



Technical Report

FlexPod Datacenter for Epic EHR Deployment Guide

Healthcare Industry Solution

Brian O'Mahony and Ganesh Kamath, NetApp
Mike Brennan, Cisco Systems, Inc.
June 2018 | TR-4693

In partnership with



Abstract

FlexPod® provides solutions for the healthcare industry. Epic develops software for the healthcare industry. Healthcare providers increasingly implement 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. The combined strengths of this prevalidated FlexPod converged infrastructure from Cisco and NetApp enables healthcare organizations to improve patient care using a fast, agile, highly scalable and cost-effective solution.

TABLE OF CONTENTS

1	Overview	4
1.1	Abstract	4
1.2	Overall Solution Benefits	4
1.3	FlexPod	5
1.4	Epic Overview	10
1.5	Comprehensive Management Tools and Automation Capabilities	14
2	Design	17
2.1	Storage Architecture	17
2.2	Storage Design and Layout	19
3	Deployment and Configuration	21
3.1	Cabling Diagram	22
3.2	Infrastructure Hardware and Software Components	22
3.3	Base Infrastructure Configuration	23
3.4	Cisco UCS Blade Server and Switch Configuration	24
3.5	ESXi Configuration Best Practices	28
3.6	NetApp Configuration	28
3.7	GenIO Tool	32
	Where to Find Additional Information	32
	FlexPod Design Zone	32
	NetApp Technical Reports	32
	ONTAP Documentation	32
	Cisco Nexus, MDS, Cisco UCS, and Cisco UCS Manager Guides	33
	Acknowledgements and Version History	33

LIST OF TABLES

Table 1)	Test infrastructure hardware and software component details	22
Table 2)	Storage provisioning details	29

LIST OF FIGURES

Figure 1)	FlexPod cooperative support model	7
Figure 2)	How NetApp deduplication works at the highest level	9
Figure 3)	FlexPod for Epic workloads	10
Figure 4)	Epic workloads	11
Figure 5)	Six-node reference architecture example	18
Figure 6)	Six-node high-level storage design and layout example	19
Figure 7)	FlexPod for Epic physical topology diagram	22

Figure 8) Cisco UCSM HTML5 UI26

1 Overview

1.1 Abstract

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.

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

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

Customers want a way to manage data that is consistent and integrated. It provides a way to manage data that puts IT in control and simplifies ever-increasing IT complexity.

1.3 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.

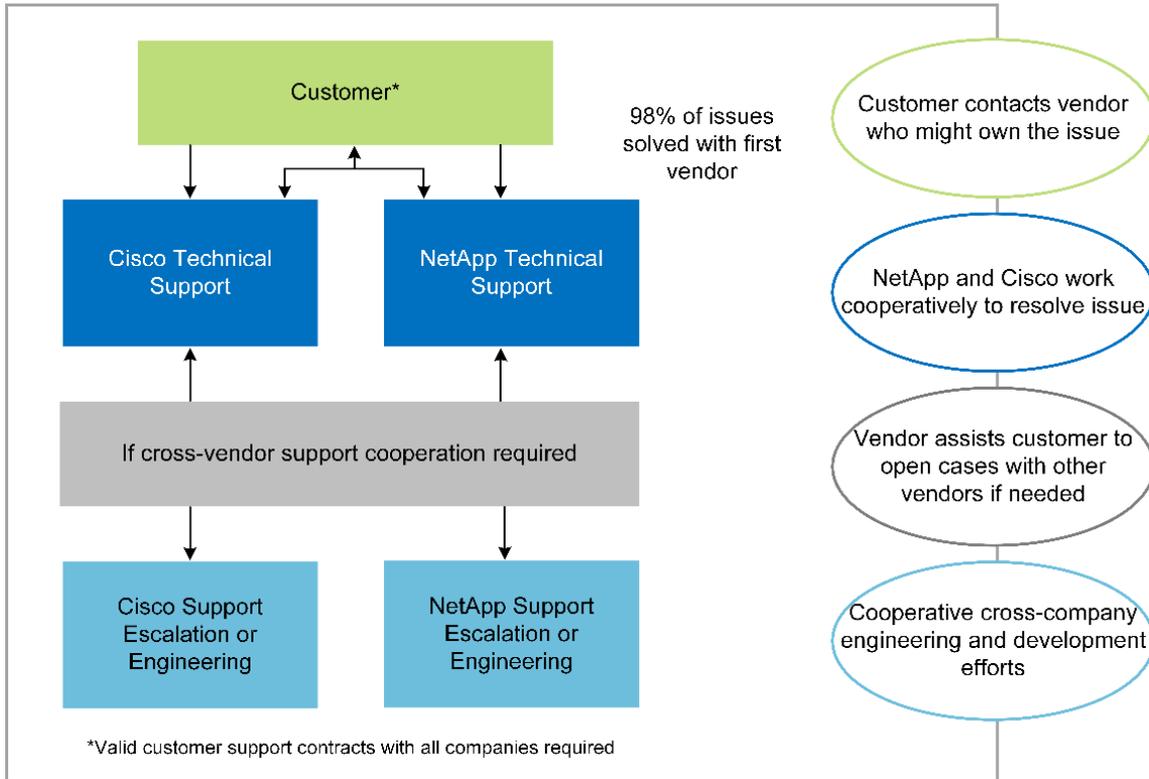
Proven FlexPod Architecture and Cooperative Support

