

Best Practices for VMware vSAN with Epic

Table of contents

Executive Summary	3
Overview and Purpose	3
Key Takeaways	3
Audience	3
Technology Overview	4
VMware vSphere 6.7 U3	4
VMware vSAN 6.7 U3	4
InterSystems Caché and IRIS	4
Microsoft SQL Server 2016	5
Solution Configuration	5
Hardware Resources	5
Network Configuration	6
Architecture Diagram	7
vSAN Storage Policy Configuration	9
Storage Policies and Epic EHR Workloads	11
Software Resources	11
Operational Database Caché Linux VM Best Practices and Layout	12
SQL Cogito VM Best Practices and Layout	13
Solution Validation	13
Conclusion	14
About the Author and Contributors	14

Executive Summary

A typical EHR environment usually consists of several different storage arrays coupled with dozens of hosts. Healthcare IT departments must not only administer and maintain this environment, they must also publish new applications, provision new hardware, and investigate new technologies. Given these mandates, many Healthcare IT CTOs are looking to lower costs. Expanding IT departments to host large scale data centers isn't a goal for most hospitals. This presents a challenge to most healthcare providers today: how does a hospital provide its clinicians and patients cutting-edge technology on a robust and responsive infrastructure without increasing costs and impacting patient care?

VMware can answer the challenge by hosting these core business functions on VMware vSAN™. Adopting vSAN minimizes the need to maintain legacy infrastructure and is capable of delivering mission-critical response times. vSAN is enterprise-class, storage virtualization software that, when combined with *vSphere*, allows administrators to manage compute and storage with a single platform. There is no need to deploy or maintain separate arrays and storage networking hardware with VMware vSAN. Its policy-based management removes the burden of provisioning and modifying numerous LUNs and data services. With its consistency and flexibility, VMware vSAN architecture provides the simplest path from server virtualization to hyperconverged infrastructure and a true *hybrid cloud* architecture.

In this document, we showcase VMware best practices and design guidelines for the Epic Operational and Analytical databases on VMware vSAN.

Overview and Purpose

VMware and Epic have conducted the initial testing that has indicated that vSAN is acceptable performance at small to medium size environments.

This best practice guide is focused on deploying and managing Epic Operational and Analytical Databases on VMware vSAN. These best practices are determined by:

- Following Epic Operational Database, and Cogito storage best practices.
- Using Epic provided test tools to simulate ODB and Cogito production workloads.
- Determining the optimal performance and configuration for Operational Database and Cogito in a single 6-node vSAN cluster.

For the most recent status of VMware vSAN with Epic, contact your Epic representative for the most current Storage Products and Technology Status Guide.

Key Takeaways

- VMware vSAN can deliver predictable and consistent performance for Epic Operational and Analytical Databases in well prescribed architectures and scenarios.
- A 6-node vSAN cluster, Ready Nodes with (2) five-disk-group configuration, is capable of hosting both the Epic Operational and Analytical databases.
- vSAN storage policy-based management provided the key to consistent performance during run-away report testing.
- A non-blocking, high-buffer count 10/25GbE or greater switch is required.
- Hosting both databases on a single 6-node cluster provides cost savings in both administration and acquisition the infrastructure.
- Consult with your VMware and Epic representatives prior to procuring any hardware or starting your vSAN and Epic project. Selecting the wrong hardware may result in an unsupported vSAN configuration and possibly impact performance.
- Refer to your Epic Hardware Configuration Guide for hardware and storage specifications tailored to your organization.

Audience

This document is intended for architects, application developers, and CTOs in the healthcare ecosystem who are involved in the early phases of planning, design, and deployment of Epic in their environment or upgrading their existing Epic infrastructure. It is assumed that the reader is familiar with the concepts and operations of running Epic software.

Technology Overview

This chapter describes the software components discussed in this Best Practice Guide. Epic maintains a list of specific recommended platforms that they have focused development, functional and performance testing, and release efforts on for each version of Epic, called the Target Platform. For a more up-to-date list of recommended software components, refer to the most recent Epic Target Platform Roadmap.

