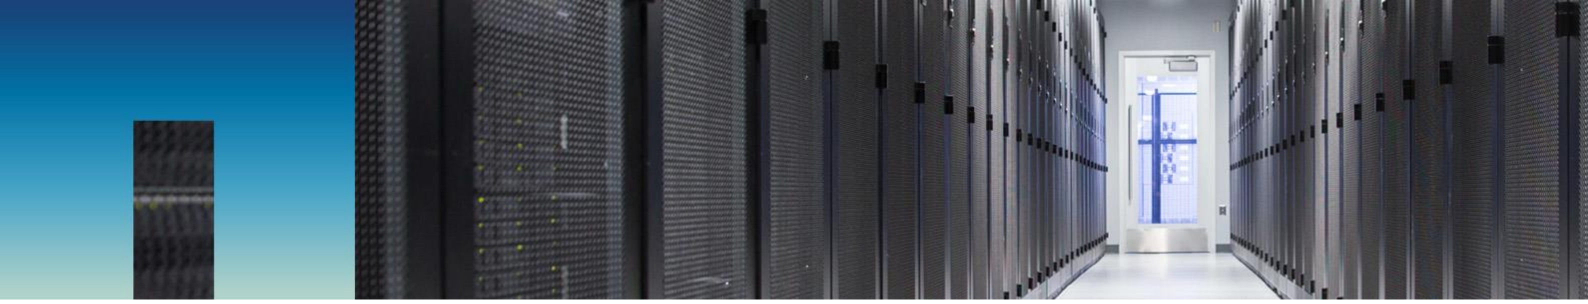

Technical Report

FlexPod for Epic

Directional Sizing Guide

Brian O'Mahony, Ganesh Kamath, Atul Bhalodia, NetApp
Mike Brennan, Jon Ebmeier, Cisco
August 2018 | TR-4707

In partnership with

Abstract

Epic develops software for the healthcare industry. A growing number of healthcare providers are implementing FlexPod®, a next-generation data center platform, to deliver high availability and sustained high performance for Epic EHR application software while increasing infrastructure efficiency and agility. This prevalidated FlexPod converged infrastructure from Cisco and NetApp enables healthcare organizations to improve patient care with a fast, agile, highly scalable, and cost-effective solution.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Introduction

1.1 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:

Section 2, “Reference Architecture,” describes the small, medium, and large compute storage architectures that can be used to host the Epic production database workload.

Section 3, “Technical Specifications,” details a sample bill of materials for the storage architectures described in section 2. The configurations that are described are only for general guidance. Always size the systems according to your workload and tune the configurations as necessary.

1.2 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.[1]
- **Agility.** The industry-leading workflow automation, orchestration, and management tools offered by FlexPod systems allow IT to be far more responsive to business requests. These 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.

1.3 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

1.4 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.

1.5 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

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

2 Reference Architecture

2.1 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.

Note: 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.

Note: 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, it is recommended to see the *Epic All-Flash Reference Architecture Strategy Handbook* for the Epic storage layout requirements.

The three storage reference architectures are:

- **Small.** 4-node architecture with 2 nodes in production and 2 nodes in DR (fewer than 5M global references)
- **Medium.** 6-node architecture with 4 nodes in production and 2 nodes in DR (more than 5M global references)
- **Large.** 12-or-more node architecture with 6 to 10 nodes in production (5M-10M global references)

Note: 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.

Note: 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.

Note: 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 will 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 4 nodes and move to 6 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: 4-Node Reference Architecture for Fewer Than 5M Global References (Up to ~22K Total IOPS)

The small reference architecture is a 4-node architecture with 2 nodes in production and 2 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 NetApp's unique multiprotocol support, 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's 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 of NetApp's best practices. Basically, pool1 runs on node prod-01 and pool2 runs on prod-02, as shown in Figure 1. 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.

Figure 1) Four-node reference architecture for Epic workloads.

