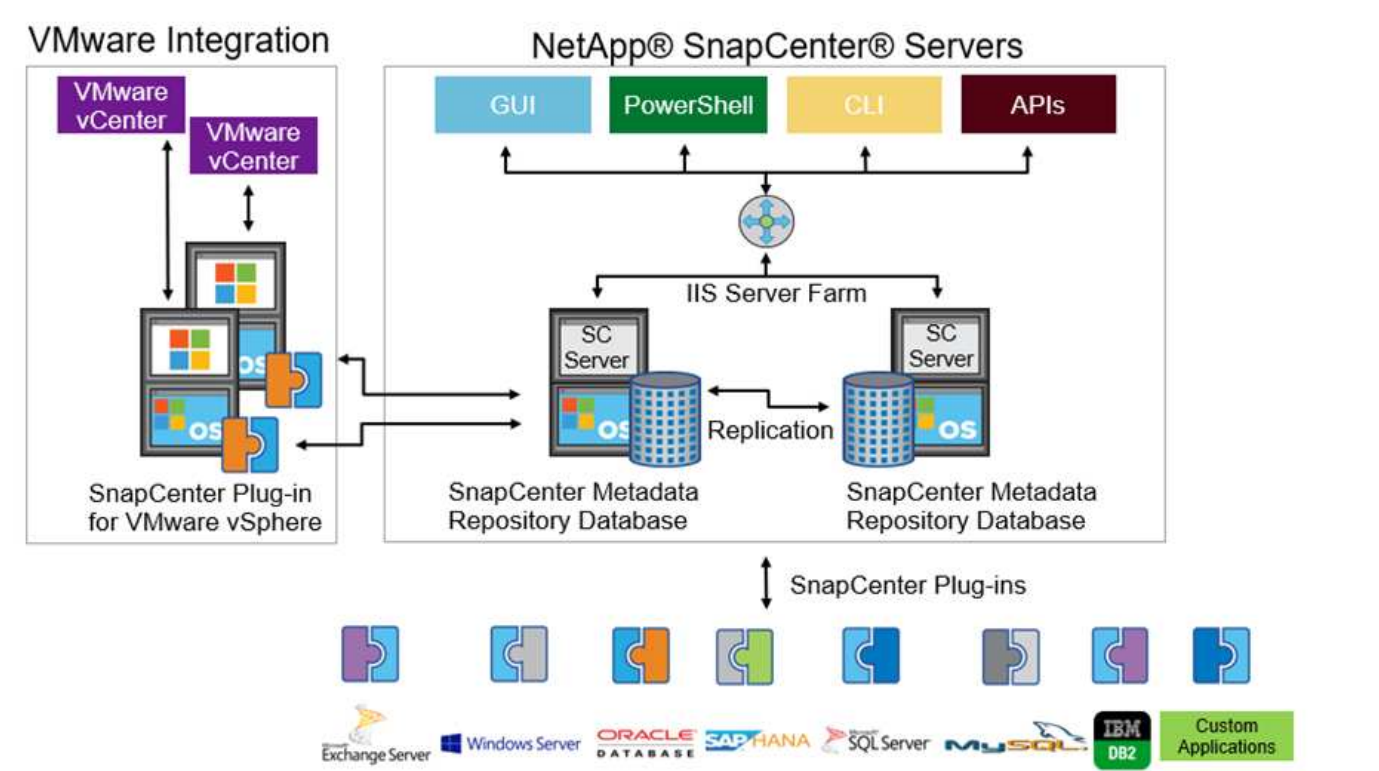
For Oracle database backup, Snapshot copies yield incredible time savings. For example, a backup that took 26 hours to complete by using RMAN alone can take less than 2 minutes to complete by using SnapCenter software.

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.

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

The following figure presents the SnapCenter architecture.



