


MULTIPLE ACTIVE DATACENTERS SOLUTION



Victor Wu

Senior Solution Expert, Business Consultation

BoardWare Information System Limited

Victor.wu@boardware.com.mo



The Dell Technologies Proven Professional Certification program validates a wide range of skills and competencies across multiple technologies and products.

From Associate, entry-level courses to Expert-level, experience-based exams, all professionals in or looking to begin a career in IT benefit from industry-leading training and certification paths from one of the world's most trusted technology partners.

Proven Professional certifications include:

- Cloud
- Converged/Hyperconverged Infrastructure
- Data Protection
- Data Science
- Networking
- Security
- Servers
- Storage
- Enterprise Architect

Courses are offered to meet different learning styles and schedules, including self-paced On Demand, remote-based Virtual Instructor-Led and in-person Classrooms.

Whether you are an experienced IT professional or just getting started, Dell Technologies Proven Professional certifications are designed to clearly signal proficiency to colleagues and employers.

Learn more at www.dell.com/certification

Table of Contents

Preface	5
Introduction	6
Environment.....	7
Hardware Requirements.....	11
Software Requirements	12
Network Requirements	13
SAN Fabric Requirements.....	18
Active-Active-Passive Solution.....	19
Active-Active-Active Solution.....	21
Failure Scenarios	23
Scenario One	24
Scenario Two	25
Scenario Three.....	26
Scenario Four.....	27
Scenario Five	28
Scenario Six	29
Scale-Out.....	30
Summary	32
Bibliography	33

Table of Figures

Figure 1 - The sample environment of multiple active Datacenters	7
Figure 2 - The rack layout diagram of multiple active Datacenters.....	8
Figure 3 - The network diagram of vSAN node.....	14
Figure 4 - The network diagram of the witness host.....	15
Figure 5 - The SAN fabric diagram for AAP solution	18
Figure 6 - The Active-Active-Passive solution diagram.....	20
Figure 7 - The Active-Active-Active solution diagram.....	21
Figure 8 - The architecture of vCenter Server Appliance HA.....	22
Figure 9 - Scenario One in AAP solution	24
Figure 10 - Scenario Two in AAP solution	25
Figure 11 - Scenario Three in AAP solution.....	26
Figure 12 - Scenario Four in AAA solution	27
Figure 13 - Scenario Five in AAA solution	28
Figure 14 - Scenario Six in AAA solution	29
Figure 15 - Configuration of vSAN cluster in vCenter Server.....	31
Figure 16 - vSphere Storage Policy-Based Management.....	31

Disclaimer: The views, processes or methodologies published in this article are those of the author. They do not necessarily reflect Dell Technologies' views, processes or methodologies.

Preface

In the digital economy, the service of mission-critical applications cannot be disrupted because it may impact the business service. System architects ensure that the performance, availability and service level agreements (SLAs) are fulfilled. When considering the high-availability solution and its performance design, many factors are considered. Most of the elements are:

- What is the best high-availability solution, e.g. storage level or application level?
- Are we using the traditional server architecture or hyper-converged infrastructure (HCI) architecture?
- Which hypervisor platform are we using, e.g. Microsoft Hyper-V or VMware vSphere?
- What are the service level agreements of business/application service?
- What is the Recovery Time Objective (RTO)?
- What is the Recovery Point Objective (RPO)?
- How many datacenters are allocated for the business service?
- What is the performance requirement?
- What are the business requirements and constraints?
- How to scale-up and scale-out the existing infrastructure?
- What is the service impact during maintain windows?

When you need continuous application availability and increased efficiency for your modern datacenter, you need to consider the different product solutions based on the business requirements. Dell EMC VPLEX is one of the solutions for you. Providing continuous availability and active-active infrastructure across two datacenters, it can support on the traditional server and Dell EMC storage. Ensure your mission-critical applications are always on and the service is non-disruptive if the service of one of two datacenters is down. However, if there are multiple datacenters across three locations, Dell EMC VPLEX does not fulfil the multiple active (Active-Active-Active) datacenter requirements across three separate areas.

Introduction

This Knowledge Sharing article provides information on multiple active Datacenters solution. We will consider how to plan and design the hyper-converged infrastructure (HCI) as multiple active Datacenters solution across three separate locations; will discuss a sample solution for multiple active Datacenters based on different SLA requirements and application tiers conditions. This solution includes software-defined storage (SDS) and Active-Active-Passive (AAP) infrastructure. In the design phase of this solution, takes into consideration the following:

1. The datacenter allocation requirement: Design how to allocate the servers and network equipment into each datacenter for a multiple active datacenter solution.
2. The software-defined storage network design: Design SDS network architecture across three separate datacenters.
3. The software-defined storage design: Plan and design hyper-converged infrastructure architecture.
4. Active-Active-Passive storage architecture design. Plan and design the traditional storage architecture for AAP solution across three separate datacenters.
5. Multiple active vSAN architecture design. Plan and design VMware vSAN architecture for a multiple active Datacenter solution.
6. SLA requirement: Plan and design the data protection based on the various SLAs and application tier across three separate datacenters.
7. The failure scenarios for multiple active datacenters: List the business/application service status based on different failure scenarios across three datacenters.
8. Scale-up and scale-out management: Plan and design the business service move across the datacenters. Scale-up within the local datacenter and scale-out extra datacenters expansion.

Figure 1 is a high-level diagram for multiple active datacenters; we will discuss the details of this environment in the next section.

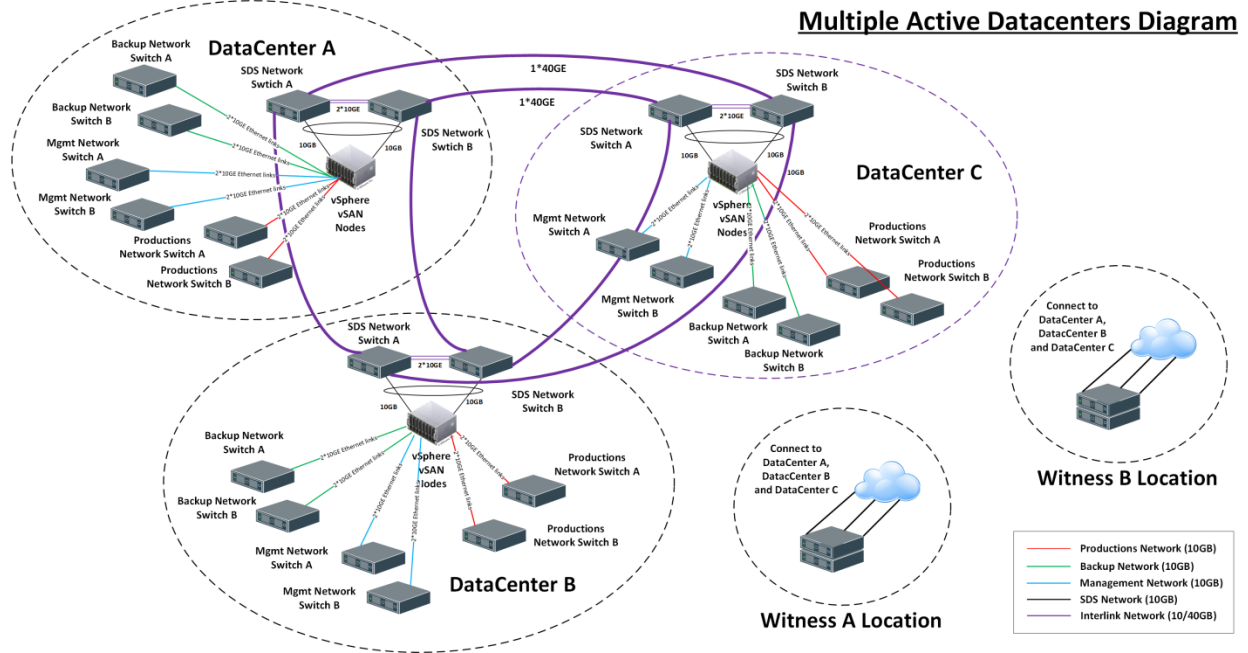


Figure 1 - The sample environment of multiple active Datacenters

Environment

In this section, we will discuss the details of hardware equipment for multiple active data centers; it includes the followings in Table 1. Figure 2 is the physical rack layout diagram for this solution. This environment includes Active-Active-Passive (AAP) and Active-Active-Active (AAA) features.