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.

Figure 1) 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.

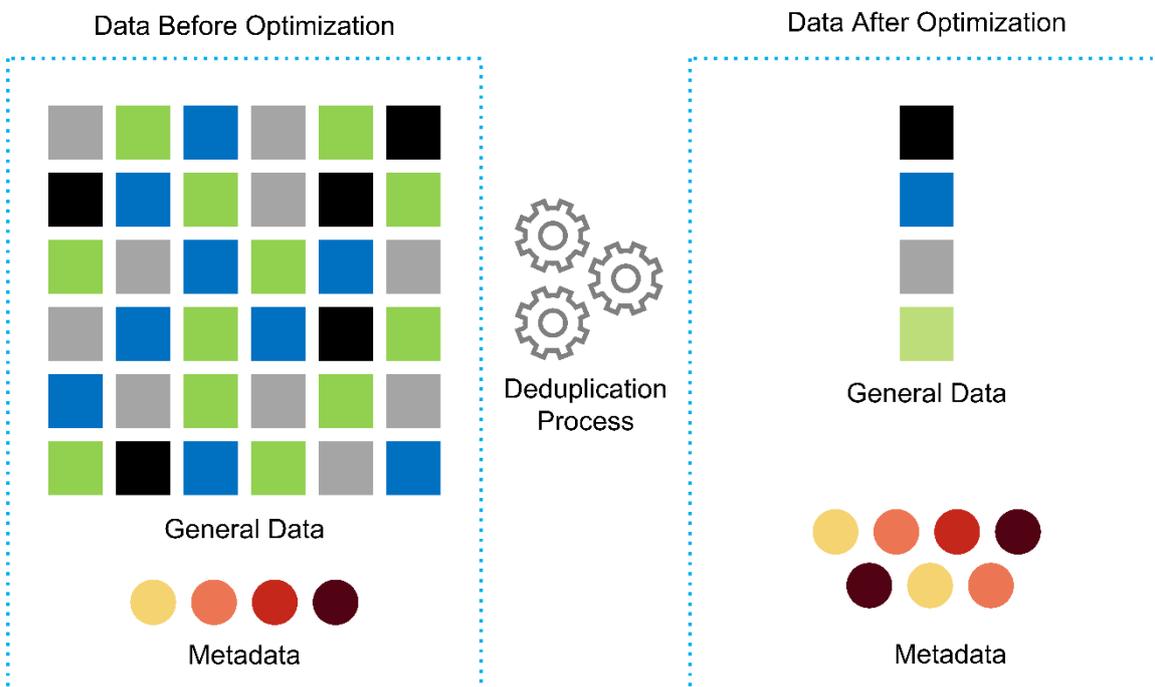
ONTAP offers features that are extremely useful in Epic environments, simplifying management, increasing availability and automation, and reducing the total amount of storage needed:

- **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.