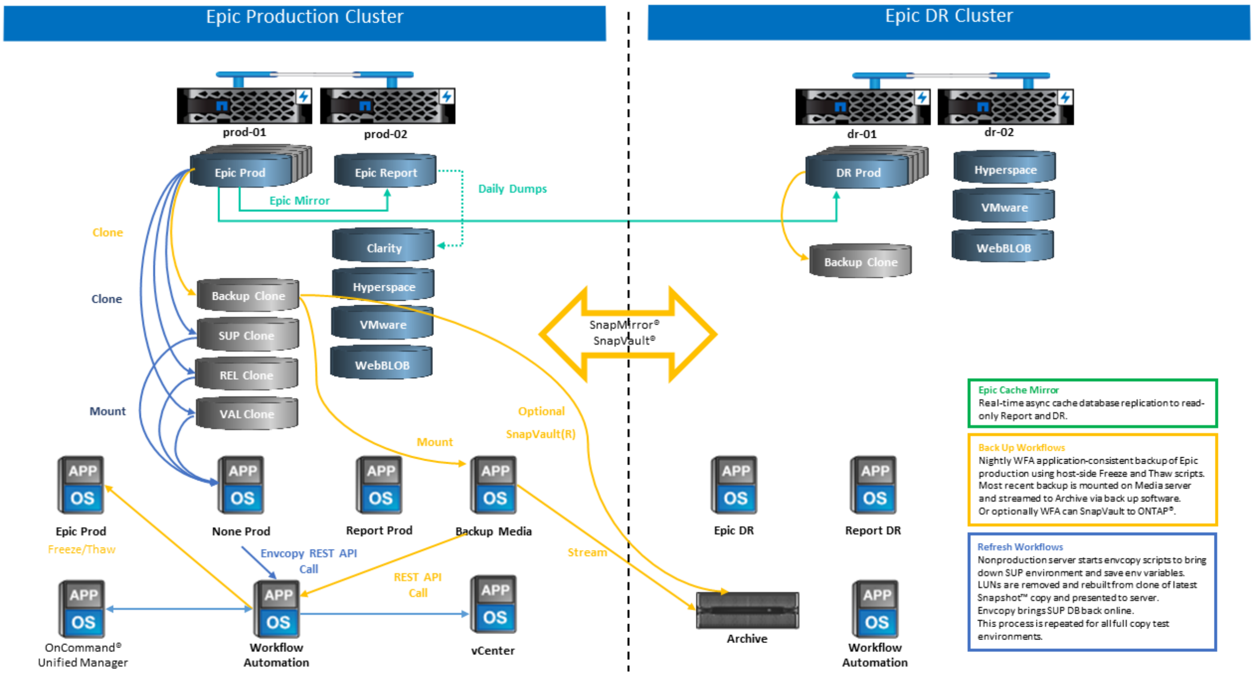

From a storage design and layout perspective, Figure 2 shows a high-level storage layout for the production database and the other constructs that comprise the Epic workload.

Figure 2) Four-node high-level storage design and layout.

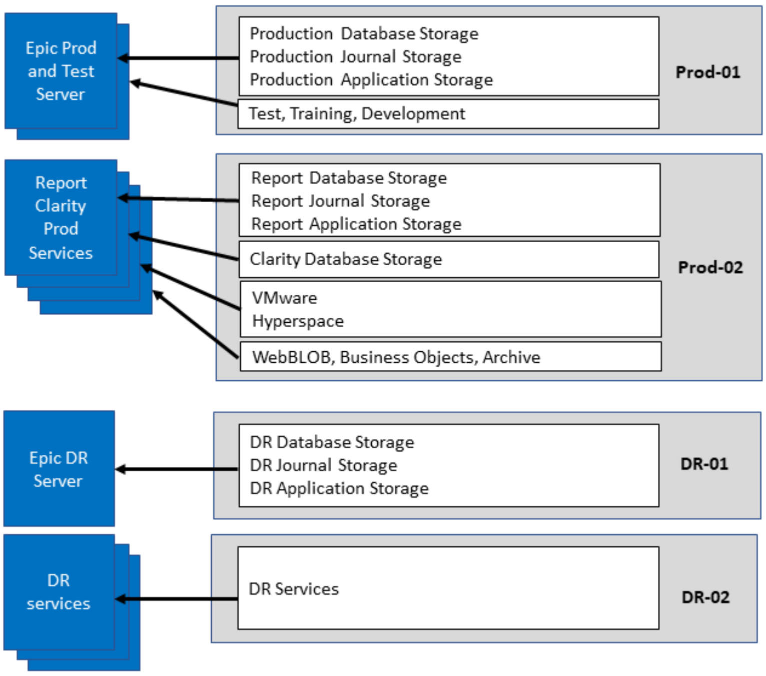

Medium Configuration: 6-Node Reference Architecture for Greater Than 5M Global References (22K-50K Total IOPS)

The medium reference architecture is a 6-node architecture with 4 nodes in production and 2 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 6-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 6 nodes rather than 4 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. Figure 3 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 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.