- VMware vSphere 6.7 U3
- VMware vSAN 6.7 U3
- InterSystems Caché 2018.1.2 and IRIS 2019.1
- Microsoft SQL Server 2016

VMware vSphere 6.7 U3

VMware vSphere 6.7 is the next-generation infrastructure for next-generation applications. It provides a powerful, flexible, and secure foundation for business agility that accelerates the digital transformation to cloud computing and promotes success in the digital economy. vSphere 6.7 supports both existing and next-generation applications through its:

- Simplified customer experience for automation and management at scale
- Comprehensive built-in security for protecting data, infrastructure, and access
- Universal application platform for running any application anywhere

With vSphere 6.7, customers can run, manage, connect, and secure their applications in a common operating environment, across clouds and devices.

VMware vSAN 6.7 U3

VMware vSAN is the industry-leading software powering VMware's software defined storage and HCI solution. vSAN helps customers evolve their data center without risk, control IT costs and scale to tomorrow's business needs. vSAN, native to the market-leading hypervisor, delivers flash-optimized, secure storage for all of your critical vSphere workloads. vSAN is built on industry-standard x86 servers and components that help lower TCO in comparison to traditional storage. It delivers the agility to easily scale IT and offers the industry's first native HCI encryption.

In vSAN 6.7 U3 release, it provides performance improvements and availability SLAs on all-flash configurations with deduplication enabled. Latency sensitive applications have better performance in terms of predictable I/O latencies and increased sequential I/O throughput. Rebuild times on disk and node failures are shorter, which provides better availability SLAs.

The 6.7 U3 release also support cloud native storage that provides comprehensive data management for stateful applications. With Cloud Native Storage, vSphere persistent storage integrates with Kubernetes.

vSAN 6.7 U3 simplifies day-1 and day-2 operations, and customers can quickly deploy and extend cloud infrastructure and minimize maintenance disruptions. Stateful containers orchestrated by Kubernetes can leverage storage exposed by vSphere (vSAN, VMFS, NFS) while using standard Kubernetes volume, persistent volume, and dynamic provisioning primitives.

InterSystems Caché and IRIS

Epic licenses and uses InterSystems Caché and IRIS for the operational database to store and manage patient records, IRIS being the newer database. Within Caché and IRIS, data can be modeled and stored as tables, objects, or multidimensional arrays. Different models can seamlessly access data—without the need for performance—killing mapping between models. All three access methods can be simultaneously used on the same data with full concurrency. Thus, making it ideal for use in hospital environments where thousands of clinicians could access and manipulate the same data.

Epic leverages many features and functions from Caché and IRIS, however, one important note are the two main deployment architectures that Epic uses:

- Symmetric Multiprocessing (SMP)—The most commonly deployed architecture for Epic customers. The data server is accessed directly.

- Enterprise Cache Protocol (ECP)—A tiered architecture in which users access the data server via a pool of application servers. All data still resides on the data server. ECP is used in some of Epic’s largest customers.

Microsoft SQL Server 2016

Epic Cogito can use either Oracle or Microsoft SQL Server for the underlying database. In our solution we use Microsoft SQL Server 2016 which enables users to build modern applications either on-premises or in the cloud. Microsoft has added Always Encrypted, which encrypts data in use and at rest and enhanced SQL Server auditing capabilities. Those capabilities along with many others are the reasons why so many organizations chose to deploy Microsoft SQL Server.

Solution Configuration

- Hardware resources
- Network configuration
- Architecture diagram
- vSAN storage policy configuration
- Storage policies and Epic EHR workloads
- Software resources

Hardware Resources

While most solutions focus on the software first, Epic predefines the software in this solution. Thus, the critical variable is the hardware for vSAN and so we call out the hardware resources first due to the importance of the performance criteria for Epic on vSAN.