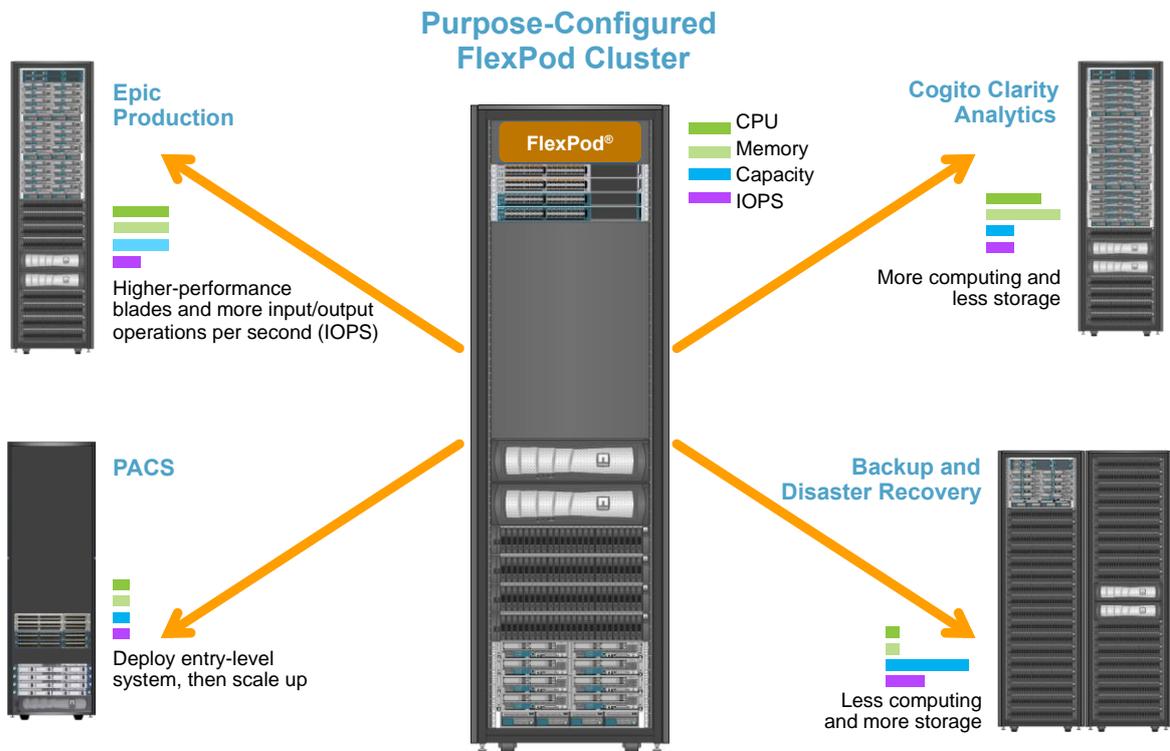
Figure 2) How NetApp deduplication works at the highest level.



- **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.

Figure 3) FlexPod for Epic workloads.



1.4 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. Figure 4 outlines five distinct modules, each with unique workload characteristics.

Figure 4) Epic workloads.



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.

Note: 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.

1.5 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 NetApp 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.

2 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.

2.1 Storage Architecture

Figure 5 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.

Note: 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.

The six-node architecture contains four nodes for production and two nodes for DR. With this architecture, with four-node production, the Epic All-Flash Reference Architecture Strategy Handbook states that you can separate Epic report workloads from Clarity.

Going with six nodes has the following key advantages:

- You can offload backup archive process from production
- You can offload all test environments from production