Figure 3) Six-node reference architecture for Epic workloads.

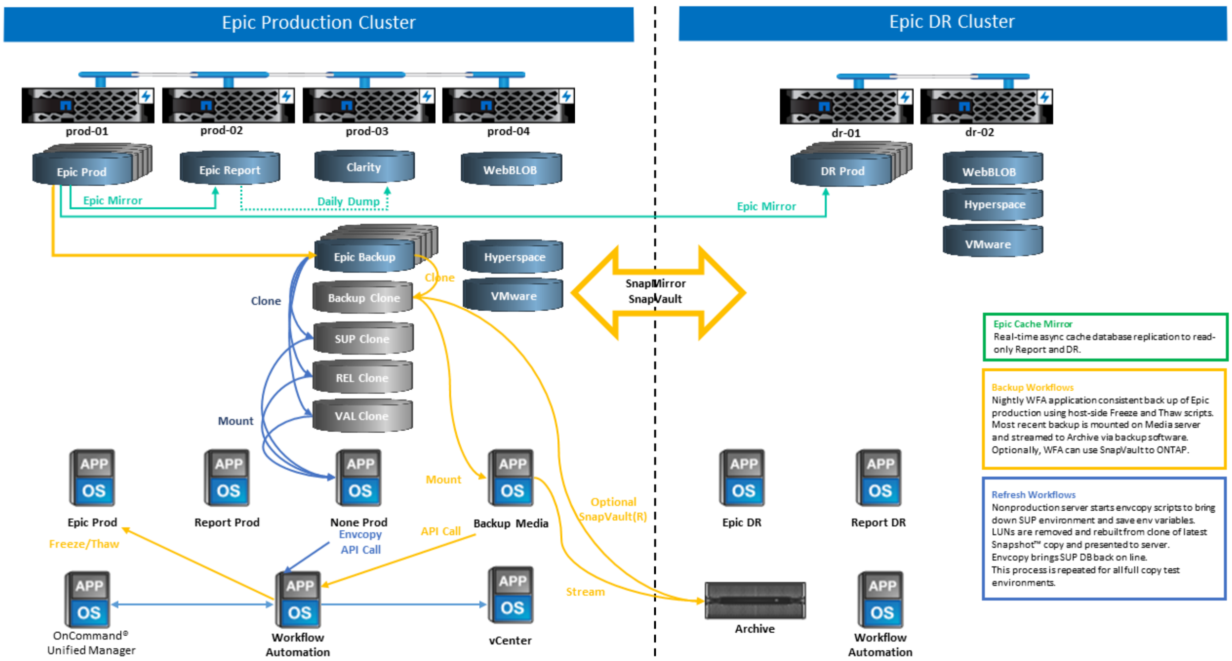

Figure 4 shows the storage layout for a 6-node architecture.

Figure 4) Six-node high-level storage design and layout.