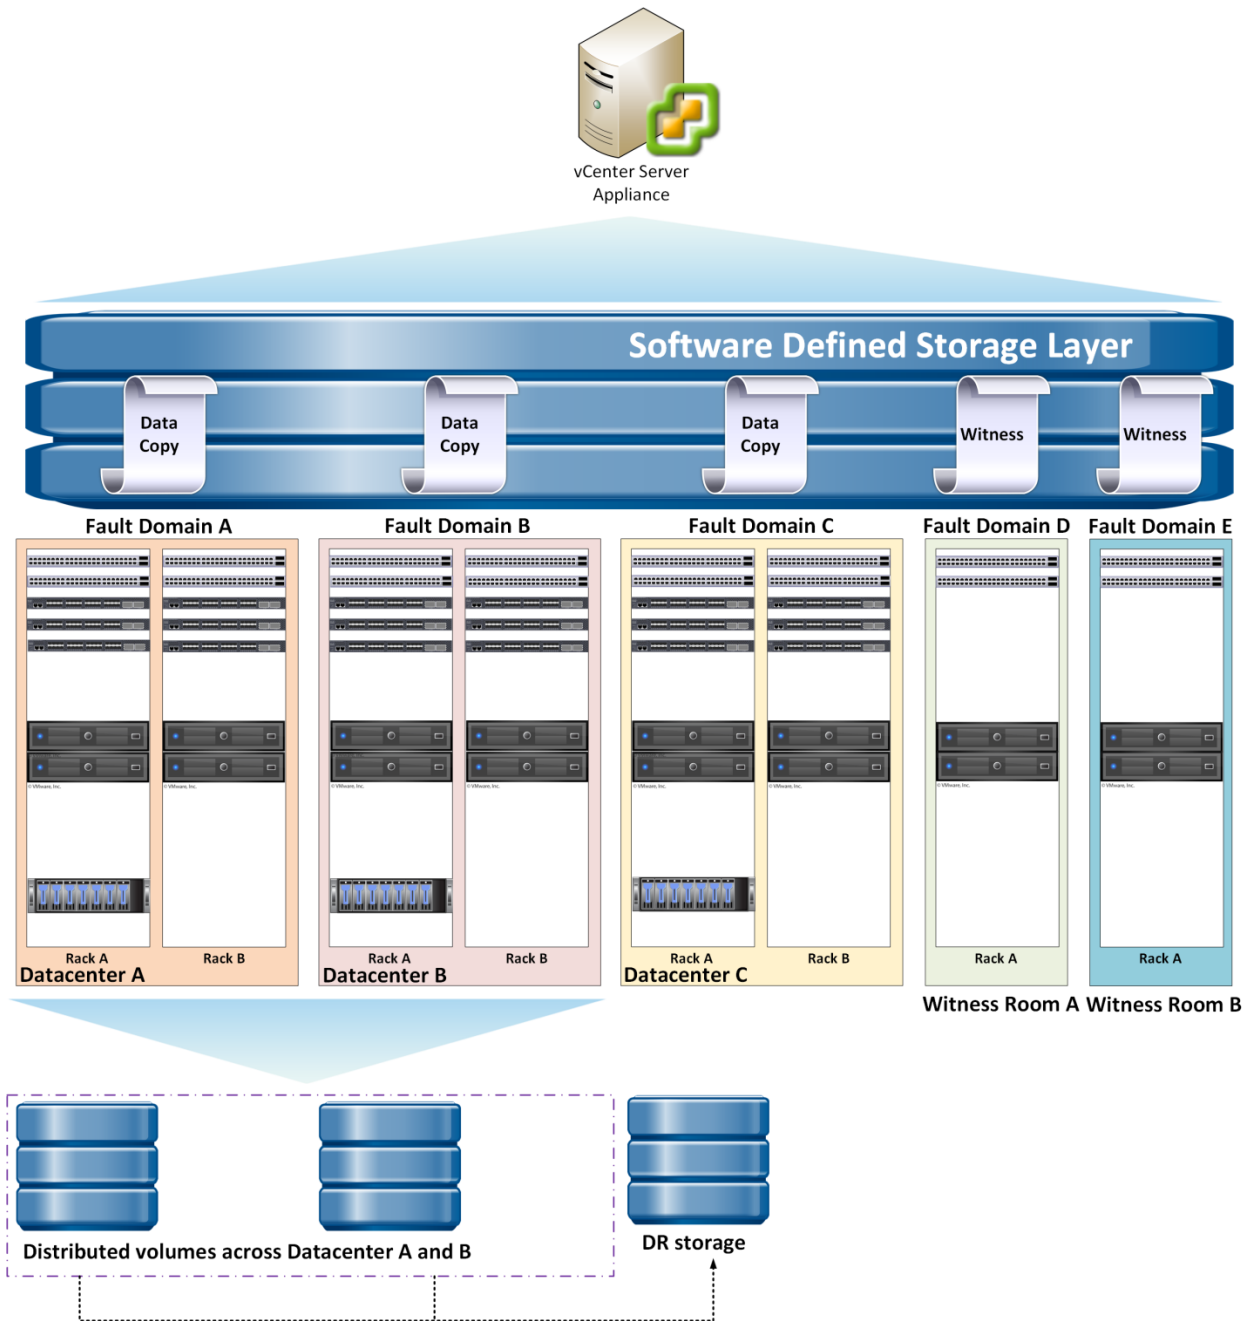


Figure 2 - Rack layout diagram of multiple active Datacenters

The Active-Active-Passive solution is the active-active storage platform across Datacenter A and B. The business service is still running across Datacenter A and B if one or both Datacenter A and B is faulted. VMware Site Recovery Manager (SRM) with storage replication or vSphere Replication enables the business service to perform the failover operation to Datacenter C when both Datacenter A and B are faulted. We will discuss the details in the section on Active-Active-Passive solution.

The Active-Active-Active solution is the active-active-active storage platform across the Datacenter A, B and C. The business service is still running across Datacenters if one of three Datacenters is faulted. VMware vSAN performs the architecture of this Active-Active-Active solution; we will discuss the details in the section of Active-Active-Active solution.

Now we will discuss the details of all hardware installed on each datacenter.

Location	1Gb Network Switch	10Gb Network Switch	Servers	SAN Storage and SAN Switch
Datacenter A	2 x network switch for Management Network.	2 x network switch for Production Network. 2 x network switch for Backup Network. 2 x network switch for SDS Network.	4 x servers	1 x SAN storage and 2 x SAN Switch
Datacenter B	2 x network switch for Management Network.	2 x network switch for Production Network. 2 x network switch for Backup Network. 2 x network switch for SDS Network.	4 x servers	1 x SAN storage and 2 x SAN Switch
Datacenter C	2 x network switch for Management Network.	2 x network switch for Production Network. 2 x network switch for Backup Network. 2 x network switch for SDS Network.	4 x servers	1 x SAN storage and 2 x SAN Switch
Witness Room A	2 x network switch for Management Network.	2 x network switch for SDS Network.	2 x servers	N/A
Witness Room B	2 x network switch for Management Network.	2 x network switch for SDS Network.	2 x servers	N/A

Table 1 - The hardware equipment for multiple active Datacenters

In Datacenter A, there are eight Ethernet switches for different network zones connection, e.g. production, backup, management and SDS. Four servers were running in VMware vSphere hypervisor and enabled the vSAN feature. For active-active-passive solution, the metro storage cluster must be enabled across Datacenter A and B; then we replicate the production data into DR SAN storage at Datacenter C. For active-active-active solution, we configure a vSphere vSAN cluster across five separate locations.

In Datacenter B, there are eight Ethernet switches for different network zones connection, e.g. production, backup, management and SDS. Four servers were running in VMware vSphere hypervisor and enabled the vSAN feature. For active-active-passive solution, the metro storage cluster must be enabled across Datacenter A and B; then we replicate the production data into DR SAN storage at Datacenter C. For active-active-active solution, we configure a vSphere vSAN cluster across five separate locations.

In Datacenter C, there are eight Ethernet switches for different network zones connection, e.g. production, backup, management and SDS. Four servers were running in VMware vSphere hypervisor and enabled the vSAN feature. For active-active-passive (AAP) solution, we need to install SAN storage (Dell SC5020 storage) at Datacenter C and configure as DR storage for data recovery if both SAN storage of Datacenter A and B are faulted.

In Witness Room A, there are four Ethernet switches for different network connections, e.g. management and SDS. Two servers were running in VMware vSphere hypervisor and enabled the vSAN feature.

In Witness Room B, there are four Ethernet switches for different network connections, e.g. management and SDS. Two servers were running in VMware vSphere hypervisor and enabled the vSAN feature.

Hardware Requirements

This section lists the details of server equipment for multiple active Datacenters; it includes:

Location	vSAN Node	Servers	CPU/CPU cores	RAM	Cache disks/Cache disk size	Capacity disks/Capacity disk size	Number of network ports
Datacenter A	All-Flash	4	2 x CPU / 24 cores per CPU	256GB	2 x disk / 200GB	6 x disk / 1.92TB	8 x 10GB ports
Datacenter B	All-Flash	4	2 x CPU / 24 cores per CPU	256GB	2 x disk / 200GB	6 x disk / 1.92TB	8 x 10GB ports
Datacenter C	All-Flash	4	2 x CPU / 24 cores per CPU	256GB	2 x disk / 200GB	6 x disk / 1.92TB	8 x 10GB ports
Witness Room A	All-Flash	2	1 x CPU / 24 cores	128GB	1 x disk / 100GB	3 x disk / 100GB	6 x 10GB ports
Witness Room B	All-Flash	2	1 x CPU / 24 cores	128GB	1 x disk / 100GB	3 x disk / 100GB	6 x 10GB ports

Table 2 - Server requirements for multiple active datacenters

Software Requirements

This section lists the details of software equipment for multiple active Datacenters; it includes:

Description	Detail	Remark
VMware vSphere 6.7 or above	Hypervisor	
vCenter Server 6.7 or above	Virtual Appliance base	
Platform Services Controller (PSC)	Virtual Appliance base	The part of vCenter Server Appliance
VMware Upgrade Manager (VUM)	It enables centralized, automated patch and version management for VMware vSphere. It also supports virtual machine and virtual appliances.	The part of vCenter Server Appliance
VMware Site Recovery Manager 8.x or above (SRM)	Virtual Appliance base	
VMware Replication Appliance 8.x or above (VR)	Virtual Appliance base	
VMware vSAN Service	The software-defined storage feature which is embedded on vSphere.	

Table 3 – The software requirements for multiple active datacenters

VMware vSphere Hypervisor is a bare-metal hypervisor; it enables customers to virtualize hardware servers and consolidate hardware resources (CPU and memory) into the application.

VMware vCenter Server is advanced server management software that provides a centralized platform for managing your VMware vSphere environments, e.g. creation of virtual machines, management of Software-Defined Network (SDN) and Software-Defined Storage (SDS), etc.

VMware Update Manager is bundled service of vCenter Server Appliance; it enables centralized, automated patch and version management for VMware vSphere, virtual machines, and virtual appliances.

VMware Site Recovery Manager (SRM) is a virtual appliance that provides the disaster recovery solution for your VMware vSphere environments; it enables the SRM recovery plan with both storage base replication and virtual machine replication base.

VMware vSphere Replication is an extension to VMware vCenter Server that provides hypervisor-based virtual machine replication and recovery in local site/across the site.

VMware vSAN is a software-defined storage solution that supports hyper-converged infrastructure (HCI) systems. vSAN is fully integrated with VMware vSphere as a distributed layer of software within the ESXi hypervisor.

Network Requirements

This section discusses the network requirement of each server installed at each Datacenter. Details of each server are shown in Table 3.

Location	Compute Node	Network Ports	Fibre Channel Ports	Out-of-band Management
Datacenter A	Server A	10 x 10Gb SFP+ ports	2 x 16GB FC ports	1 x 1Gb port
	Server B	10 x 10Gb SFP+ ports	2 x 16GB FC ports	1 x 1Gb port
	Server C	10 x 10Gb SFP+ ports	2 x 16GB FC ports	1 x 1Gb port
	Server D	10 x 10Gb SFP+ ports	2 x 16GB FC ports	1 x 1Gb port
Datacenter B	Server A	10 x 10Gb SFP+ ports	2 x 16GB FC ports	1 x 1Gb port
	Server B	10 x 10Gb SFP+ ports	2 x 16GB FC ports	1 x 1Gb port
	Server C	10 x 10Gb SFP+ ports	2 x 16GB FC ports	1 x 1Gb port
	Server D	10 x 10Gb SFP+ ports	2 x 16GB FC ports	1 x 1Gb port
Datacenter C	Server A	10 x 10Gb SFP+ ports	2 x 16GB FC ports	1 x 1Gb port
	Server B	10 x 10Gb SFP+ ports	2 x 16GB FC ports	1 x 1Gb port
	Server C	10 x 10Gb SFP+ ports	2 x 16GB FC ports	1 x 1Gb port
	Server D	10 x 10Gb SFP+ ports	2 x 16GB FC ports	1 x 1Gb port
Witness Room A	Server A	6 x 10Gb SFP+ ports	N/A	1 x 1Gb port
	Server B	6 x 10Gb SFP+ ports	N/A	1 x 1Gb port
Witness Room B	Server A	6 x 10Gb SFP+ ports	N/A	1 x 1Gb port
	Server B	6 x 10Gb SFP+ ports	N/A	1 x 1Gb port

Table 4 – Network requirements on each server for multiple active datacenters

Figure 3 is the network connectivity of vSAN node for different network zones at Datacenter A, B and C.

Remark: “vmnic” stands for VMware Network Interface Card. “HBA” stands for Host Bus Adapter.

