

VMware® Stretched Metro Cluster with VMstore Synchronous Replication & Automatic Failover

VMstore™ Best Known Methods

Revision History

Version	Date	Author	Notes
1.0	June 2020	Harrison Waller	Initial Release

Table of Contents

Introduction.....	4
Target Audience.....	4
Solution Overview	4
Infrastructure Architecture	6
VMstore Synchronous Replication.....	6
Witness Setup	7
vSphere Configuration	12
vSphere HA	12
Host Status and Isolation	13
Permanent Device Loss.....	16
All Paths Down Scenario	16
vSphere DRS.....	17
vSphere Storage DRS	20
Failure Scenarios	21
Single-Host Failure in Site A	21
Single-Host Network Isolation in Site A	22
Storage Partition	23
Data Center Partition	24
Full Storage Failure in Site A	25
Permanent Device Loss.....	25
Full Compute Failure in Site A.....	26
Loss of Site A	27
Summary.....	28

Introduction

Disaster recovery is just as important in a data center as is the data itself. There are a number of various solutions to address recovery from any type of outage. A well-known method is with VMware's SRM product which is a robust site and application recovery software suite. This paper, however, looks at a unique concept called a metro stretched cluster which allows you to split an infrastructure across two sites but maintain unified network and storage. At the core of this solution is Tintri VMstore's implementation of synchronous replication.

In the VMware authored paper - [VMware vSphere® Metro Storage Cluster Recommended Practices](#) VMware does an outstanding job of laying out the framework of the solution and discusses it in detail. It's highly recommended that you have read or have that document handy as you use this one. An effort has been made to parallel that document to make understanding and implementing the solution as easy as possible.

When considering the stretched cluster's major benefits, they are:

- *Downtime avoidance*
- *Site workload balancing and mobility*

Note: VMware has an excellent document on helping to decide whether a Stretched Metro Cluster or using VMware SRM is the best solution for your environment - [Stretched Clusters & VMware Site Recovery Manager](#).

The main focus of this document will be around the storage configuration with a strong emphasis on aligning VMware DRS, HA and affinity settings.

Target Audience

The focus of this document is for administrators wanting to implement a VMware metro stretched cluster in their own datacenter with a Tintri VMstore. They need to be familiar with VMware ESXi, vCenter, networking and the VMstore as well as VMware clustering techniques.

Solution Overview

Following the approach of the VMware vMSC best practices guide we'll be looking at a uniform configuration of the stretched cluster since it's the most commonly implemented type. In this design both the management and storage networks are stretched across both datacenters, so all hosts have simultaneous access to both Tintri VMstores (Shown in Figure 1).

In a steady state each host accesses the network and storage locally while each VMstore replicates to the partner array. This minimizes traffic between data centers to avoid the performance impact of reads traversing sites.




Figure 1 - VMstore in a VMware Metro Stretched Cluster

So, with Site A hosts they would have active access to the VMstore physically located in site A. The datastore, through VMstore’s Synchronous Replication, replicates over to the VMstore in Site B. In Figure 1 it shows the replicated VMs as in standby mode. The status and access to the replicated datastore is invisible to all hosts and they are unaware that there is a duplicate copy. Furthermore, the process of automatic failover to the replicated copy is virtually transparent.

The same applies for the hosts in Site B having active access to the datastore on the local VMstore in Site B and having uninterrupted access to the replicated datastore in Site A with an automatic failover.

The vCenter server is assigned site A affinity (datastore is located locally). The quorum server should be placed in a third location and is discussed in the next section along with VMstore’s synchronous replication configuration.

Infrastructure Architecture

In the following sections we'll discuss VMstore's synchronous replication and describe in detail how to configure it. We'll then look at vSphere HA's features on maintaining VM uptime and failure responses. We'll also describe how to configure them accordingly.

VMstore Synchronous Replication

Tintri introduced synchronous replication to its VMstore product several years ago and works between any model array with appropriate firmware. Customers have been using it since and have enjoyed several phases of improvement along the way.