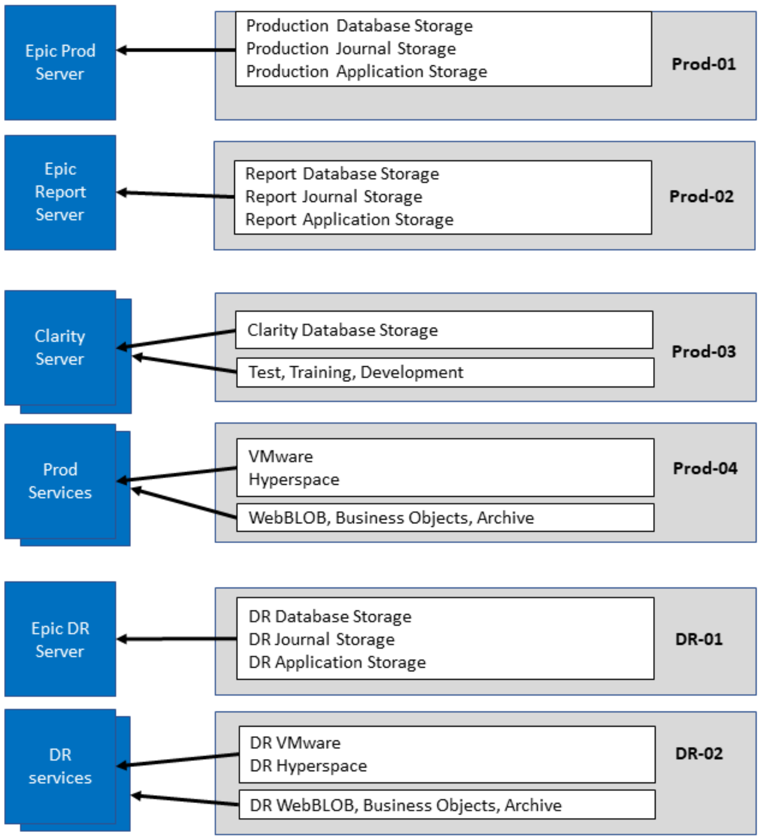

Large Configuration: Reference Architecture for Greater Than 10M Global References (More Than 50K IOPS)

The large architecture is typically a 12-or-more-node architecture with 6 to 10 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 Figure 5.

Larger customers have two options:

- Retain the 6-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.

Figure 5) Large reference architecture for Epic workloads.

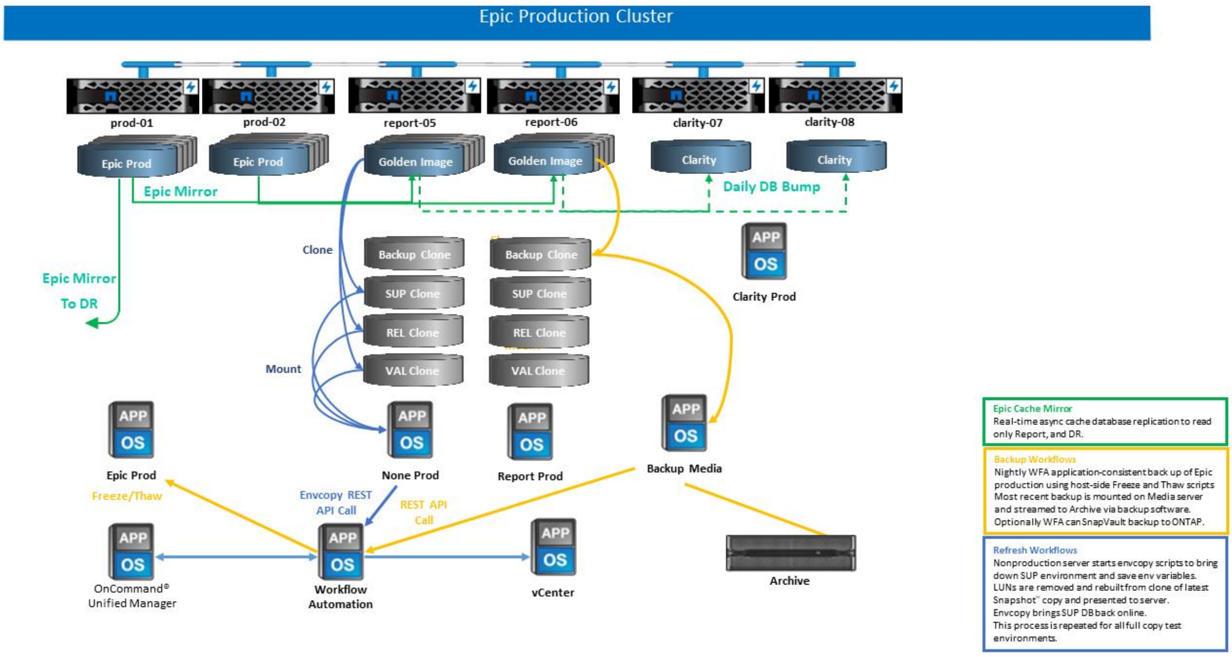

Figure 6 shows the storage layout for a large reference architecture.

Figure 6) Large configuration high-level storage design and layout.