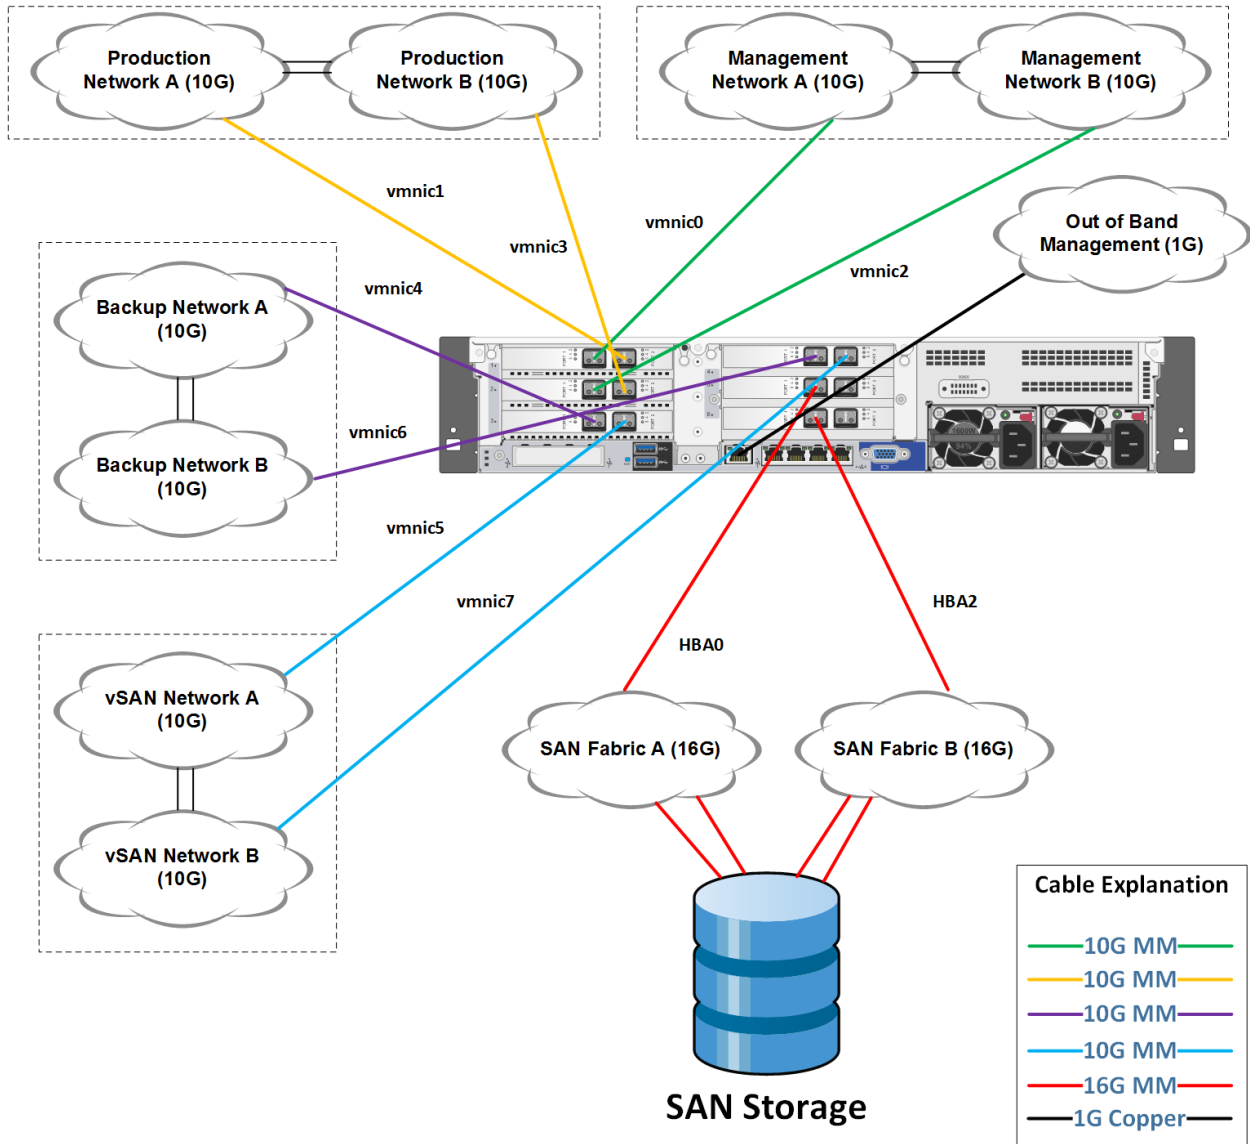


Figure 3 - Network diagram of vSAN node

Figure 4 is the network connectivity of the Witness node for different network zones at Witness room A and B.

Remark: “vmnic” stands for VMware Network Interface Card. “HBA” stands for Host Bus Adapter.

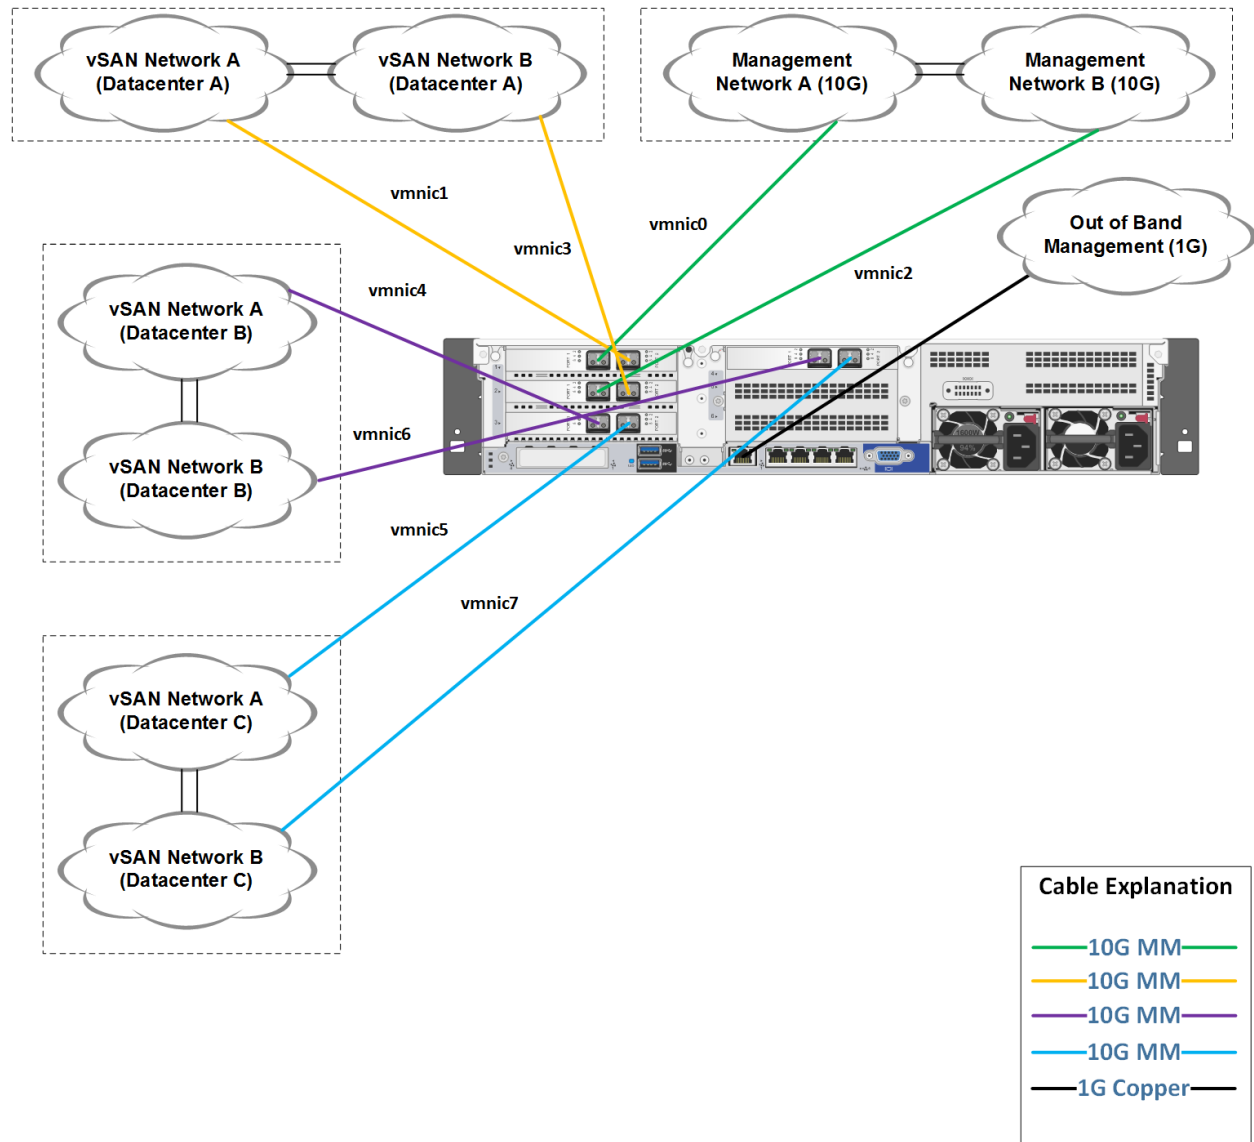


Figure 4 - Network diagram of the Witness host

When we create virtual networking on each VMware vSphere host, we need to define the different virtual network switch and its uplinks. VMware vSphere can support either vSphere Standard Switch (VSS) or Distributed Virtual Switch (DVS). We highly suggest defining DVS, which can provide central network management and flexible control on your VMware vSphere environments.

We need to create the isolated network for vSAN across each Datacenter, and dual 10Gb Ethernet links for redundancy on each SDS Network Switch. There is a pair of dedicated 10Gb network switches for vSAN network that installs at isolated physical Racks. Datacenter-to-Datacenter connectivity is dual 10Gb links for redundancy with <1ms network latency. Table 4 shows the network requirements of vSAN network across three Datacenters and two Witness rooms.

Location	vSAN Network	Network connections
Datacenter A	2 x 10GB Ethernet uplinks	Datacenter A to Datacenter B
	2 x 10GB Ethernet uplinks	Datacenter A to Datacenter C
Datacenter B	2 x 10GB Ethernet uplinks	Datacenter B to Datacenter C
	2 x 10GB Ethernet uplinks	Datacenter B to Datacenter A
Datacenter C	2 x 10GB Ethernet uplinks	Datacenter C to Datacenter A
	2 x 10GB Ethernet uplinks	Datacenter C to Datacenter B
Witness Room A	2 x 10GB Ethernet uplinks	Witness Room A to Datacenter A
	2 x 10GB Ethernet uplinks	Witness Room A to Datacenter B
	2 x 10GB Ethernet uplinks	Witness Room A to Datacenter C
Witness Room B	2 x 10GB Ethernet uplinks	Witness Room B to Datacenter A
	2 x 10GB Ethernet uplinks	Witness Room B to Datacenter B
	2 x 10GB Ethernet uplinks	Witness Room B to Datacenter C

Table 5 – Network requirements of vSAN network across each datacenter

According to the above sample environment, we will define four network zones for each vSAN node, ESXi management network, production network, vSAN network and backup data network.