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

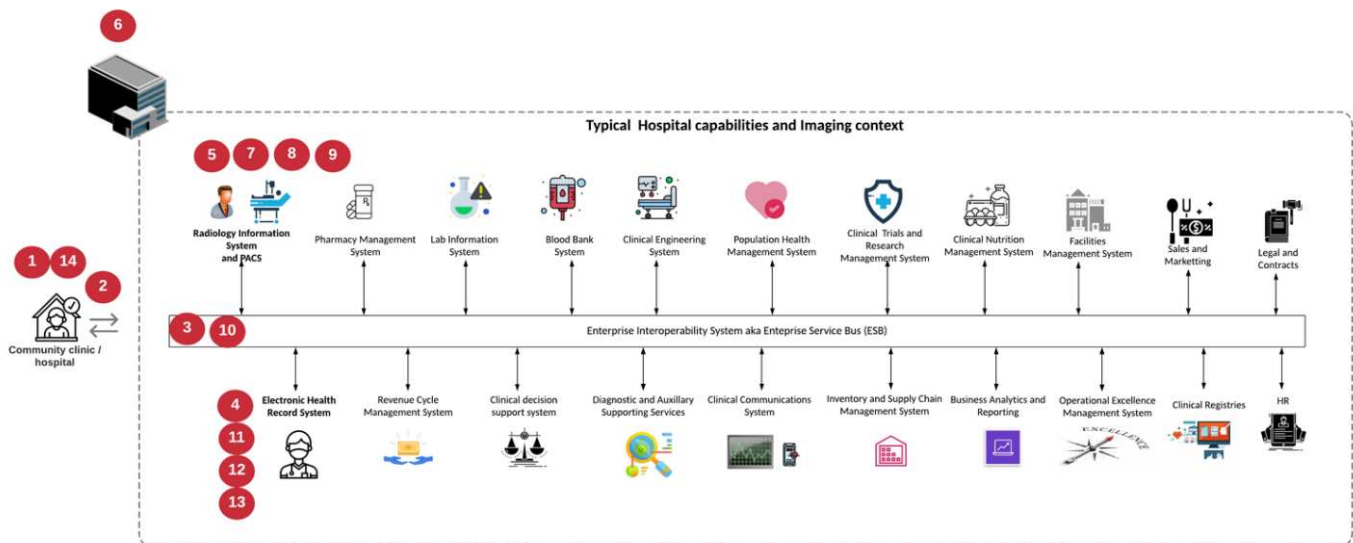
- Standard switch or distributed virtual switch (as needed)
- NFS datastore for workload VMs
- Management, client traffic network, and storage network port groups for VMs
- Network adapter for management, client traffic, and storage access (NFS, iSCSI, or SMB/CIFS) for each VM
- VMware DRS enabled
- 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 use 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/URLs to the DICOM media directly or dynamically generated links/URLs 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.

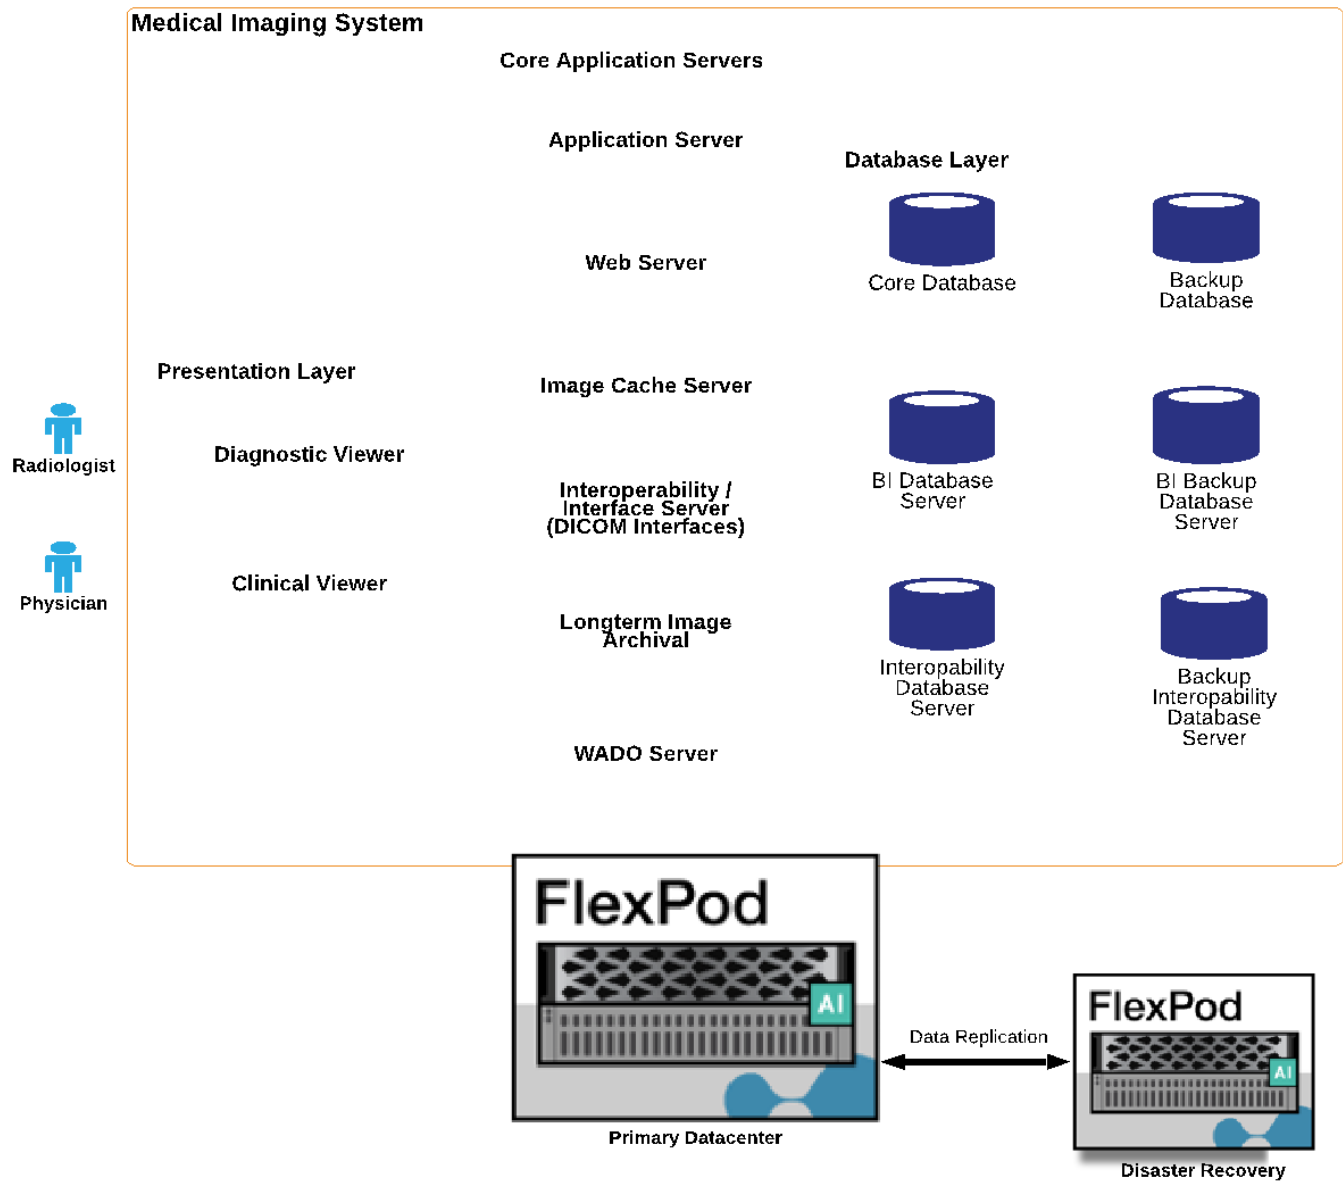
The medical imaging system can provide image access to a community of physicians or to CIN-participating physicians. Typical medical imaging systems include components that enable image interoperability with other health IT systems within and outside of your healthcare organization. Community physicians 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.