As part of the research for this guide, we tested with the below hardware. Note that this hardware simply serves as a baseline. Refer to the Epic Hardware Guide and the vSAN HCG for the most recent hardware recommendations. Each vSAN node has the following configuration:

TABLE 1. Hardware Configuration for ODB and Analytical Database vSAN Cluster Cluster—Small Customer

PROPERTY	SPECIFICATION	PURPOSE
Server model name	6 x Dell EMC PowerEdge 640	
CPU	2 x Intel(R) Xeon(R) Platinum 6254 CPU @ 3.10GHz, 18 core each	
RAM	576GB	
Network adapter	1 x Mellanox Technologies MT27710 (dual port) 2 x Intel X550 (dual port)	vSAN Network Traffic VM Traffic vMotion and Management Traffic
Storage adapter	1 x Dell HBA330 Adapter	
Disks	Cache - 2 x 1600GB Write Intensive NVMe Capacity - 8 x 3.84TB Mixed Use SAS SSDs	

vSAN Best Practice:

Although VMware vSAN supports a wide range of hardware for optimal performance with applications such as EHRs we recommend:

- Minimum Dual Intel Gold processors with minimum of 18 cores per socket and 2.6GHz frequency base¹
- Minimum of 576GB RAM per vSAN node
- Minimum (4) 10GbE, preferred 25GbE Network Interface Cards
- Two disk groups minimum per vSAN node, with a minimum of (5) disks per disk group
- Mixed Use or Write Intensive NVMe SSDs for the vSAN cache tier disks
- Mixed Use NVMe or 12Gb SAS SSDs for the vSAN capacity tier disks
- Network Switches must be a non-blocking architecture and with high-buffers
- Ensure all components are on the VMware vSAN Hardware Compatibility Guide

These recommendations should be cross-referenced with the Epic Hardware Configuration Guide. Hardware configurations that do not meet the VMware vSAN requirements for Epic EHR may impact performance.

Network Configuration

As shown in Figure 1, we created a VMware vSphere Distributed Switch™ for each vSAN cluster to act as a single virtual switch across all associated nodes for the cluster. The one dual 25GbE NIC on each of the nodes were configured for vSAN traffic. The NICs were trunked together as a Link Aggregation Group (LAG) for bandwidth and redundancy.

The two dual 10GbE NICs were configured on the VDS for VM, vMotion and Management traffic. The NICs were trunked as a Link Aggregation Group (LAG) for bandwidth and redundancy; one LAG using two 10GbE ports for VM traffic and another two 10GbE ports LAG for vMotion and Management traffic. All the networks are configured on different VLAN IDs.

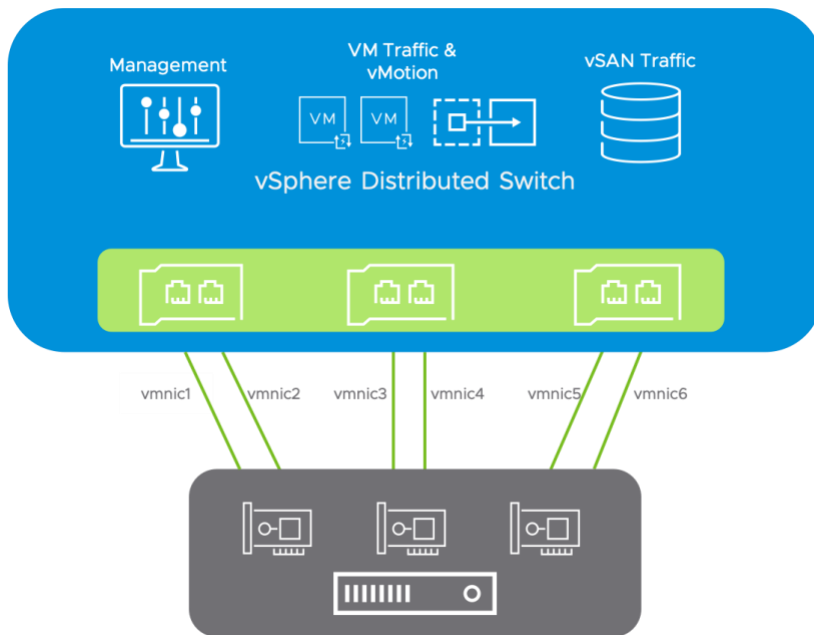


FIGURE 1: VDS Configuration

vSAN Best Practice:

- Configure a LAG for vSAN and a separate LAG for VM, vMotion, and Management traffic
- Configure MTU 9000 for vSAN traffic

¹ Follow Epic’s Hardware config and processor guidelines

Architecture Diagram

In this solution, we use a 6-node vSAN cluster with vSAN Ready Nodes to validate the Epic database workloads. We consulted with their lead architects and followed Epic best practices. Epic testing and recommendations methodology are quite strict and with the onset of Hyperconverged Infrastructure they asked VMware to test additional workloads and scenarios.

TABLE 2. VM Type and Purpose

DATABASE TYPE	DATABASE PURPOSE	VM COUNT
ODB	Production	1
ODB	Non-Production	1
Report	Production	1
Support Release	Production	1
Clarity	Production	1
Clarity	Non-Production	1
Caboodle	Production	1
Caboodle	Non-Production	1
Cube	Production	1
Cube	Non-Production	1
MST ACE	Production	1
vCenter	Production	1

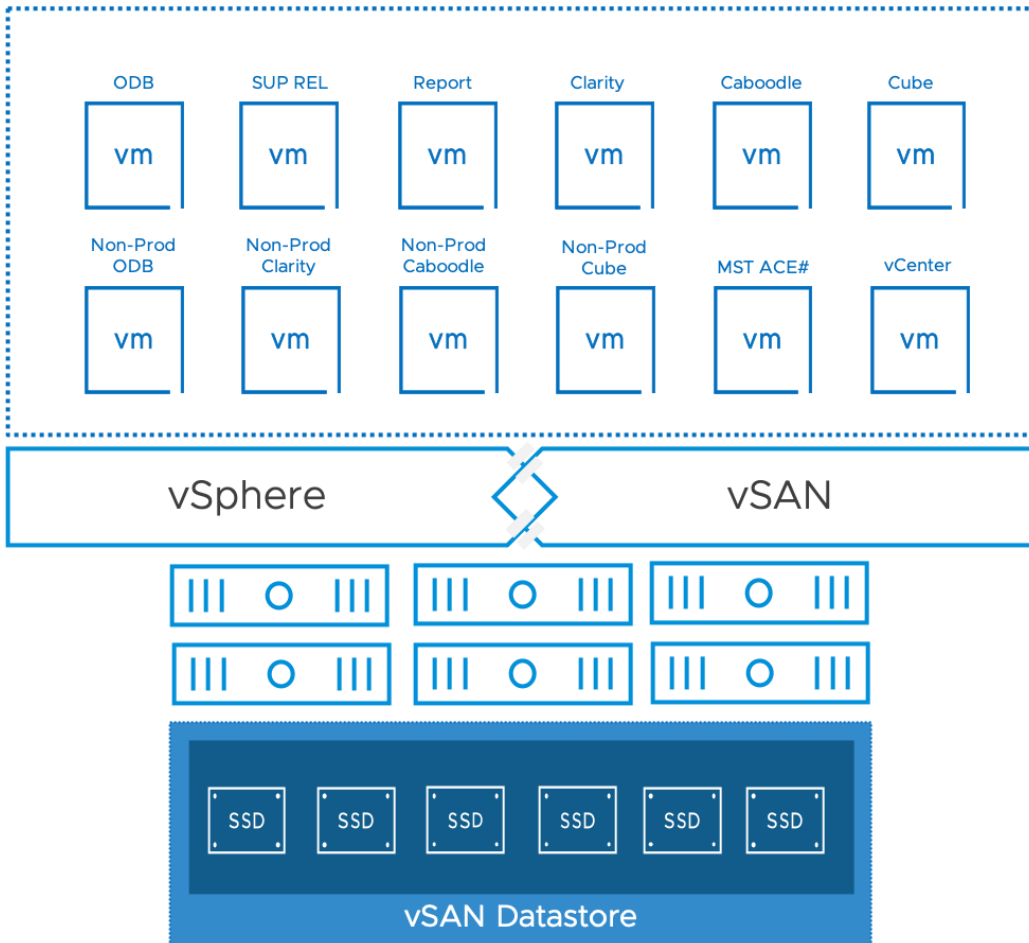


FIGURE 2: Epic ODB and Analytical Database vSAN Cluster

With our Database vSAN cluster, both performance and availability are the driving factors; thus, we use (6) nodes that allow us to use Failures to Tolerate of 2 (FTT=2). The architecture illustrated above will deliver the required performance for Epic ODB. We found that while using RAID1 for the ODB's VMDK's and running an aggressive IO test that vSAN maintains the required IO response time. This configuration is designed for a small Epic customer. This cluster size can scale to what is required in Epic's Hardware configuration guide.

vSAN Best Practice:

- Designate a vSAN cluster for the Database workloads
- Leverage SPBM (Storage Policy Based Management) for performance and availability
- Use RAID1 and FTT=2 for all ODB VMDKs
- Use RAID1 and IOPs Limiter set to 3000 IOPS for all other VMs on the Database vSAN Cluster
- Enable Checksum
- Disable Dedupe and Compression

vSAN Storage Policy Configuration

In our design, we use different storage policies for the Epic workloads. Table 3 shows the detailed configuration.