1. vSphere Standard Switch (vSwitch) A is used for the ESXi management and vMotion network; each network port group has two physical uplinks.
2. Distributed Virtual Switch (DvSwitch) A is used for vSAN network, which has two physical uplinks.
3. DvSwitch B is used for the production network, which has two physical uplinks.
4. DvSwitch C is used for the Backup Data Network, which has two physical uplinks.

Virtual Switch	Port Group	# of vNICs	vNIC Teaming
vSphere Standard Switch A	ESXi Management vMotion Network	2	vmnic0, vmnic2
Distributed Virtual Switch A	vSAN Network	2	vmnic5, vmnic7
Distributed Virtual Switch B	production Network	2	vmnic1, vmnic3
Distributed Virtual Switch C	Backup Data Network	2	vmnic4, vmnic6

Table 6 - Virtual network configuration of vSAN node

For the witness host, we will define two network zones for each vSAN node and ESXi management network.

1. vSphere Standard Switch (vSwitch) A is used for the ESXi management and vMotion network; each network port group has two physical uplinks.
2. Distributed Virtual Switch (DvSwitch) A is used for vSAN network, which has six physical uplinks.

Virtual Switch	Port Group	# of vNICs	vNIC Teaming
vSphere Standard Switch A	ESXi Management vMotion Network	2	vmnic0, vmnic2
Distributed Virtual Switch A	vSAN Network A	2	vmnic1, vmnic3
	vSAN Network B	2	vmnic4, vmnic6
	vSAN Network C	2	vmnic5, vmnic7

Table 7 - Virtual network switch configuration of the witness node

SAN Fabric Requirements

This section discusses the SAN Fabric network requirement for the Active-Active-Passive solution. Figure 5 is the SAN fabric diagram for three Datacenters.

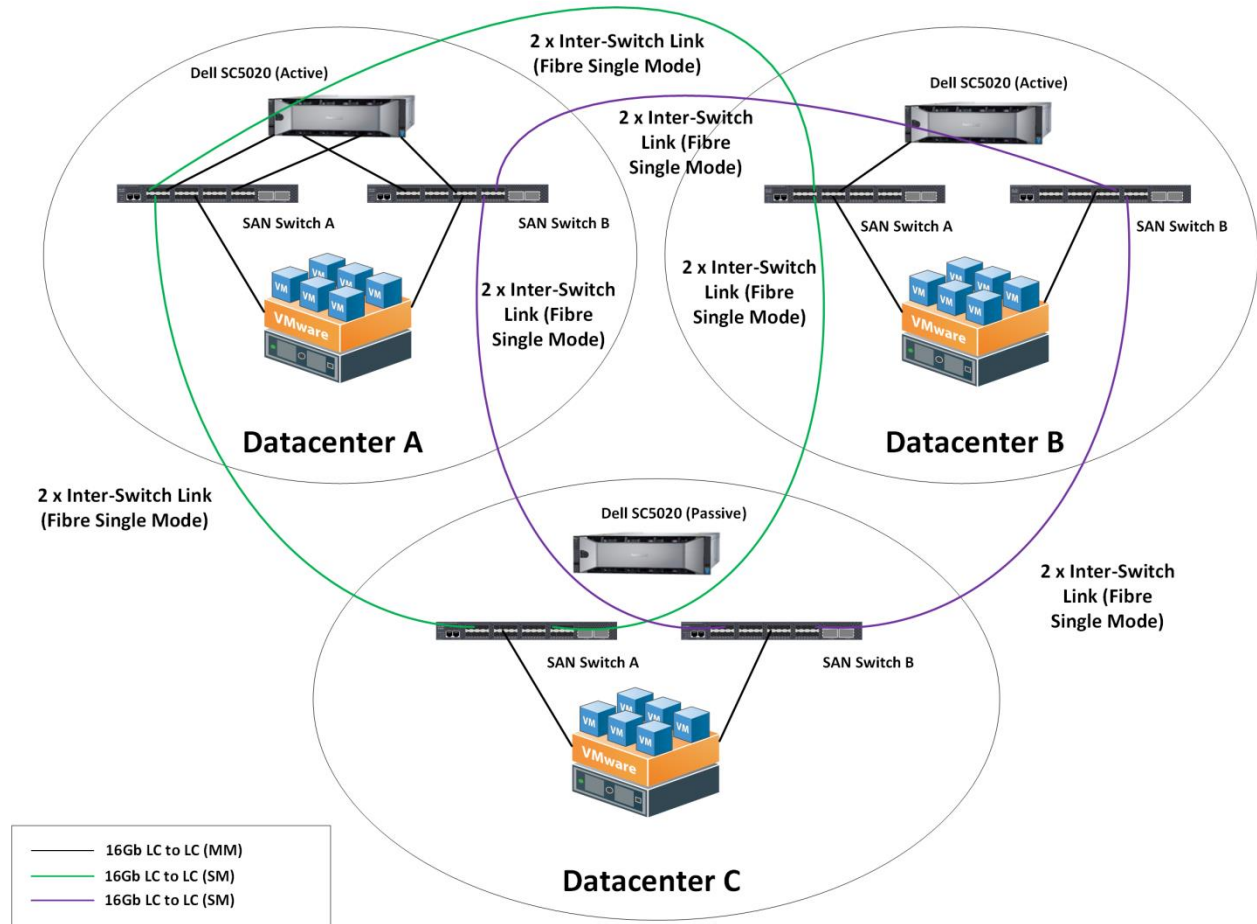


Figure 5 - SAN fabric diagram for AAP solution

The Active-Active-Passive solution requires two SAN switches and a SAN storage installed at each Datacenter. In this environment, the recommended hardware is Connectrix DS-6510B SAN switch and Dell SC5020 SAN storage. There are two inter-switch links (ISL) that connect to each SAN switch A across three Datacenter so that you can see it has two SAN Fabric network.

Location	SAN switch model	Number of ISL	SAN Fabric
Datacenter A	Dell EMC Connectrix DS-6510B	2 x 16G FC single mode connection on each SAN switch	SAN Switch A is Fabric A SAN Switch B is Fabric B
Datacenter B	Dell EMC Connectrix DS-6510B	2 x 16G FC single mode connection on each SAN switch	SAN Switch A is Fabric A SAN Switch B is Fabric B
Datacenter C	Dell EMC Connectrix DS-6510B	2 x 16G FC single mode connection on each SAN switch	SAN Switch A is Fabric A SAN Switch B is Fabric B

Table 8 - SAN Fabric requirement of each SAN switch

Active-Active-Passive Solution

This solution is running across three Datacenters. Figure 6 shows three Dell SC5020 storages and some vSphere hosts running at each Datacenter; the live volumes of Dell SC5020 are enabled between two Dell SC5020 storages across Datacenter A and B and the vSphere Stretched Cluster with live volumes as VMFS datastore across Datacenter A and B. All virtual machines of vSphere Stretched Cluster can keep running on live volumes and no service interruption when Dell SC5020 storage is faulted at either Datacenter A or B. The vSphere Stretched Cluster is managed by one vCenter Server, the other vSphere cluster is managed by the other vCenter Server at Datacenter C. When both SC5020 storage A and B are faulted, how can the virtual machines be recovered? We also configure the asynchronous replication from a consistent group of SC5020 storage A and B to SC5020 storage C.

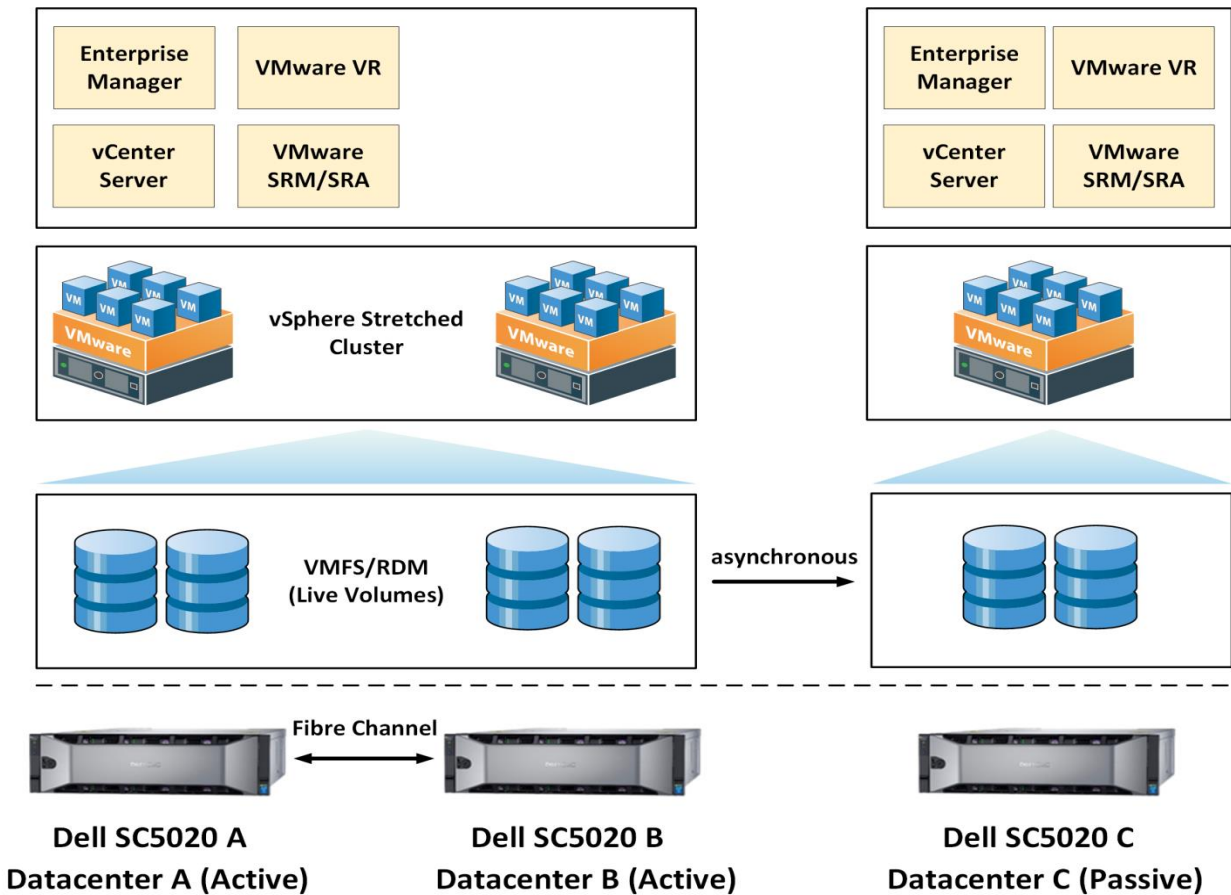


Figure 6 - Active-Active-Passive solution diagram

VMware vSphere Site Recovery Manager (SRM) is a disaster recovery and planned migration tool for virtual machines. SRM can support recovery protection with both array-based replication and virtual machine replication. We need to define the SRM recovery plan for disaster recovery. When both SC5020 storage A and B are faulted, we can execute VMware SRM recovery plan to recover the virtual machines into DR vSphere cluster at Datacenter C. According to the above, the virtual machines are running on an Active-Active-Passive platform across three Datacenters. Table 9 is the summary of Active-Active-Passive solution.