The benefits are unmatched flexibility in:

- *Recovery Time Objectives (RTO) as low as 30 seconds or less for planned and manual failovers*
- *Recovery Point Objective (RPO) is 0 seconds. By definition synchronous replication is an exact copy of the original data.*
- *Support for inter-data center distances up to 100 km or round-trip time of 10ms across a metro cluster network*
- *Support for synchronous and asynchronous replication on the same array simultaneously*
- *Bi-directional as well as one to many synchronous replication where an array can host primary VMs and replica (secondary) VMs for different groups of VMs simultaneously. In other words, a VMstore can be the primary and replication target at the same time.*
- *Increased ROI as replication can be achieved between any model type running appropriate FW.*

With the latest update starting in TXOS release 4.5.2 Tintri VMstore offers synchronous replication with **automatic failover**. Now, recovery or avoidance from a disaster can happen automatically adding to the list of benefits:

- *Disaster avoidance with synchronous replication and automatic failover*

The basic functionality is achieved by creating a Service Group (SG) through the Tintri Global Center (TGC) UI. The service group contains the primary (source) and secondary (target) arrays and the proper replication network settings. A cluster IP is also assigned to the service group. This cluster IP (and folder) is then mounted as a datastore within the ESXi cluster. VMs that are desired for replication are simply created or migrated over to that datastore. During a planned failover or an unplanned failure, the cluster IP easily moves between VMstores.

To assist in the coordination of a failover is the introduction of a quorum device called the witness. The quorum is the arbitrator in negotiating failover and determining who the primary location should be so as to avoid a split-brain situation. This is a downloadable rpm package that runs on Centos 6.x with at least 100MB in its root partition. It's easily installable and configurable.

Note: It's Best Practice to run the witness at a separate site. If the witness is too closely co-located to one of the VMstores, then it won't have the safety or independence in case of site perturbations. The cloud would be an ideal location.

Witness Setup

Setting up the witness is very simple with the following commands from a command shell. Here are some of the basic commands:

Installation

```
# rpm -i <witness-server.rpm>
# initctl start tintri-witness
```

Setting the AuthKey

```
# /opt/tintri/bin/quorumcmd set-authkey tintri123
# quorumcmd: setAuthKey { tintri123 } succeeded

# /opt/tintri/bin/quorumcmd print-authkey
# quorumcmd: getAuthKey succeeded: local authkey: tintri123
```

Changing the Port (default 9095)

```
# /opt/tintri/bin/quorumcmd set-server-port 4315
# set-server-port: setIncomingPortNumber( 4315 ) succeeded

# /opt/tintri/bin/quorumcmd print-server-port
# quorumcmd: getIncomingPortNumber succeeded; server port : 4315
```

Note: Do not use port 9097 as that's reserved internally for the quorum server process

Opening Firewall Ports

```
iptables -I INPUT 1 -p tcp --dport 9095 -m conntrack --ctstate NEW,ESTABLISHED -j ACCEPT

iptables -I OUTPUT 1 -p tcp --sport 9095 -m conntrack --ctstate NEW,ESTABLISHED -j ACCEPT
```

With the witness configured now you can configure a Service Group (SG) that will manage the synchronous replication of the datastore containing the desired VMs. To do this simply login to the Tintri Global Center (TGC) UI and click on the **Service Groups** tab at the top. Click on the **Add Service Group** button from the **Actions** pull down menu and enter in a name for the service group as shown in Figure 2.

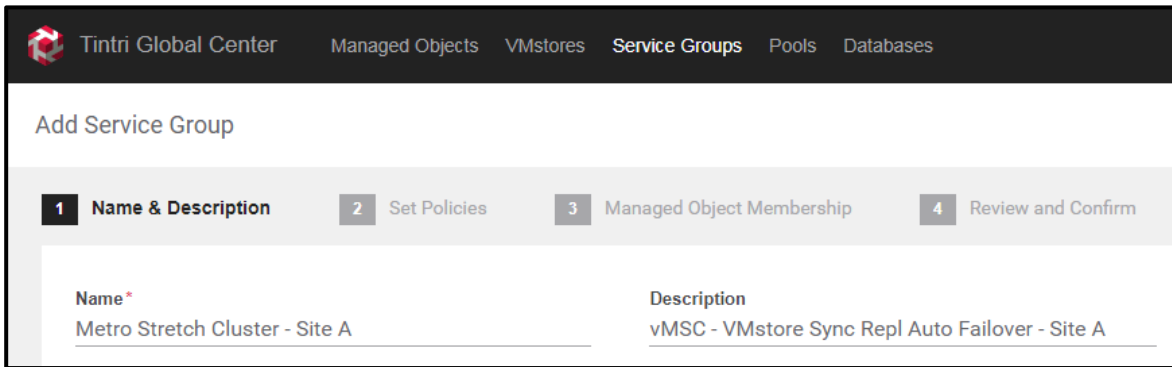


Figure 2 - Create Service Group Name

Note: Take special care in naming service groups so that they are clearly understandable. One method is to have the name match the folder/directory of the underlying datastore.

On the next screen select Option 2 which is the option for synchronous replication (Figure 3).

