



Best Practices for VMware Cloud Foundation with Epic

VMware Application Acceleration

Table of contents

Best Practices for VMware Cloud Foundation with Epic	3
Executive Summary	3
Introduction	4
Technology Overview	5
VMware Cloud Foundation	5
VMware vSphere	5
VMware vSAN	5
VMware NSX Data Center	5
InterSystems Caché and IRIS	5
Microsoft SQL Server 2016	6
Test Tools	7
Monitoring Tools	7
vSAN Performance Service	7
VMware Skyline Health	7
Workload Generation and Testing Tools	7
Solution Configuration	8
Hardware Resources	8
Network Configuration	9
Architecture Diagram	10
vSAN Storage Policy Configuration	11
Storage Policies and Epic EHR Workloads	13
Software Resources	14
.....	15
Operational Database Caché Linux VM Best Practices and Layout	16
SQL Cogito VM Best Practices and Layout	17
Solution Validation	18
Conclusion	19
About the Author and Contributors	20

Best Practices for VMware Cloud Foundation with Epic

Executive Summary

With the list of priorities growing daily, hospital IT departments struggle to just maintain critical systems. IT projects like migrating workloads to the cloud are often postponed or delayed indefinitely due to lack of resources or cloud expertise. Unfortunately, it is these very workloads that would benefit from cloud like features such as workload isolation and dynamic network segmentation.

Continuing to delay these projects comes at cost, either a hardware failure or mandatory upgrade reshuffles the project to an escalation. This escalation cycle vs. project repeats itself. How can hospitals manage these projects before they become escalated priorities and is there a way to migrate to the cloud without the additional recourses and expertise?

Introduction

VMware Cloud Foundation™ is an integrated software platform that automates the deployment and lifecycle management of a complete software-defined data center (SDDC) on a standardized hyperconverged architecture. It can be deployed on premises on a broad range of supported hardware or consumed as a service in the public cloud (VMware Cloud™ on AWS or a VMware Cloud Provider™). With the integrated cloud management capabilities, the result is a hybrid cloud platform that can span private and public environments, offering a consistent operational model based on well-known VMware vSphere® tools and processes, and the freedom to run apps anywhere without the complexity of app re-writing.

This document outlines the best practices for Epic Operational and Analytic Databases on VMware Cloud Foundation 4.2.

Technology Overview

The technology components in this solution are:

- VMware Cloud Foundation
 - VMware vSphere
 - VMware vSAN
 - NSX Data Center
- InterSystems Caché and IRIS
- Microsoft SQL Server 2016

VMware Cloud Foundation

VMware Cloud Foundation is an integrated software stack that combines compute virtualization (VMware vSphere®), storage virtualization (VMware vSAN™), network virtualization (VMware NSX®), and cloud management and monitoring (VMware vRealize® Suite) into a single platform that can be deployed on-premises as a private cloud or run as a service within a public cloud. This documentation focuses on the private cloud use case. VMware Cloud Foundation bridges the traditional administrative silos in data centers, merging compute, storage, network provisioning, and cloud management to facilitate end-to-end support for application deployment.

VMware vSphere

VMware vSphere is VMware's virtualization platform, which transforms data centers into aggregated computing infrastructures that include CPU, storage, and networking resources. vSphere manages these infrastructures as a unified operating environment and provides operators with the tools to administer the data centers that participate in that environment. The two core components of vSphere are ESXi and vCenter Server. ESXi is the virtualization platform where you create and run virtual machines and virtual appliances. vCenter Server is the service through which is used to manage multiple hosts connected in a network and pool host resources.

VMware vSAN

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, and 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 Hyperconverged Infrastructure (HCI) encryption.

vSAN 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.

VMware NSX Data Center

VMware NSX Data Center is the network virtualization and security platform that enables the virtual cloud network, a software-defined approach to networking that extends across data centers, clouds, and application frameworks. With NSX Data Center, networking and security are brought closer to the application wherever it's running, from virtual machines to containers to bare metal. Like the operational model of VMs, networks can be provisioned and managed independent of underlying hardware. NSX Data Center reproduces the entire network model in software, enabling any network topology—from simple to complex multitier networks—to be created and provisioned in seconds. Users can create multiple virtual networks with diverse requirements, leveraging a combination of the services offered via NSX or from a broad ecosystem of third-party integrations ranging from next-generation firewalls to performance management solutions to build inherently more agile and secure environments. These services can then be extended to a variety of endpoints within and across clouds.