Datacenter	ESXi Cluster	Management host	Disaster Recovery Software	Storage Replication
Datacenter A	vSphere Stretched Cluster	vCenter Appliance Server	Site Recovery Manager/SRA vSphere Replication	Enabled Synchronous Replication
Datacenter B				
Datacenter C	vSphere Cluster	vCenter Appliance Server	Site Recovery Manager/SRA vSphere Replication	Enabled Asynchronous Replication

Table 9 - Summary of Active-Active-Passive solution

Active-Active-Active Solution

Figure 7 shows five separate locations, with the production data allocated into Datacenter A, B and C. Two small rooms are used for allocating the vSAN witness hosts. Create a vSphere cluster with all ESXi hosts in five locations, and then enable the vSAN feature on this vSphere cluster that can deliver a software-defined storage. The virtual machines of vSAN cluster can be running across Datacenter A, B and C.

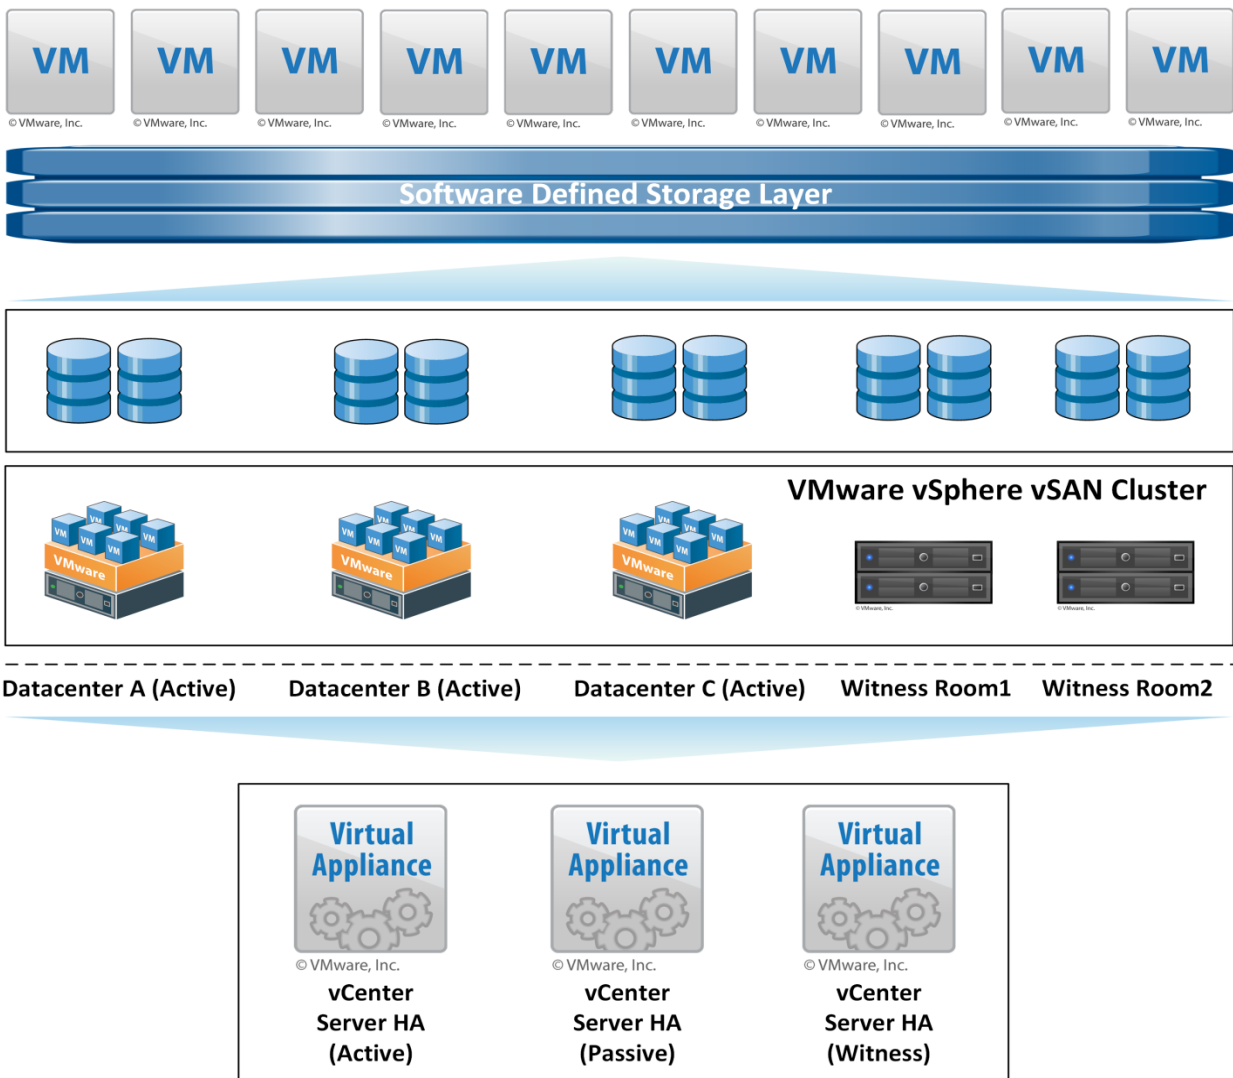


Figure 7 - Active-Active-Active solution diagram

When the vSAN cluster configuration is ready, you need to define five “Fault Domains” for the Active-Active-Active platform (refer to Table 10). In vSphere settings, it includes Storage Policy-Based Management (SPBM), and you can define the different storage policies for each ESXi host in SPBM. In vSAN storage policies, you must choose with FTT = 2 (mirror protection) which is three equal data components and then distribute these across Datacenter A, B and C.

Sites	vSphere Cluster	Fault Domain	Storage Policy	vSAN components
Datacenter A	vSphere vSAN Cluster	A	FTT = 2, FTM=RAID-1	Replica
Datacenter B		B	FTT = 2, FTM=RAID-1	Replica
Datacenter C		C	FTT = 2, FTM=RAID-1	Replica
Witness Room1		D	FTT = 2, FTM=RAID-1	Witness
Witness Room2		E	FTT = 2, FTM=RAID-1	Witness

Table 10 - Summary of vSAN storage policy for Active-Active-Active solution

For the central management of Active-Active-Active platform, there is a vCenter Server Appliance to manage the vSphere vSAN cluster. You can see the vCenter Server Appliance is also enabled with the High available feature. Figure 8 is the architecture of vCenter Server Appliance HA.

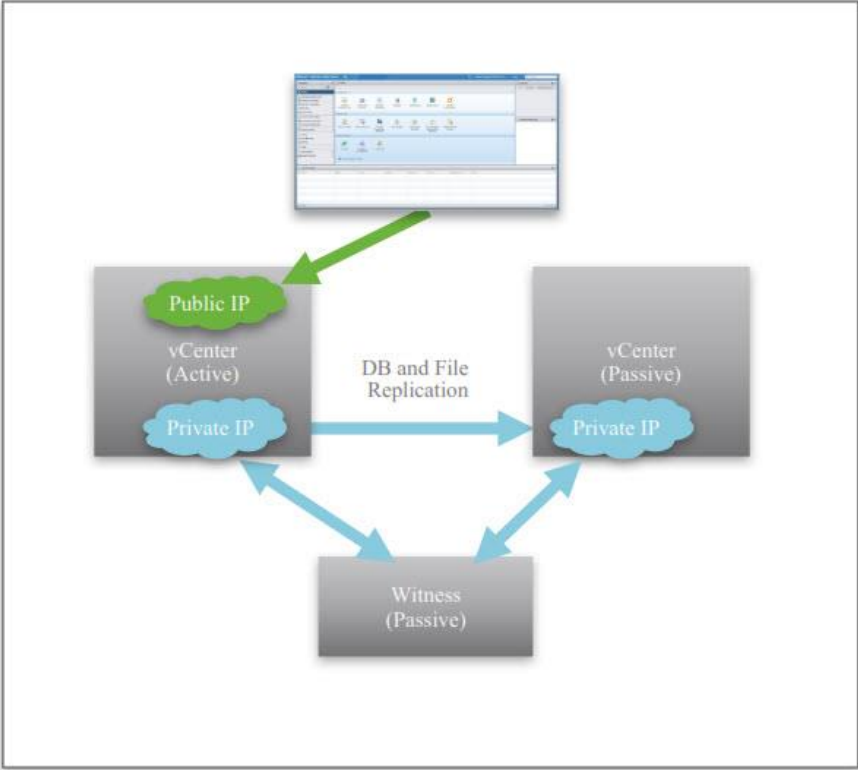


Figure 8 - Architecture of vCenter Server Appliance HA

A vCenter HA cluster consists of three vCenter Server Appliance instances; active node, passive node and witness node. When the vCenter Server HA is enabled, it includes two IP address, one is the public network, and the other is the private network. You can access the public network of the vCenter Server for management. Table 11 shows software and hardware requirements of vCenter Server HA.

Component	Requirements
ESXi	ESXi 6.0 or above. A minimum of three ESXi hosts is highly recommended; each vCenter HA node should be running on a different ESXi host for better protection. Suggest enabling VMware DRS feature to protect the set of ESXi hosts.
vCenter Server Appliance	VMware vCenter 6.5 or above, 4 CPU and 16GB RAM or above.
Network connectivity	vCenter HA network latency between Active, Passive, and Witness nodes must be less than 10 ms. The vCenter HA network must be on a different subnet than the management network.
Licenses	vCenter HA requires a single vCenter Server license. vCenter HA requires a Standard edition.

Table 11 - The software and hardware requirement of vCenter Server HA

If the service of vCenter HA active node is stopped, the vCenter service can failover to the passive node allocated at Datacenter B. The vCenter witness node provides a quorum to protect against a split-brain situation.

Failure Scenarios

This section discusses some failure scenarios for Active-Active-Passive and Active-Active-Active solution. Both Active-Active-Passive and Active-Active-Active solutions can coexist on a vSphere vSAN cluster, which is managed by a vCenter Server Appliance. The customer can flexibly migrate the virtual machines to different types of Datastore in a vSphere cluster. Each type of Datastore is based on the service level agreements (SLAs), e.g. vSAN datastore, distribution volume and local datastore, etc.

Scenario One

What is the status of virtual machines running at Datacenter A when the Dell SC5020 is faulted at Datacenter A?