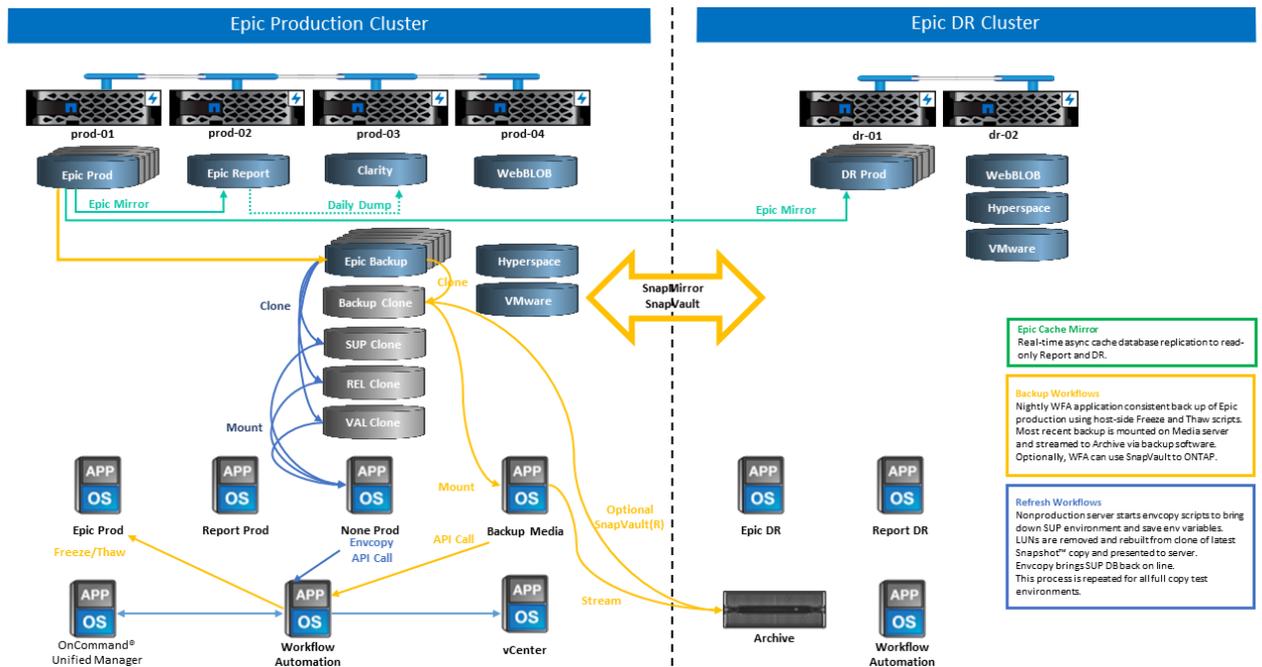
Production runs on node prod-01. Report runs on node prod-02, which is an up-to-the-minute Epic mirror copy of production. Test environments such as support, release, and release validation (SUP, REL, and RELVAL) can be cloned instantaneously from either Epic production, report, or DR. Figure 5 shows 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. SnapMirror backup copies can be used to create 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 made from the report Snapshot copy backup and run directly on node 2. With this design, a SnapMirror destination copy is not required to be saved on disk.

Figure 5) Six-node reference architecture example.



2.2 Storage Design and Layout

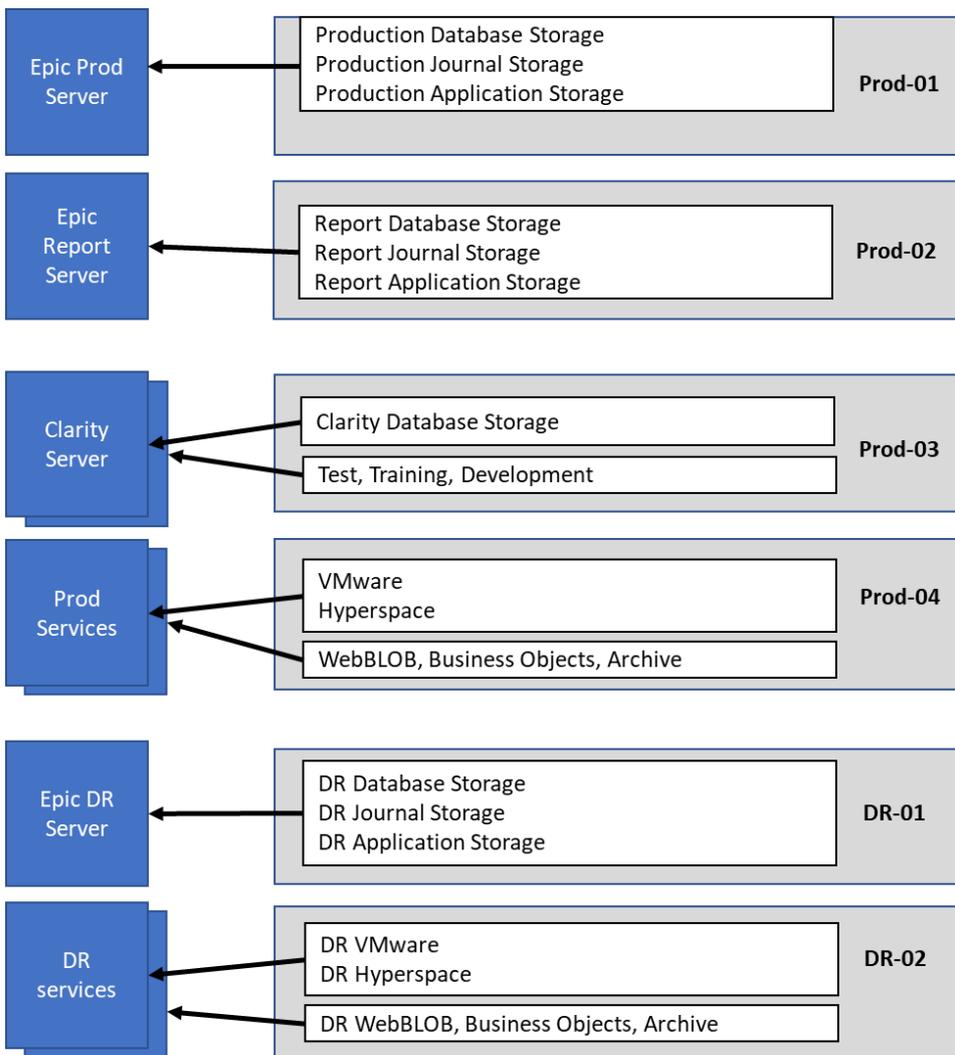
The first step toward satisfying Epic's HA and redundancy requirements is to design the storage layout specifically for the Epic software environment, including isolating disk pool 1 from disk pool 2 onto dedicated high-performance storage. See the Epic All-Flash Reference Architecture Strategy Handbook for information about what workloads are in each disk pool.