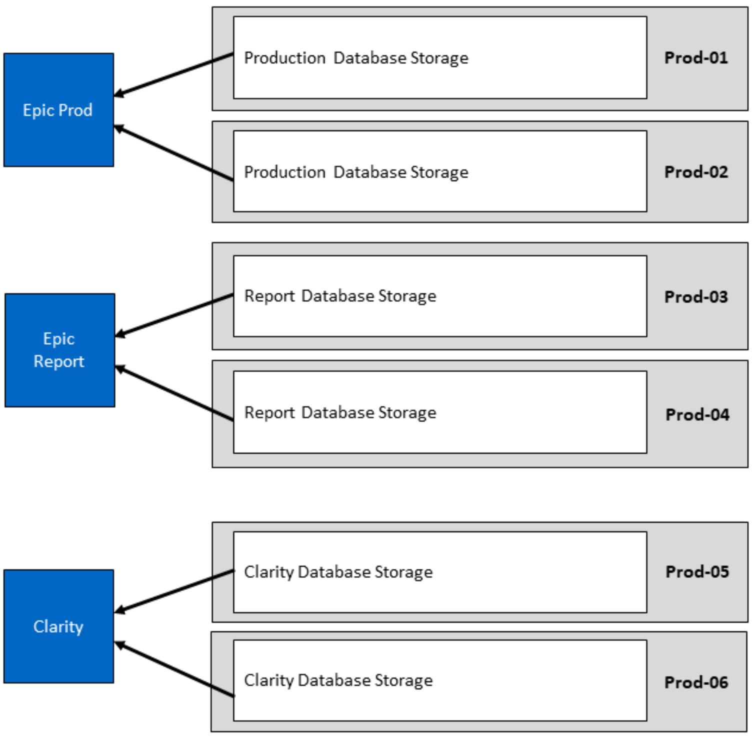

2.2 Cisco UCS Reference Architecture for Epic

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.

Figure 7) Example of a basic small Epic configuration.