Figure 9 - Scenario 1 in AAP solution

In Active-Active-Passive solution, the customer can online move the virtual machines across two Datacenters when the maintenance windows require at both Datacenters. The virtual machines are running in Stretched Cluster, and two protection methods protect each virtual machine, i.e. Local DC premises HA protection and over two DC premises HA protection. In Figure 9, if the Dell SC5020 is faulted at Datacenter A, the virtual machines are still running and experience no service interruption.

Scenario Two

Figure 10 depicts when all ESXi hosts and Dell SC5050 are faulted at the same time at Datacenter, what is the status of virtual machines running at Datacenter A?

All virtual machines will trigger the vSphere HA feature and start-up into the ESXi hosts at Datacenter B.

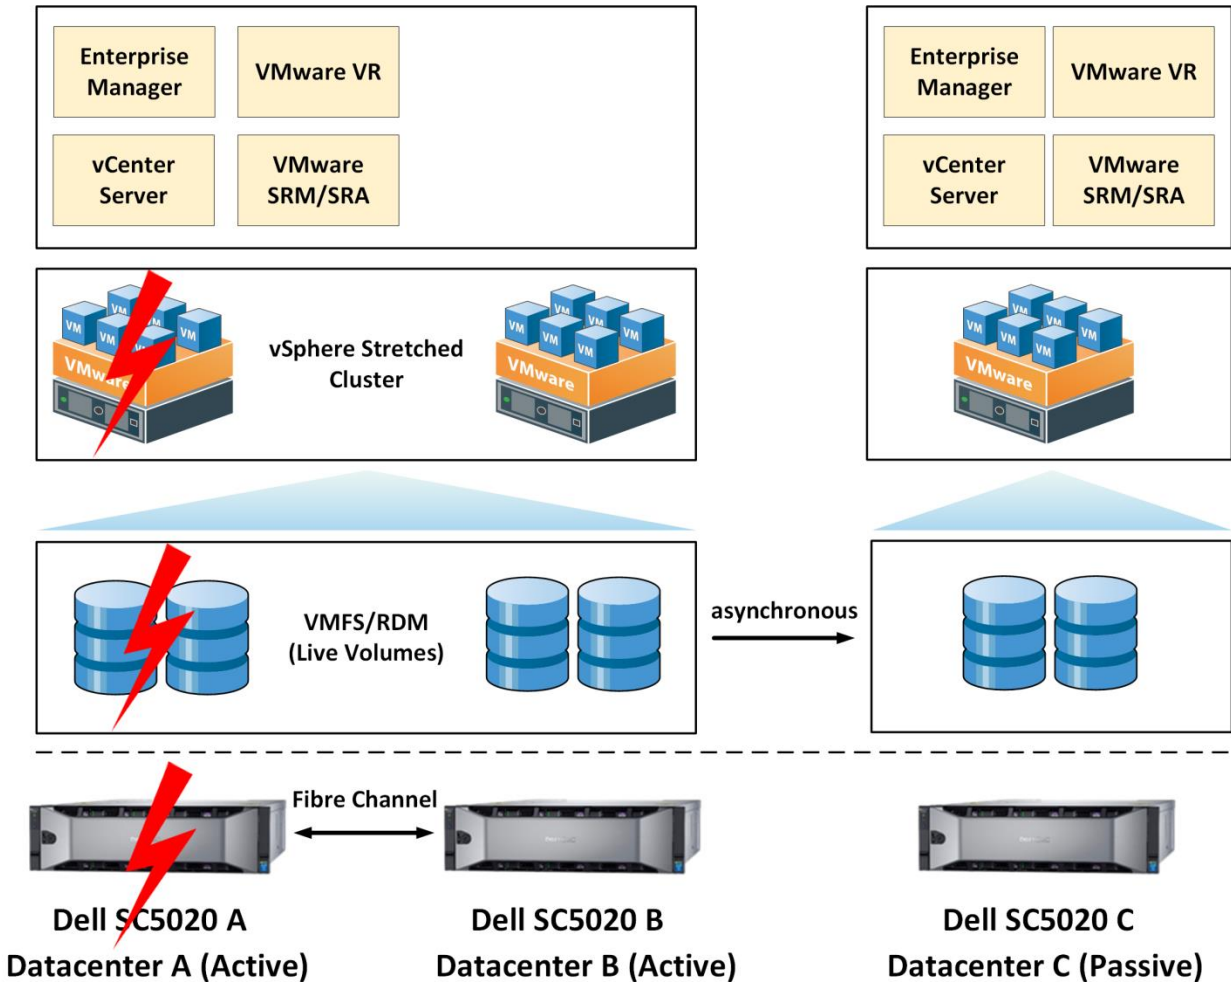


Figure 10 - Scenario 2 in AAP solution

Scenario Three

In Figure 11, when all ESXi hosts and Dell SC5050 are faulted at Datacenter A and B, what is the status of virtual machines running at both Datacenters?

All virtual machines will shut down at once. In this situation, we can execute the VMware SRM recovery plan, and all protected virtual machines can start-up based your RPO settings at Datacenter C.

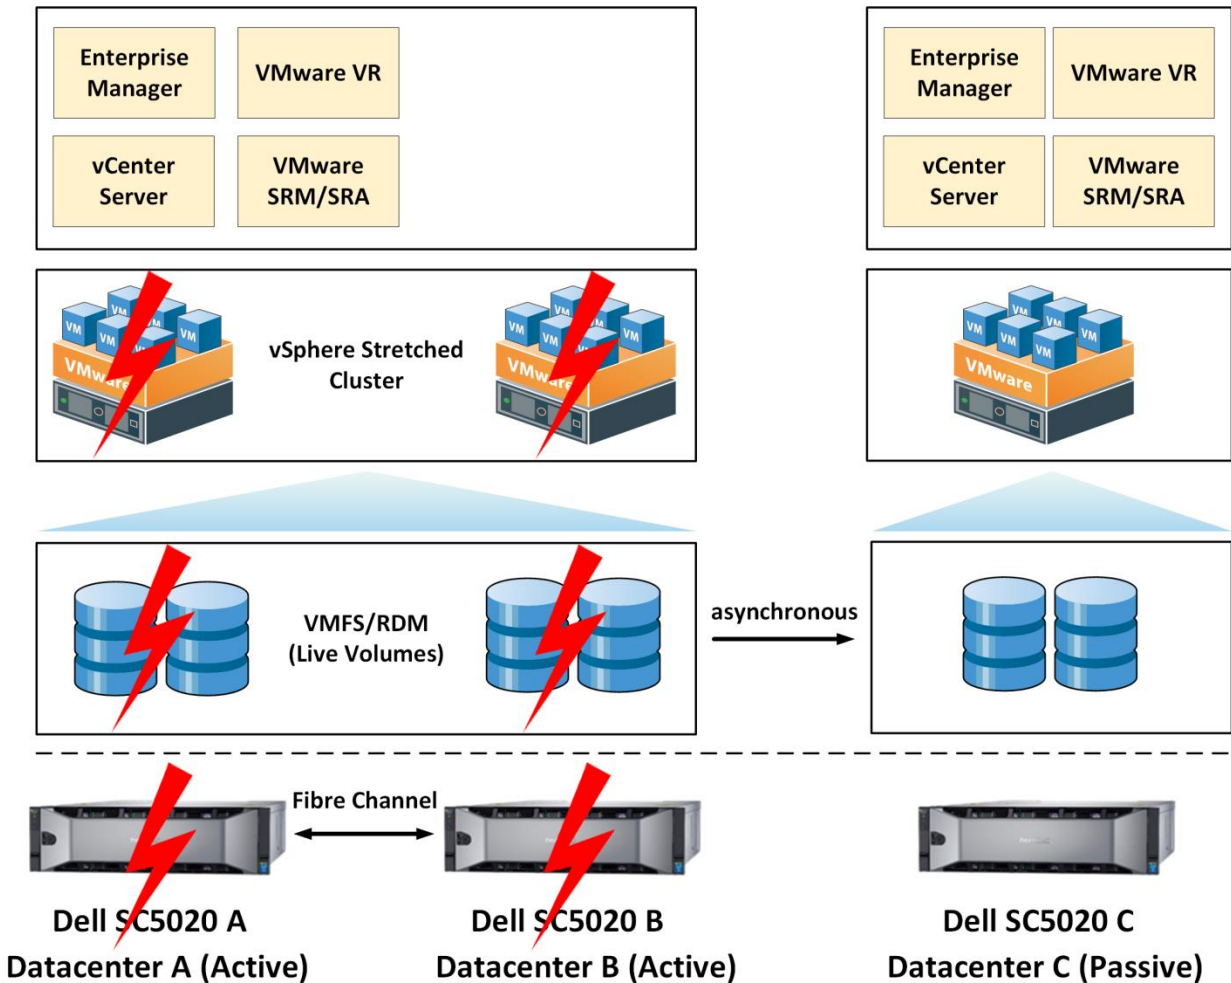


Figure 11 - Scenario 3 in AAP solution

Scenario Four

In Active-Active-Active solution, the customer can online move the virtual machines across each Datacenter when the three Datacenters require maintenance windows. The virtual machines are running in a vSphere vSAN Cluster across three Datacenters, and two protection methods protect each virtual machine, i.e. Local DC premises HA protection and over three DC premises HA protection. Each virtual machine consists own copy across each Datacenter premises.

In Figure 12, if some ESXi hosts or all hosts are faulted at Datacenter A, what is the status of virtual machines running at Datacenter A?