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 be accessing and manipulating 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 testing, we use Microsoft SQL Server 2016 that 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.

Test Tools

We leverage the following monitoring and benchmark tools for our functional testing for the operational and analytic databases on VMware Cloud Foundation.

Monitoring Tools

vSAN Performance Service

[vSAN Performance Service](#) is used to monitor the performance of the vSAN environment through the vSphere Client. The performance service collects and analyzes performance statistics and displays the data in a graphical format. You can use the performance charts to manage your workload and determine the root cause of problems.

VMware Skyline Health

[VMware Skyline Health Diagnostics](#) is a self-service diagnostics platform. It helps to detect and provide remediation to issues in the vSphere and vSAN product line. It provides recommendations in the form of a Knowledge Base article or procedure to remediate issues. It uses the log files of these products to detect issues. vSphere administrators can use this tool to troubleshoot issues before contacting VMware Global Support Service.

Workload Generation and Testing Tools

GenerateIO is a test tool provided and designed by Epic to simulate the workload I/O of the Operational Database.

[DiskSpd](#) is a Microsoft Windows based storage I/O tool designed to simulate Windows I/O. Epic provides a preconfigured DiskSpd test tool package designed to simulate workload I/O of the Analytic Database.

With both test tools, the results are gathered at the end of each test run and sent to Epic for consultation.

Solution Configuration

This section introduces the resources and configurations:

- 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 VMware Cloud Foundation and so we call out the hardware resources first due to the importance of the performance criteria for Epic on VMware Cloud Foundation and vSAN.

As part of the research for this guide, we test 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 VMware Cloud Foundation 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	

VMware Cloud Foundation and vSAN Best Practice:

Although VMware vSAN supports a wide range of hardware for optimal performance with applications such as EHRs, we recommend the following storage best practices:

- Minimum Dual Intel Gold processors with minimum of 18 cores per socket and 2.6GHz frequency base^[1]
- 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

Figure 1 shows the VMware vSphere Distributed Switches configuration for the workload domain of the VMware Cloud. Two 25 GbE vmnics were used and configured with teaming policies.

In the management workload domain, NSX-T resides on the management port group of the virtual dedicated server. In the Epic VI workload domain, we use NSX-T to add segments to the VDS. For the NSX-T design guidance, see [NSX-T Reference Design](#).