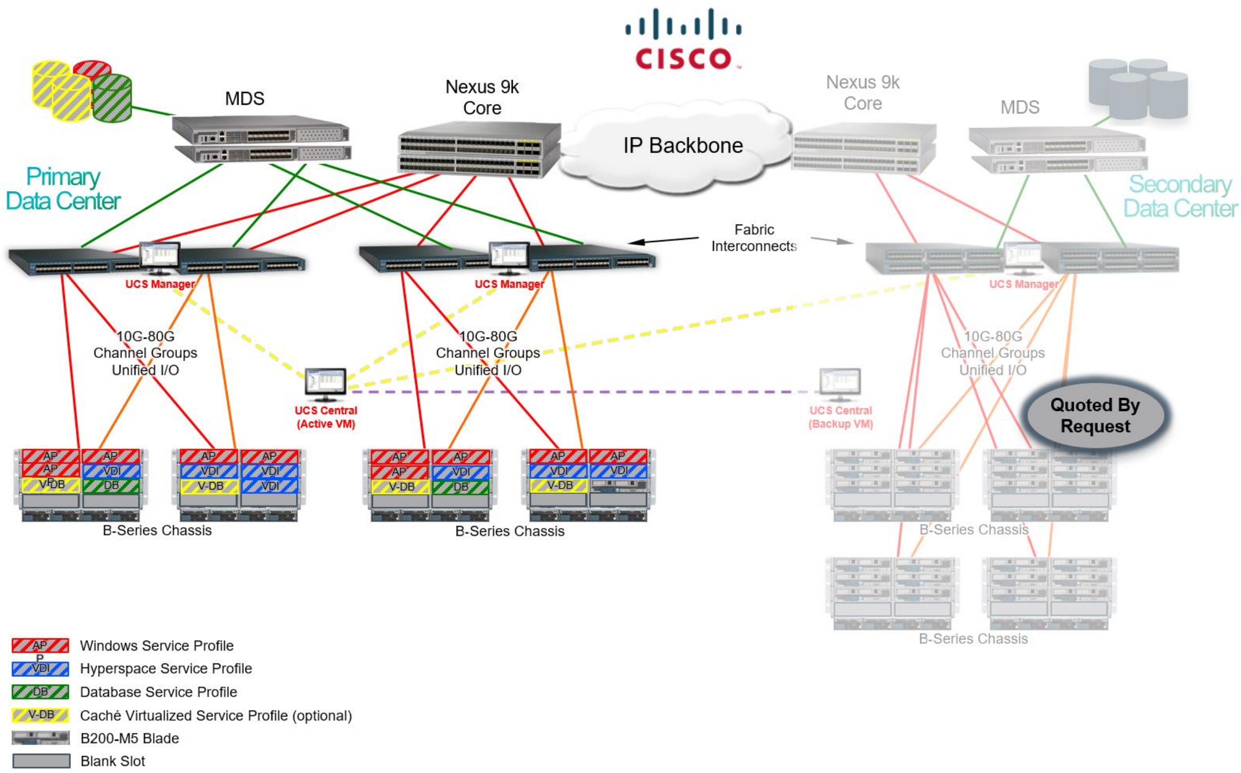

Expanded Design for Larger Implementations

An expanded design for Epic on Cisco UCS follows the same best practices as a basic design. The primary difference is in the scale of the expanded design. With larger scale there is usually a need for higher performance in the core switching, SAN, and processor requirements for Caché databases. Larger implementations typically have more Hyperspace users and need more XenApp for Hyperspace or other virtual application servers. Also, with requirements for more processing power, Cisco UCS quad-socket servers with Intel Skylake processors are used for the Chronicles Caché database and the related Production, Reporting, and Disaster Recovery Caché servers.

Figure 8) Example of an expanded Epic design.

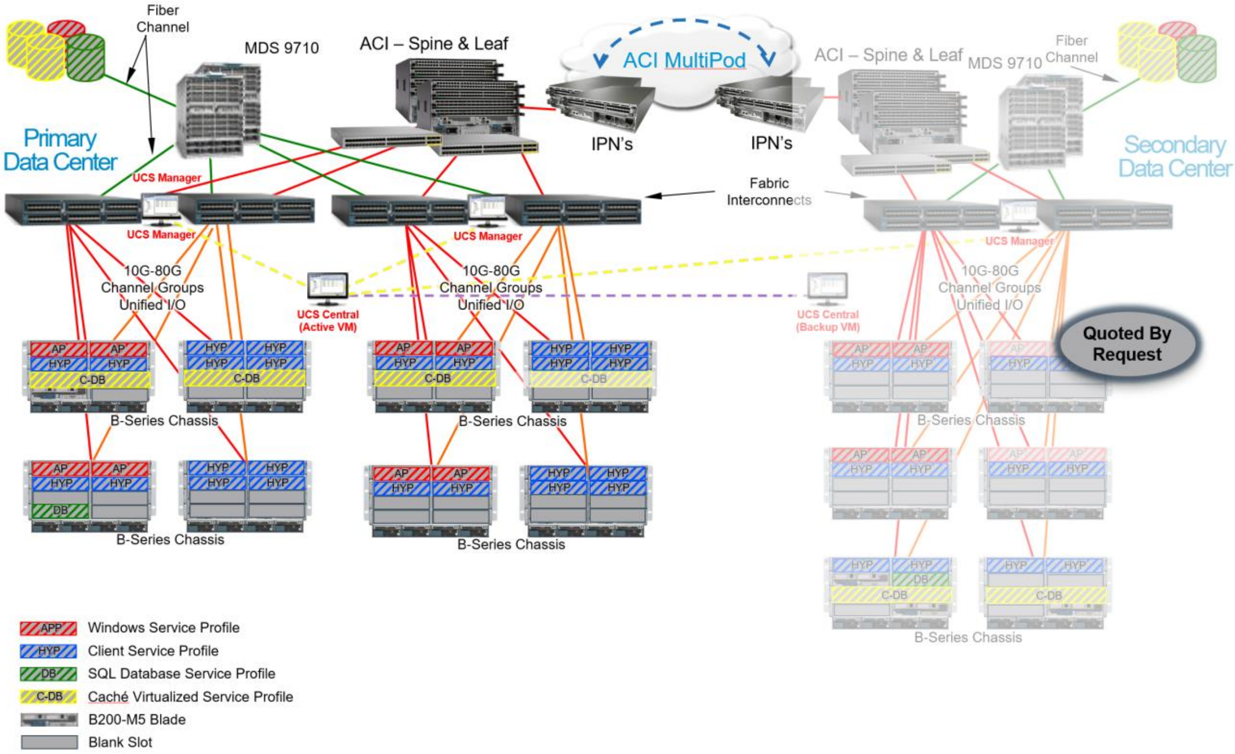

Hyperspace Active–Active Implementations

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 simply 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.

Figure 9) Example of a hyperspace active-active Epic configuration.