	Small	Medium	Large
Annual studies	<250K studies	250K–500K studies	500K–1 million studies
Tier 1 Storage			
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
Tier 2 Storage			
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.

	Small	Medium	Large
Annual studies	<250K studies	250K–500K studies	500K–1 million studies
Data Center 1			
Number of VMs	21	27	35
Total virtual CPU (vCPU) count	56	124	220
Total memory requirement	225GB	450GB	900GB

	Small	Medium	Large
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
Data Center 2			
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.

	Small	Medium	Large
Data Center 1			
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
Data Center 2			
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 ifgroup, 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 ifgroups. 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, management, and backup operations, NetApp recommends the following volume design best practices:

- Isolate databases with I/O-intensive queries throughout the day in different volumes and eventually have separate jobs to back them up.
- For faster recovery, place large databases and databases that have minimal recovery time objectives (RTOs) in separate volumes.
- Consolidate into a single volume your small-to-medium-sized databases that are less critical or that have fewer I/O requirements. When you back up a large number of databases that reside in the same volume, fewer Snapshot copies need to be maintained. NetApp also recommends that you consolidate Oracle database server instances to use the same volumes to control the number of backup Snapshot copies that are created.
- For database replicas, place the data and log files for replicas in an identical folder structure on all nodes.
- Place database files in a single FlexVol; don't spread them across FlexVols.
- Configure a volume auto size policy, when appropriate, to help prevent out-of-space conditions.
- When the database I/O profile consists mostly of large sequential reads, such as with decision support system workloads, enable read reallocation on the volume. Read reallocation optimizes the blocks for better performance.

- For ease of monitoring from an operational perspective, set the Snapshot copy reserve value in the volume to zero.
- Disable storage Snapshot copy schedules and retention policies. Instead, use NetApp SnapCenter Plug-In for Oracle Database to coordinate Snapshot copies of the Oracle data volumes.
- Place user data files and log files on separate FlexVols so that appropriate QoS can be configured for the respective FlexVols and so that different backup schedules can be created.

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

- 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_netappaffa.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

- 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

- 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 © 2021 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means-graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system-without prior written permission of the copyright owner.

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

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

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

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

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

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