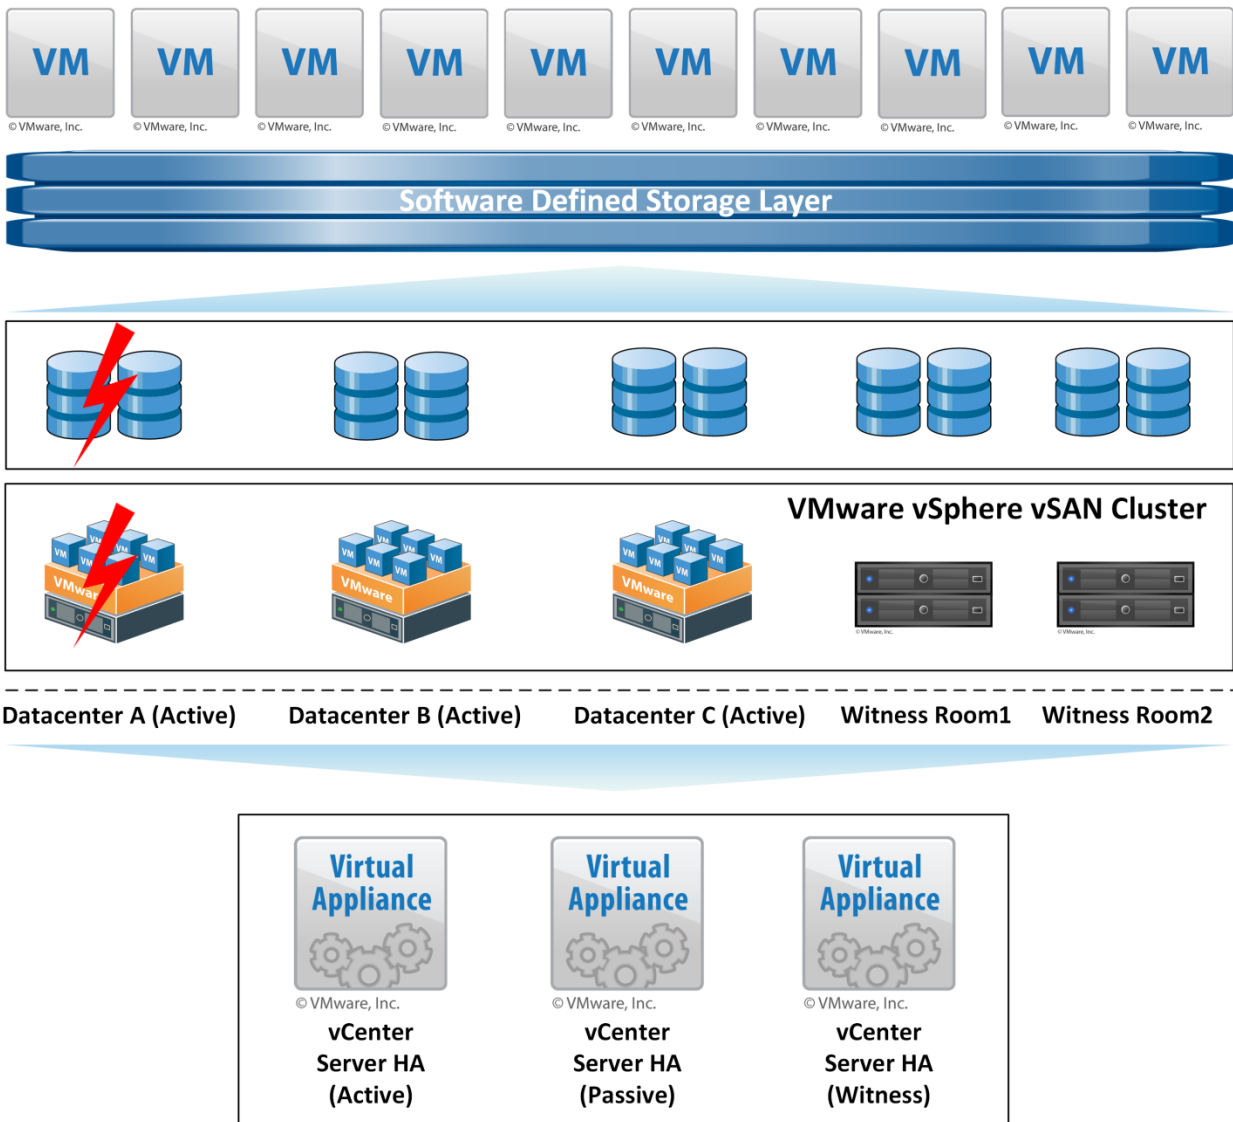


Figure 12 - Scenario 4 in AAA solution

If some ESXi hosts are faulted at Datacenter A, the virtual machines of those ESXi hosts will trigger the vSphere HA feature and then start-up on other ESXi hosts at Datacenter A. These virtual machines still consist own copy across each Datacenter premises.

If all ESXi hosts are faulted at Datacenter A, the virtual machines will trigger the vSphere HA feature and then start-up on the ESXi hosts at Datacenter B. These virtual machines consist own copy across Datacenter B and C.

Scenario Five

In Figure 13, if all hosts are faulted at Datacenter A and B. What is the status of virtual machines running at both Datacenters?

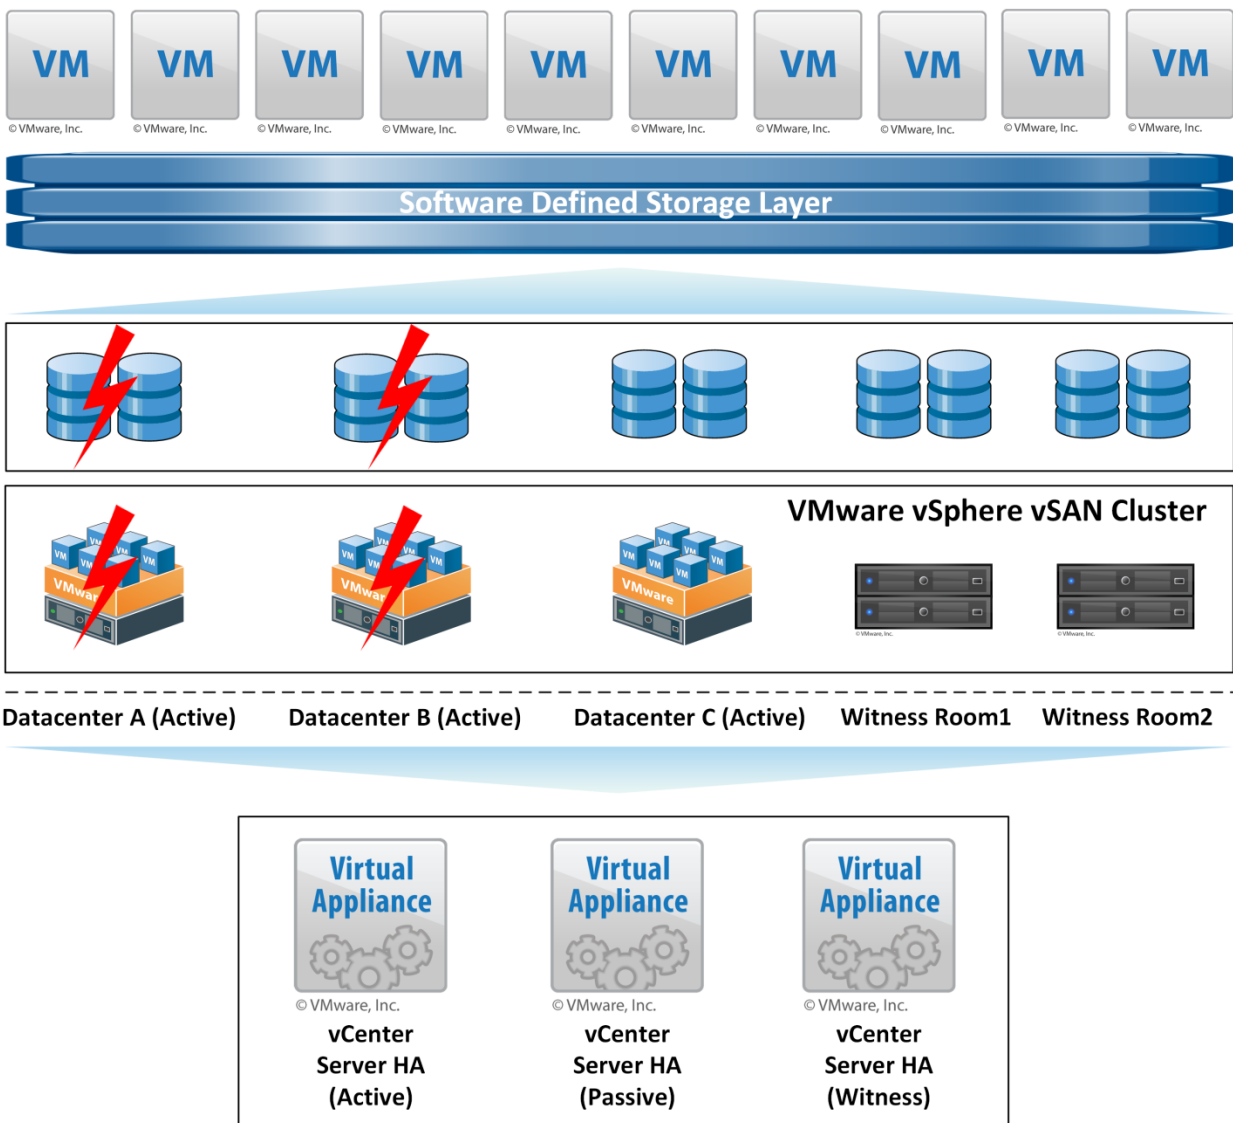


Figure 13 - Scenario 5 in AAA solution

If all ESXi hosts are faulted at Datacenter A and B, the virtual machines will trigger the vSphere HA feature and then start-up on the ESXi hosts at Datacenter C. These virtual machines consist own copy at Datacenter C.

Scenario Six

In Figure 14, if the vCenter HA active node faulted at Datacenter A, what is the status of virtual machines are running at Datacenter A?