TABLE 3: vSAN Storage Policy Configuration for Epic Workloads

Feature	Value	Applied to VMs	Description
Failure to Tolerate	2 failure – RAID-1 (Mirroring)	ODB	Defines the number of disks, host or fault domain failures a storage object can tolerate. This set for both PROD and NON-PROD ODB VMs.
Failure to Tolerate	1 failure – RAID-1 (Mirroring)	vCenter, MST ACE#	Defines the number of disks, host or fault domain failures a storage object can tolerate. This is set for all other VMs in Database cluster.
IOPS Limiter	3,000	SUP REL, Report, Clarity, Caboodle, Cube, Non-Prod ODB, Non-Prod Caboodle, Non-Prod Cube, Non-Prod Clarity	Sets the maximum number of IOPS per VMDK. This is set on all VM's VMDKS in the Database Cluster except for PROD and NON-PROD ODB.
Dedupe and Compression	Disabled		Block-level deduplication and compression for storage efficiency.
Checksum	Enabled		Checksum enabled in each vSAN cluster.

SPBM ODB, RAID1 with FTT=2 Example:

Rules VM Compliance VM Template Storage Compatibility

General

Name: ODB-FTT2

Description:

Rule-set 1: VSAN

Placement

Storage Type: VSAN

Site disaster tolerance: None - standard cluster

Failures to tolerate: 2 failures - RAID-1 (Mirroring)

Number of disk stripes per object: 1

IOPS limit for object: 0

Object space reservation: Thin provisioning

Flash read cache reservation: 0%

Disable object checksum: No

Force provisioning: No

SPBM IOPS Limiter set at 3000 IOPS per VMDK with RAID1 Example:

Rules	VM Compliance	VM Template	Storage Compatibility
General			
Name		IOPS Limiter	
Description			
Rule-set 1: VSAN			
Placement			
Storage Type		VSAN	
Site disaster tolerance		None - standard cluster	
Failures to tolerate		1 failure - RAID-1 (Mirroring)	
Number of disk stripes per object		1	
IOPS limit for object		3000	
Object space reservation		Thin provisioning	
Flash read cache reservation		0%	
Disable object checksum		No	
Force provisioning		No	

SPBM RAID1 with FTT=1 Example:

Rules	VM Compliance	VM Template	Storage Compatibility
General			
Name		RAID1-FTT1	
Description			
Rule-set 1: VSAN			
Placement			
Storage Type		VSAN	
Site disaster tolerance		None - standard cluster	
Failures to tolerate		1 failure - RAID-1 (Mirroring)	
Number of disk stripes per object		1	
IOPS limit for object		0	
Object space reservation		Thin provisioning	
Flash read cache reservation		0%	
Disable object checksum		No	
Force provisioning		No	

Storage Policies and Epic EHR Workloads

During our test cycles we concluded, with Epic’s guidance, the use of the IOPS Limiter Storage Policy would be beneficial to prevent any performance impact to the ODB environment in the event of an unplanned IO change. As noted in Table 3 we recommend the IOPS Limiter with 3,000 IOPS per VMDK with FTT=1 to be the vSAN Default Storage Policy. We also recommend using FTT=2 for both PROD ODB and NON-PROD ODB. This gives the ODB environment the ability to handle two failures within the vSAN cluster prior to the ODB failing over to the DR site. Lastly, we recommend using RAID1 for all VMs and VMDKs in the Database cluster. RAID1 will deliver the performance that the Operational and Analytical databases require.

With mission critical workloads such as an EHR, we recommend using the default setting of Checksum. vSAN uses end-to-end checksum to ensure the integrity of data by confirming that each copy of a file is exactly the same as the source file. The system checks the validity of the data during read/write operations, and if an error is detected, vSAN repairs the data or reports the error. If a checksum mismatch is detected, vSAN automatically repairs the data by overwriting the incorrect data with the correct data. Also, we recommend disabling Dedupe and Compression as it may impact application response time.

Software Resources

TABLE 4: Software Resources

SOFTWARE	VERSION	PURPOSE
VMware vSphere and vSAN	6.7 U3 - 14320388	vSphere Cluster to host virtual machines and provide vSAN Cluster. vSAN is a software-defined storage solution for hyperconverged infrastructure.
VMware vCenter Server	6.7 U3 - 14367737	Centralized platform for managing VMware vSphere environments
Centos	7.6.1810	Operating system for Operational database
InterSystems Caché	2018.1.2	Operational database platform
InterSystems IRIS	2019.1	Operational database platform
Microsoft Windows Server	2016, x64, Standard Edition	Operating system for Analytical database
Microsoft SQL Server 2016	2016 Enterprise Edition	Analytical database platform
GenIO		Epic Test Tool
DiskSpd	2.0.17	Microsoft test tool for SQL

vSAN Best Practice:

- Ensure the ESXi and vCenter builds versions are on Epic’s Target Platform
- Ensure the Linux and Windows versions are on Epic’s Target Platform and supported by VMware
- Ensure the Oracle and/or Microsoft SQL Servers versions are on Epic’s Target Platform
- Contact epic@vmware.com prior to conducting any testing or procuring hardware to ensure the success of your project

Check out the following references:

- [VMware Performance Best Practices](#)
- [VMware Best Practices for SQL Server](#)

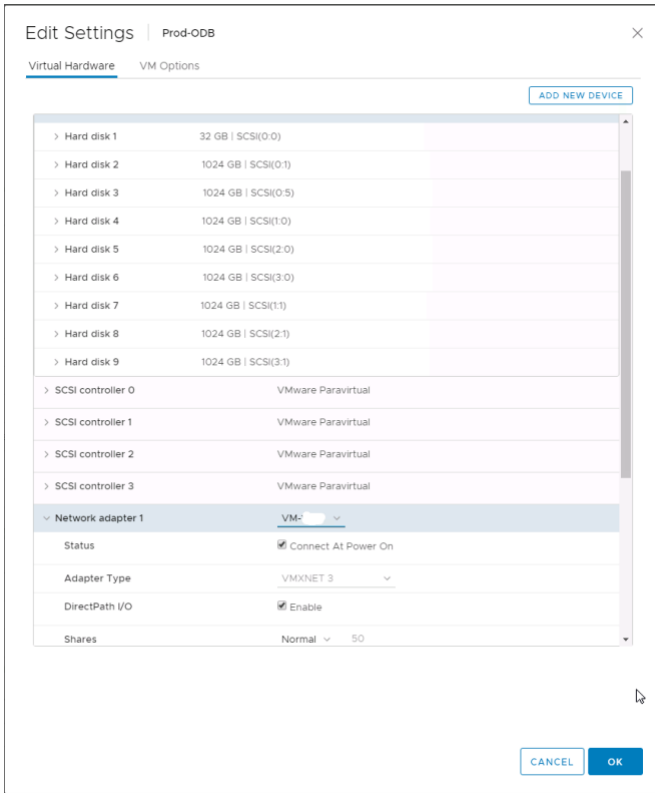
Operational Database Caché Linux VM Best Practices and Layout