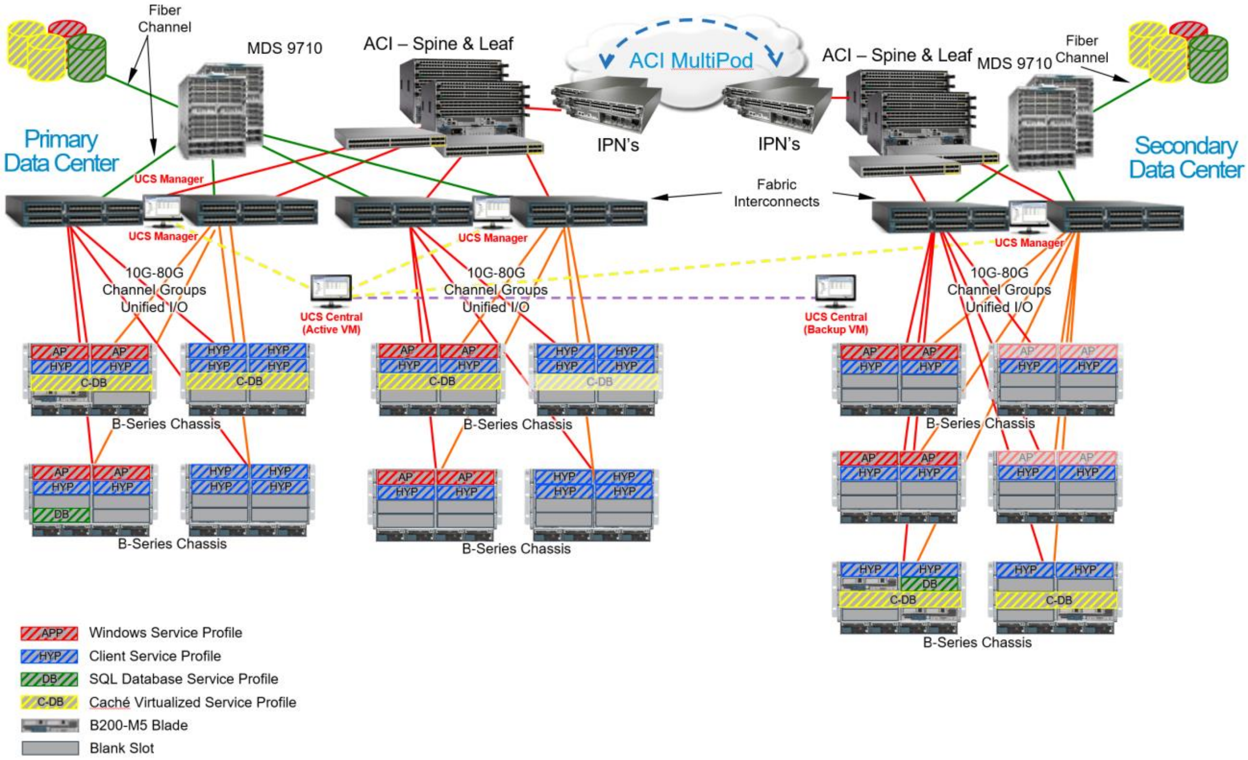

3 Technical Specifications for Small, Medium, and Large Architectures

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.

Table 1 lists the capacity configurations for the Epic production database workload. The total capacity listed accommodates the need for all Epic components.

Table 1) Sample small, medium, and large storage configurations for Epic production database workload.

	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

	Small	Medium	Large
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.

Note: 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.

Table 2 lists the standard software required for the small, medium, and large configurations.

Table 2) Software for 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

Table 3) 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 Cisco UCS Hardware and Software Compatibility matrix)

Layer	Product Family	Quantity and Model	Details
	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
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	

Note: 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.

Table 4) 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 Cisco UCS Hardware and Software Compatibility 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	

Note: 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.

Table 5) 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 Cisco UCS Hardware and Software Compatibility 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	

Note: 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.

Where to Find Additional Information

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

- Detailed deployment of FlexPod Datacenter environment: FlexPod Datacenter with FC Cisco Validated Design
- Overview of Epic software environments, reference architectures, and integration best practices guidance: TR-3928: NetApp Best Practices for Epic
- TR-3930i: NetApp Sizing Guidelines for Epic (access to Field Portal is required to view this document)
- Cisco Best practices with Epic on Cisco UCS: Epic_on_Cisco_UCS_tech_brief
- NetApp FlexPod Design Zone
- FlexPod DC with Fibre Channel Storage (MDS Switches) Using NetApp AFF, vSphere 6.5U1, and Cisco UCS Manager
- TR-4693: FlexPod Datacenter for Epic EHR Deployment Guide

- [NetApp Product Documentation]

Acknowledgements and Version History

- 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.

Version	Date	Document Version History
Version 1.0	August 2018	Initial version.

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

NetApp®