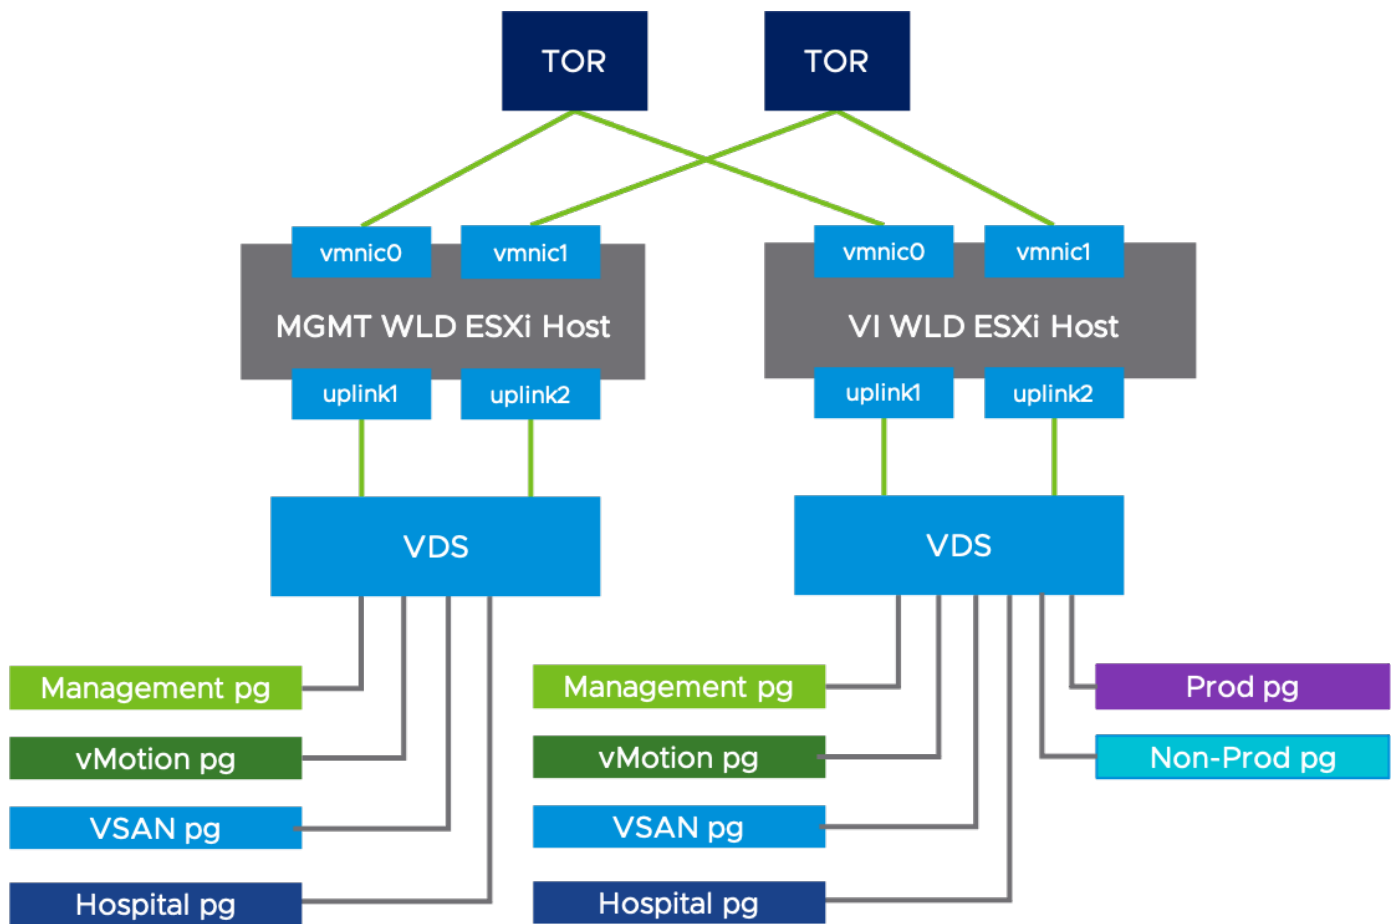


Figure 1. VMware vSphere Distributed Switches Configuration

Beyond the segments created by VMware Cloud Foundation, we created two additional segments to service production VM traffic isolated from non-production VM traffic.

Table 2. Networking Configuration

Segment	Teaming Policy	VMNIC0	VMNIC1
Management network	Route based on Physical NIC load	Active	Active
VM network	Route based on Physical NIC load	Active	Active
vSphere vMotion	Route based on Physical NIC load	Active	Active
vSAN	Route based on Physical NIC load	Active	Active
Hospital	Load Balance Source	Active	Active

VMware Cloud Foundation vSAN Best Practice:

- Configure Large MTU for vSAN traffic

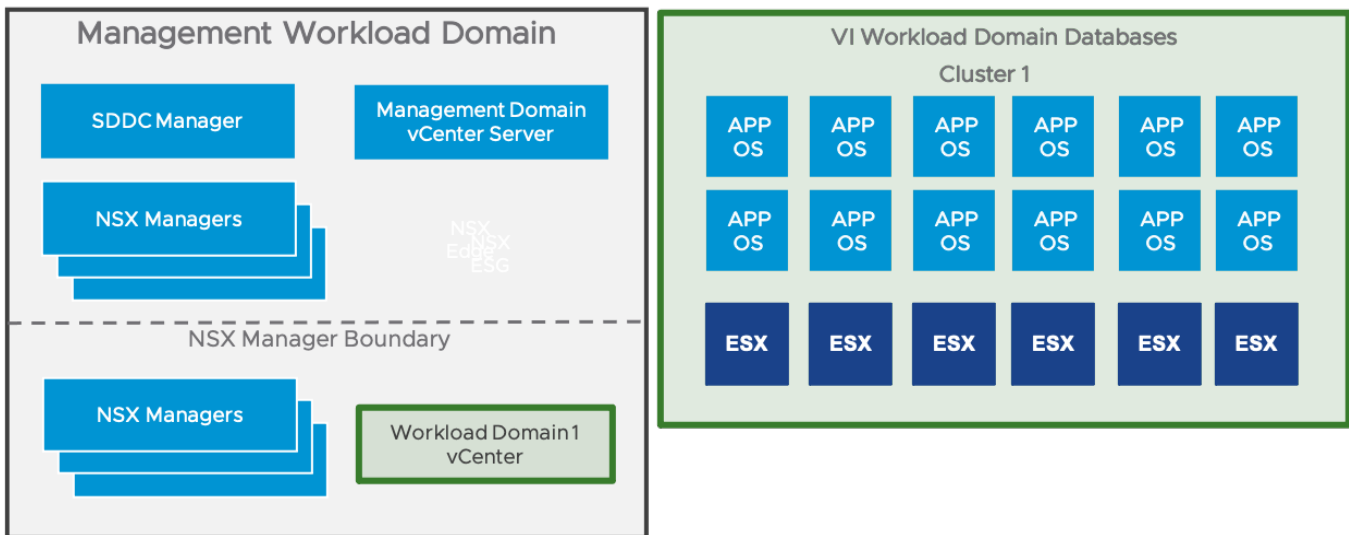
Architecture Diagram

In our test environment, we deployed a 6-node Virtual Infrastructure (VI) Workload Domain to validate the Epic database workloads and a 4-node Management Workload Domain. We consulted with Epic lead architects and followed their 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 3. VM Settings

VM Role	Purpose	VM Count
Operational Database	Production	1
Operational Database	Non-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
VMware vCenter	Production	1

VM Role	Purpose	VM Count
Management Domain vCenter Server	Mgmt. Production	1
Platform Server Controller	Mgmt. Production	2
SDDC Manager	Mgmt. Production	1
Management Domain NSX-T Manager	Mgmt. Production	3
VI Workload Domain NSX-T Manager	Mgmt. Production	3
vRealize Log Insight	Mgmt. Production	3
vRealize Lifecycle Manager	Mgmt. Production	1
vRealize Operations	Mgmt. Production	1



With our Database vSAN cluster, both performance and availability are the driving factors; thus, we used (6) nodes that allow us to set Failures to Tolerate to 2 (FTT=2). The architecture illustrated above delivers the required performance for Epic ODB. We found that while using RAID1 for the ODB’s VDMK’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.

Best Practice for VMware Cloud Foundation and vSAN:

- 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 3,000 IOPS for all other VMs on the Database vSAN Cluster
- Enable Checksum
- Disable Dedupe and Compression

vSAN Storage Policy Configuration

In our design, we used different storage policies for the Epic workloads. Table 4 shows the detailed configuration.

Table 4. vSAN Storage Policy Configuration

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.
Deduplication 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	
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		Thick provisioning	
Flash read cache reservation		0%	
Disable object checksum		No	
Force provisioning		No	
Encryption services		No encryption	
Space efficiency		No space efficiency	
Storage tier		All flash	

SPBM RAID1 with FTT=1 Example:

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
Encryption services	No preference
Space efficiency	No space efficiency
Storage tier	All flash

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		
Encryption services	No preference		
Space efficiency	No space efficiency		
Storage tier	All flash		

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 4, 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 deduplication and compression as it may impact application response time.

Software Resources

Table 5 shows the software resources used in this solution.

Table 5. Software Resources

Software	Version	Purpose	
VMware Cloud Foundation	4.2	An integrated software stack that combines compute virtualization, storage virtualization, network virtualization, and cloud management and monitoring into a single platform.	
VMware vCenter Server and ESXi	7.0 u1c/u1d	VMware vSphere is a suite of products: vCenter Server and ESXi.	
VMware vSAN	7.0 u1d	vSAN is the storage component to provide low-cost and high-performance next-generation HCI solutions.	
VMware NSX-T Data Center	3.1.0	NSX-T provides an agile software-defined infrastructure to build cloud-native application environments.	
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	
GeneratelO	1.15.2	Epic Test Tool	
DiskSpd	2.0.21a	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.

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 GenerateIO 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 Cloud Foundation enables hospital IT departments to leverage their existing subject matter expertise and migrate seamlessly into a hybrid cloud. Deploying VMware Cloud Foundation removes one more barrier for faster new-application adoption with built-in network load balancing and policy based I/O isolation. Most importantly the ever-challenging life cycle management is automated within VMware Cloud Foundation.

About the Author and Contributors

Christian Rauber, staff solutions architect in the Solutions Architecture team of the Cloud Platform Business Unit, wrote the original version of this paper.

[1] Follow Epic's Hardware config and processor guidelines:
https://www.appliedsystems.com/doc_central/ep-epic/epichswconfigurationguidelines2016.htm