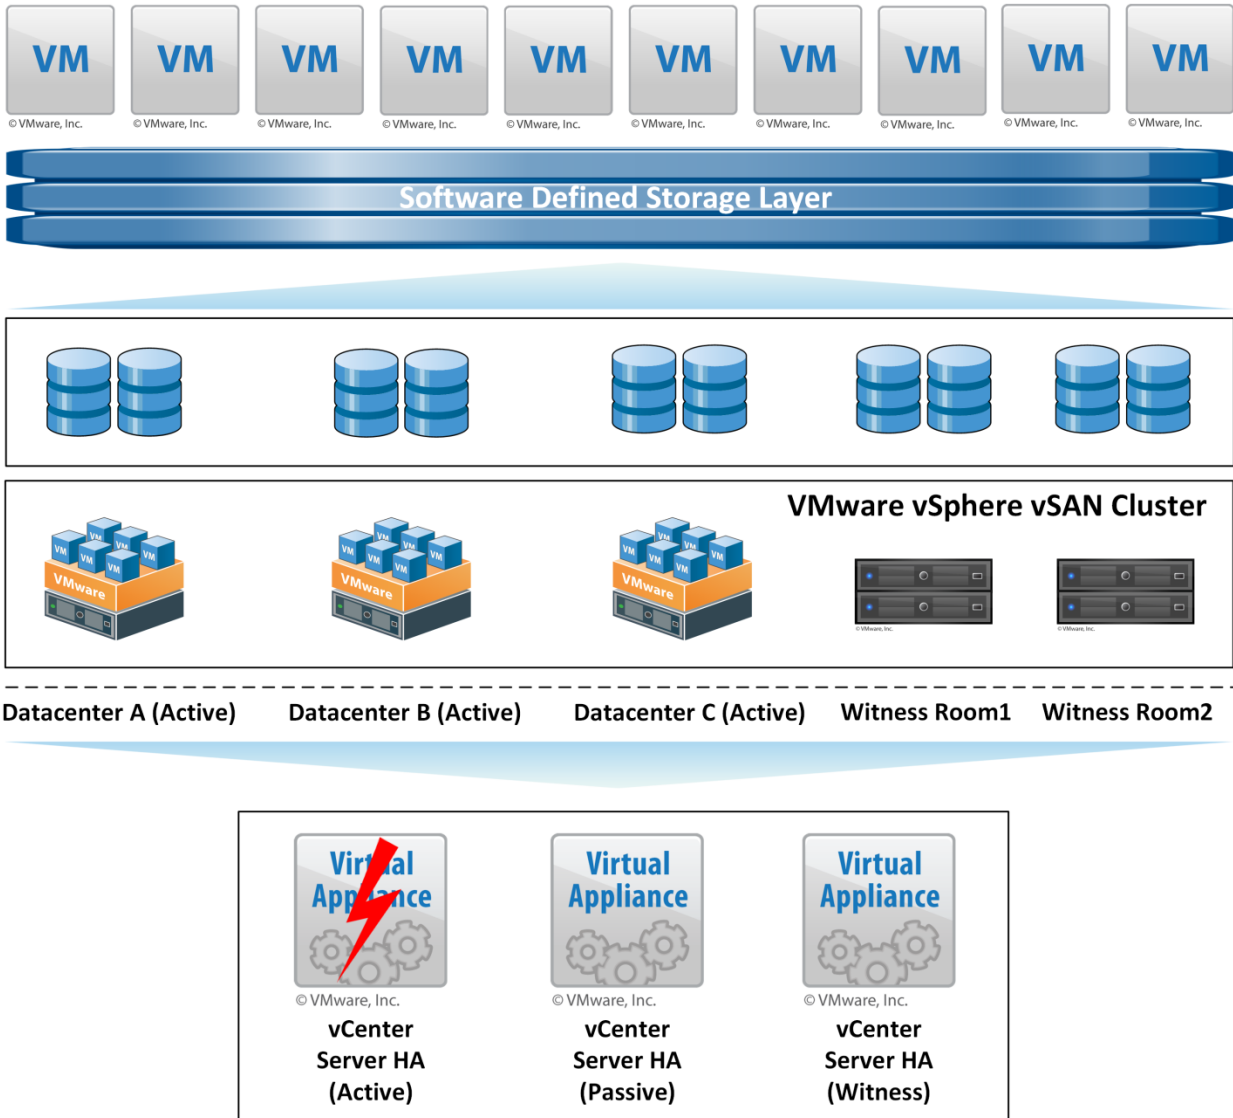


Figure 14 - Scenario 6 in AAA solution

If the vCenter HA active node faulted at Datacenter A, it doesn't affect the service of all virtual machines running at each Datacenter. It only affects the vCenter management service when vCenter HA active node is stopped. The vCenter HA passive node will take the overall service immediately at Datacenter B.

Remark: The failover time of vCenter HA depends on the resources and configuration settings on each vCenter Server Appliance.

Scale-Out

In this flexibility of vSphere vSAN cluster, you can easily set up the storage volumes into different SLA requirements at any time, e.g. Active-Active-Active, Active-Active-Passive, Active-Active-Passive and Passive-Active-Active. Table 12 is the feature summary for multiple Active Datacenters.

Solutions	Features	SLA Requirements
Active-Active-Active	vSphere vSAN Cluster enabled with 5 Fault Domain, FTT = 2 and FTM = RAID-1.	It consists own copy across each Datacenter premises.
Active-Active-Passive	vSphere Metro Storage Cluster.	It consists of two copies across two Datacenters and a second backup copy of the production data store at the 3rd Datacenter.
Local DC premises HA protection	vSphere HA and FT feature enabled on local Datacenter.	It consists of only one copy at local Datacenter.

Table 12 – Feature summary for multiple Active Datacenters

When the volume capacity of the vSphere vSAN cluster is over 80%, you need to plan the expansion of the node in the vSphere vSAN cluster. You can easily add the new vSAN nodes into existing vSAN Cluster with vCenter Server HTML5 (Figure 15); service interruption of vSAN cluster is not required.

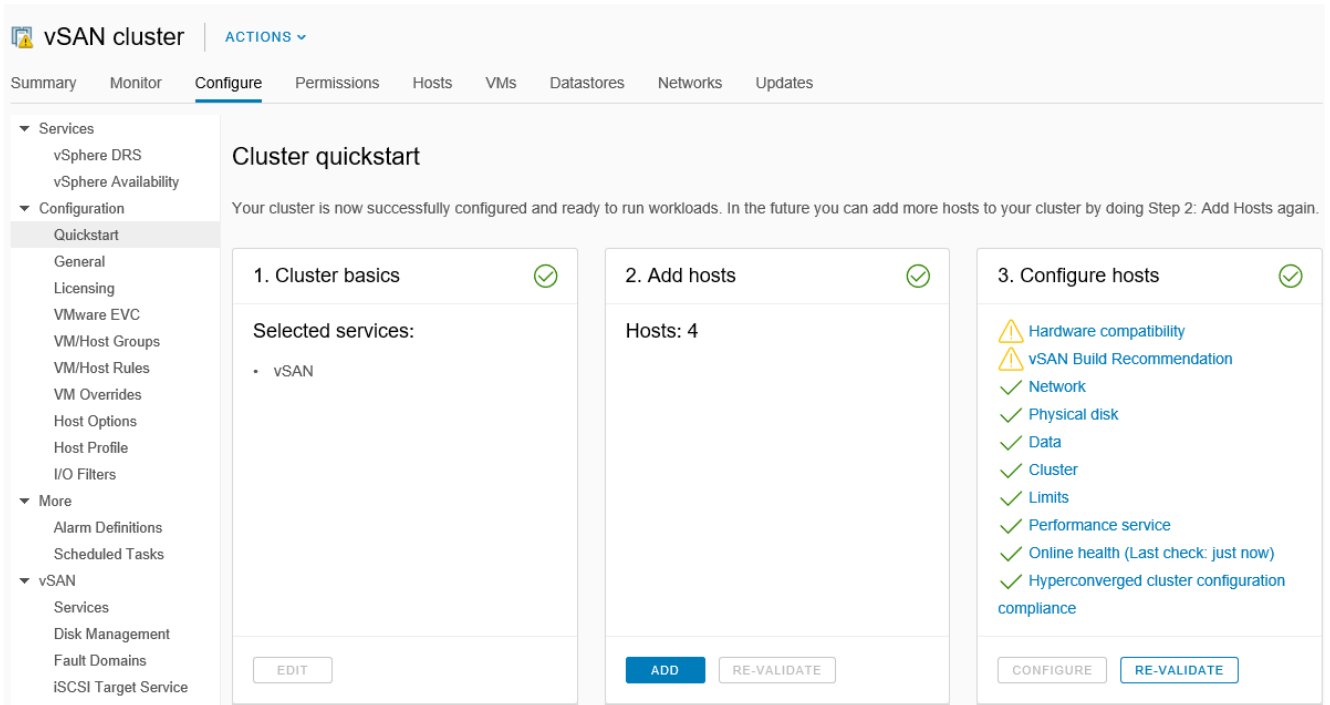


Figure 15 - Configuration of vSAN cluster in vCenter Server

If you want to change the SLA requirement for each virtual machine, you need to update the vSAN storage in Storage Policy-Based Management (Figure 13) and apply the storage policy into each virtual machine again.

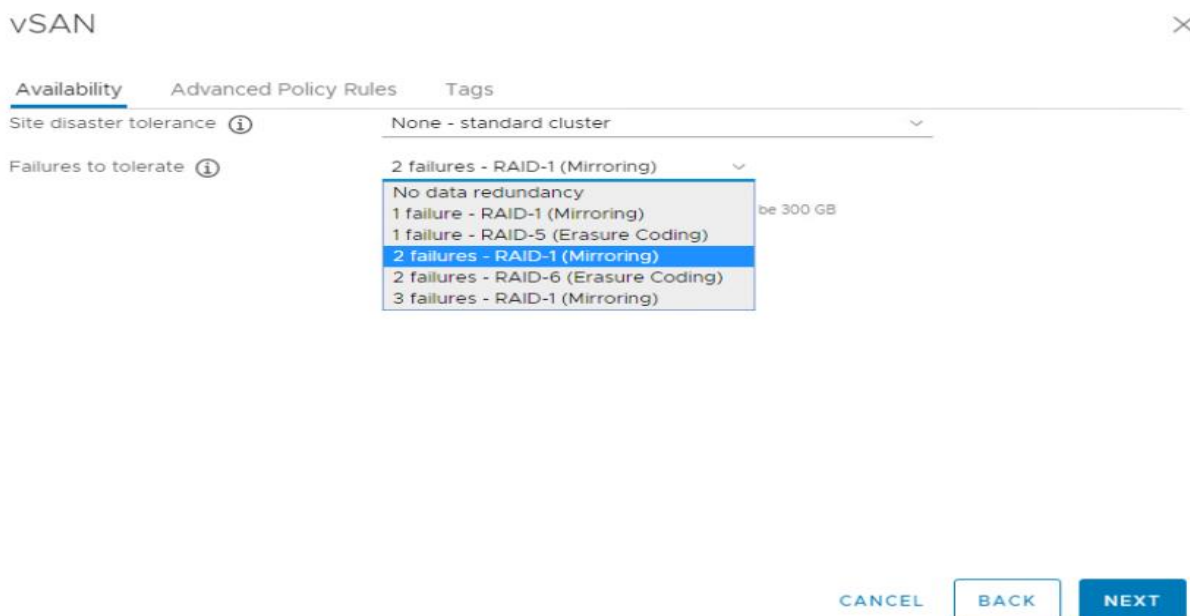


Figure 16 - vSphere Storage Policy-Based Management

For details of Failures to tolerate (FTT) in vSAN policy, refer to Table 13.

RAID Configuration	Failures to Tolerate (FTT)	Minimum Hosts Required
RAID-1 (Mirroring) This is the default setting.	1	3
RAID-5 (Erasure Coding)	1	4
RAID-1 (Mirroring)	2	5
RAID-6 (Erasure Coding)	2	6
RAID-1 (Mirroring)	3	7

Table 13 - RAID Configurations, FTT, and Host Requirements

Summary

This article described how to plan and build the multiple active Datacenters solution. Readers will gain an understanding of how to prepare the software and hardware requirements for multiple active datacenters solution as well as how to design the architecture of active-active-passive and active-active-active solutions. Also discussed were the main benefits, what they can provide and how to help the customers get the most value on multiple active Datacenters, e.g. performance, management, availability, flexibility, simplicity, etc. This article is intended for system architects and system engineers to plan and design software-defined storage and multiple active datacenters solutions.

Bibliography

Administering VMware vSAN

<https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.virtualsan.doc/>

vSAN Planning and Deployment

<https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.vsan-planning.doc/>

Dell EMC Technical White Paper Dell EMC SC Series Storage: Synchronous Replication and Live Volume

https://downloads.dell.com/manuals/all-products/esuprt_software/esuprt_it_ops_datcentr_mgmt/general-solution-resources_white-papers7_en-us.pdf

vSphere Availability

<https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.avail.doc/>

Managing the VMware Cloud on AWS Data Center

<https://docs.vmware.com/en/VMware-Cloud-on-AWS/services/com.vmware.vsphere.vmc-aws-manage-data-center.doc/GUID-EDBB551B-51B0-421B-9C44-6ECB66ED660B.html>

Dell Technologies believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." DELL TECHNOLOGIES MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying and distribution of any Dell Technologies software described in this publication requires an applicable software license.

Copyright © 2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.