One of the key differences between vSAN and traditional Fibre Channel SAN storage is disk type presented to the OS. With SAN storage a RAW Disk mapping is presented, however, with vSAN a VMDK is presented. This greatly reduces the complexity of both administration and troubleshooting. We follow the Epic Storage Quick Reference Guide for the storage layout starting from the VM configuration which remains the same with vSAN.

- Use Multiple PVSCSI Controllers
- Use VMXNET3 NIC
- Distribute the disks across the PVSCSI controllers
 - PVSCSI 0: OS/database VMDKs
 - PVSCSI 1: Database, /epic/prd, /epic VMDKs
 - PVSCSI 2: Database, /epic/prd, /epic VMDKs
 - PVSCSI 3: Database, Journal VMDKs
- Configure the IO Scheduler
- Create Volume Groups
- Create Logical Volumes
- Create File Systems
- Mount File Systems

The VM CPU, RAM, and storage layout will be documented in the Hardware Configuration Guide.

Below is our Prod-ODB VM layout as an example.



SQL Cogito VM Best Practices and Layout

As mentioned previously, the disk type presented to the OS is a key difference with vSAN. This greatly reduces the complexity of administration, troubleshooting, and configuration. We follow the Epic Cogito on VMware Architecture document.

- Use Multiple PVSCSI Controllers
- Use VMXNET3 NIC
- Distribute the disks across the PVSCSI controllers
- Configure Windows Disks, except OS, for 64K Block Size

Solution Validation

There are three typical profiles of Epic customers: small, medium, and large.

- Small customer can generate up to 5M global references and <25K IOPs.
- Medium customer generates between 5M to 10M global references and between 25K – 50K IOPs.
- Large customer generates more than 10M global references and >50K IOPs.

Epic requires consistent and predictable response times for their applications. Notably the ODB has the following requirements:

- Random reads to the ODB, using the file system response time:
 - Average read latencies must be 2ms or less

- 99% of read latencies must be below 60ms
- 99.9% of read latencies must be below 200ms
- 99.99% of read latencies must be below 600ms
- Random writes to the ODB, using the file system response time:
 - Average write latencies must be 1ms or less
- Average write cycle must be completed <45 seconds

Using the GenIO test tool vSAN performs at expectable levels for both small and medium size Epic customer environments. While vSAN also performs well for large size customer profiles, Epic has restricted HCI to small and medium customers for initial support.

Conclusion

VMware vSAN proves to be on-par with traditional storage for mission-critical workloads such as Epic EHR. However, performance is only one factor for such a demanding and dynamic ecosystem; vSAN also delivers both cost-effectiveness and agility. vSAN powered HCI organically changes the EHR ecosystem and exposes the business application owners to architectural efficiency and simplicity by empowering them with deployment choices without sacrificing performance.

By adopting vSAN infrastructure and application, owners can more readily collaborate and architect for the business problems of today and tomorrow.

About the Author and Contributors

Christian Rauber, a staff solutions architect in the Hyperconverged Infrastructure, Product Enablement team wrote the original version of this paper.