Placing each disk pool on a separate node creates the fault domains required for 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.

Figure 6 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.

Figure 6) Six-node high-level storage design and layout example.



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.

Note: 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.

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

3 Deployment and Configuration

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

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

This document does not cover:

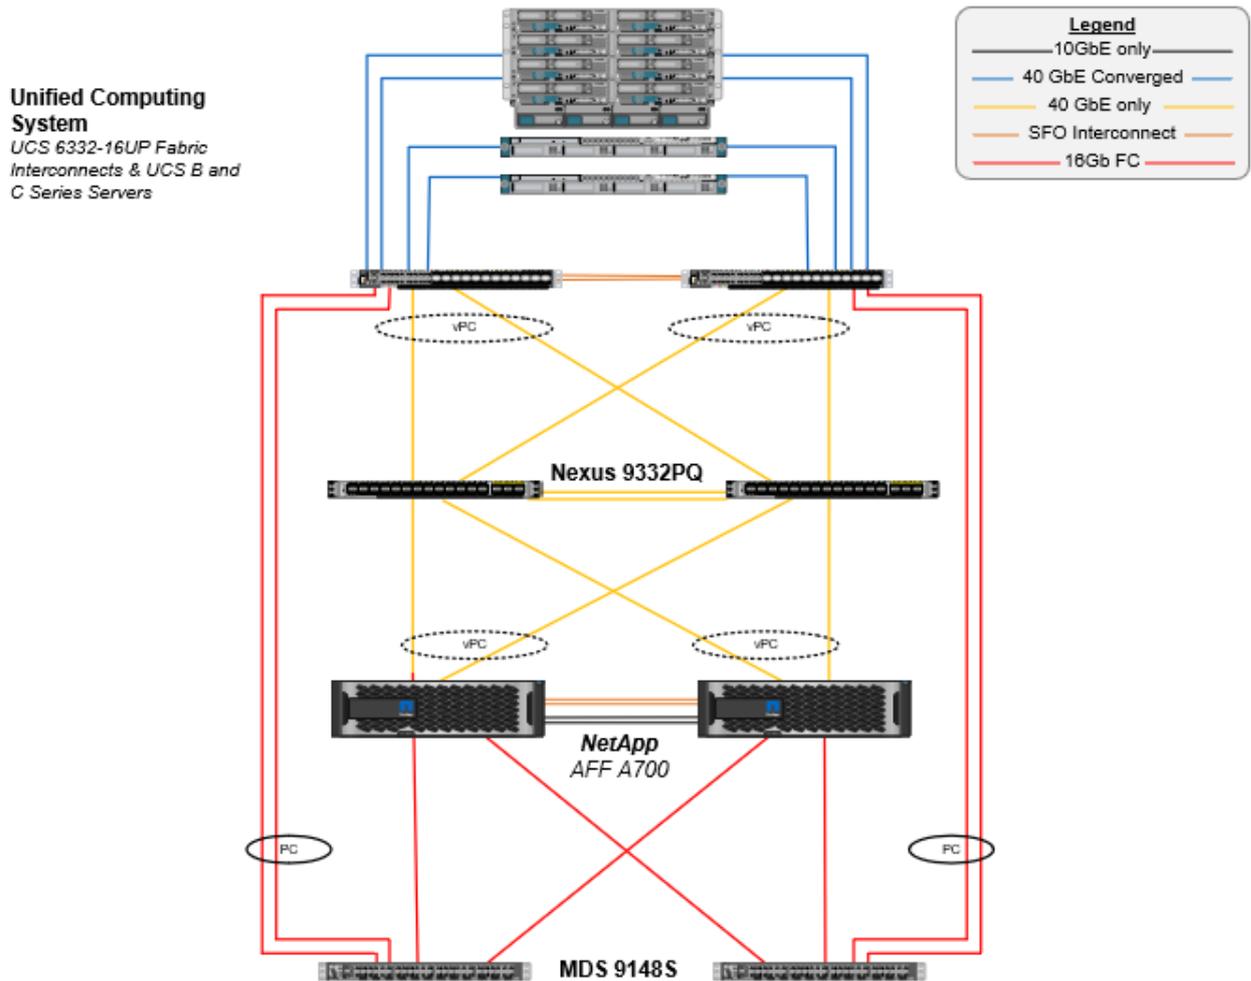
- Detailed deployment of FlexPod Datacenter environment. See [FlexPod Datacenter with FC Cisco Validated Design](#).
- Overview of Epic software environments, reference architectures, and integration best practices guidance. See NetApp [TR-3928: NetApp Best Practices for Epic](#).
- Quantitative performance 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.

3.1 Cabling Diagram

Figure 7 illustrates the 16Gb FC/40GbE topology diagram for an Epic deployment.

Figure 7) FlexPod for Epic physical topology diagram.



3.2 Infrastructure Hardware and Software Components

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

Table 1) Test infrastructure hardware and software component details.

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) *	

3.3 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.

3.4 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 section 3.1, “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 per [FlexPod Datacenter with FC Cisco Validated Design](#).

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 section 3.1, “Cabling Diagram.”

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

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

2. Configure individual ports, port channels, and descriptions as per the FlexPod Cisco MDS switch configuration section in [FlexPod Datacenter with FC Cisco Validated Design](#).
3. To create the necessary VSANs for the Epic solution, complete the following steps while in global configuration mode:
 - For 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 fcl/1
vsan <vsan-a-id> interface fcl/2
vsan <vsan-a-id> interface port-channel110
vsan <vsan-a-id> interface port-channel112
```

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

- a. For 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 fcl/1
vsan <vsan-b-id> interface fcl/2
vsan <vsan-b-id> interface port-channel111
vsan <vsan-b-id> interface port-channel113
```

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

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

```

– For 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

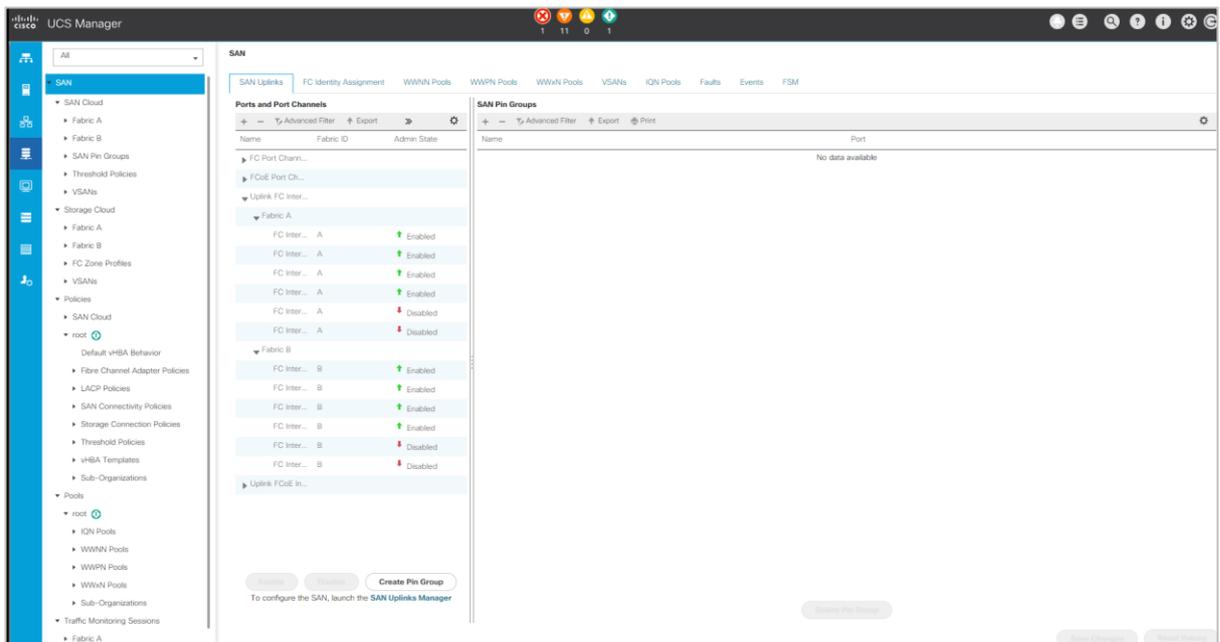
Cisco UCS allows Epic customers to leverage their subject matter experts in network, storage, and compute to create policies and templates that tailor the environment to their specific needs. After being created, these policies and templates can be combined into service profiles that deliver consistent, repeatable, reliable, and fast deployments of Cisco blade and rack servers.

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

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

Figure 8 shows a sample screenshot of the SAN node in Cisco UCS Manager.

Figure 8) Cisco UCSM HTML5 UI.



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 section 3.1, “Cabling Diagram.”

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

3.5 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:

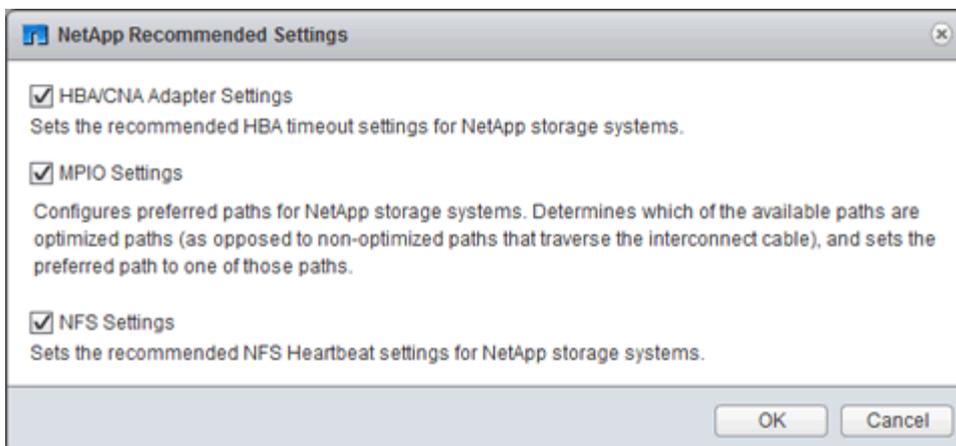
- Virtual Storage Console (VSC) for VMware vSphere checks and sets the ESXi host multipathing settings and HBA timeout settings that work best with NetApp storage systems. The values that VSC sets are based on rigorous internal testing by NetApp.
- For the best storage performance, customers should consider 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.

3.6 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 Table 2.

Table 2) Storage provisioning details.

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_report_db1/epic_report_db1 (1.4TB) /epic_report_db2/epic_report_db2 (1.4TB) /epic_report_db3/epic_report_db3 (1.4TB) /epic_report_db4/epic_report_db4 (1.4TB) /epic_report_db5/epic_report_db5 (1.4TB) /epic_report_db6/epic_report_db6 (1.4TB) /epic_report_db7/epic_report_db7 (1.4TB) /epic_report_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)

Parameter	Controller 1	Controller 2
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
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 for Epic workloads would be FC, CIFS, Snapshot, SnapRestore®, FlexClone, and SnapMirror.

1. 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, complete the following step:

1. Run the following Data ONTAP commands to configure AutoSupport:

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

Hardware-Assisted Takeover Configuration

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

1. Run the following ONTAP command:

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

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

2. Run the following ONTAP command:

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

```
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).

Note: 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.

Note: 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.

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

Note: 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.

3.7 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

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

NetApp Technical Reports

- [TR-3929: Reallocate Best Practices Guide](#)
- [TR-3987: Snap Creator Framework Plug-In for InterSystems Caché](#)
- [TR-3928: NetApp Best Practices for Epic](#)
- [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 Red Hat Enterprise Linux Express 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 9148S FC Switch](#)

Acknowledgements and Version History

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

Version	Date	Document Version History
Version 1.0	May 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.

Copyright Information

Copyright © 2018 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark Information

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

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

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

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

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