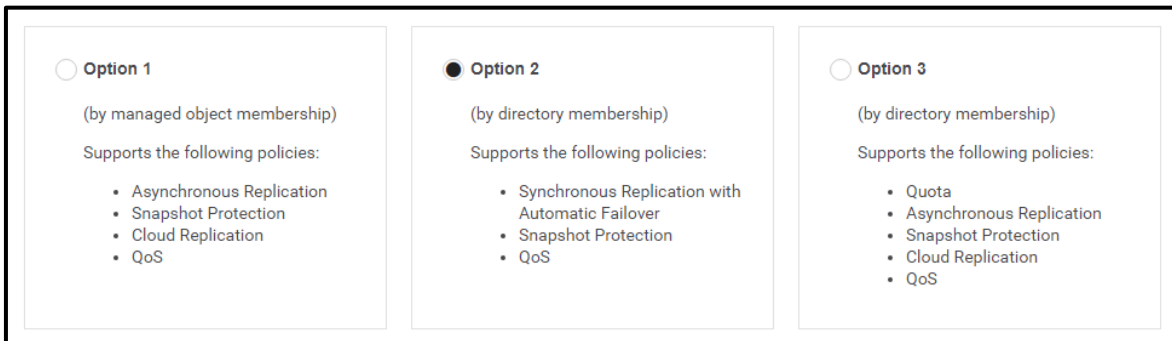


Figure 3 - Pick Synchronous Replication

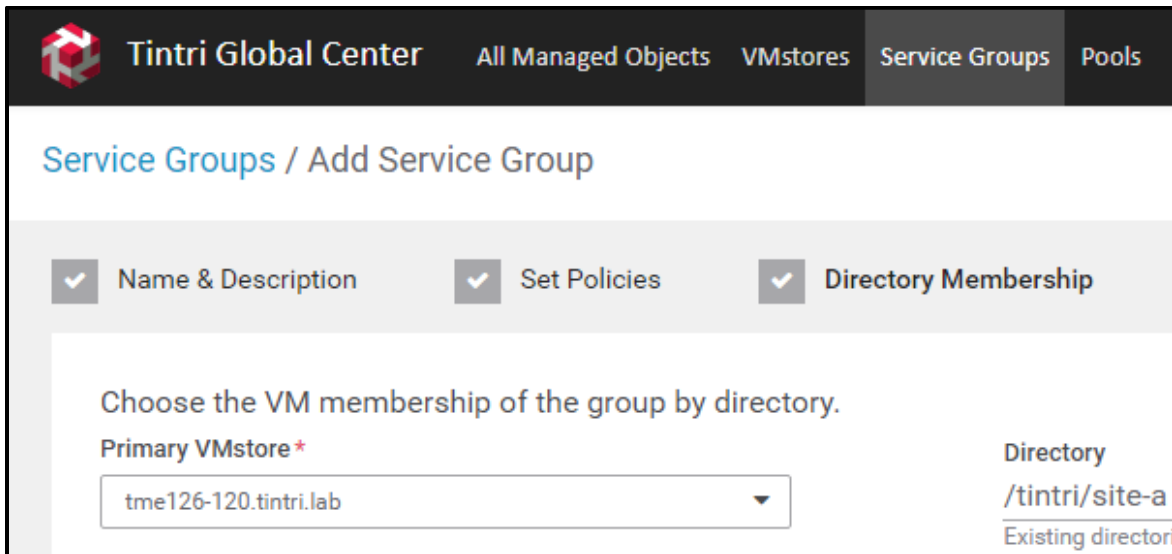


Figure 4 - Choose Primary VMstore

On the next screen pick the VMstore that is going to be the primary array in the replication. Using the pull-down menu select the appropriate VMstore (Figure 4). This can be changed later with a manual failover. Add the directory of the datastore that will contain the VMs to be replicated. The directory can be empty at this point. Simply VMotion or create any VM in this datastore that are needing replication.

To that end, synchronous replication is datastore specific, replicating only the VMs in the named datastore. On the next screen (Figure 5) select the secondary VMstore using the pull-down menu. Add an IP address that'll be used for the replication cluster. Now pick the replication network or IPs to be used for the data traffic. You may test the replication settings if desired.

The screenshot shows the 'Add Service Group' configuration page in the Tintri Global Center. The 'Protection' tab is active, indicated by a '4' in a black box. The configuration is divided into several sections:

- DESTINATION**: A dropdown menu for 'Secondary VMstore*' is set to 'tme127-120.tintri.lab'.
- CLUSTER IP**: 'Cluster IP (Datastore)*' is '4.15.1.100' and 'Netmask*' is '255.255.255.0'.
- REPLICATION**: 'Source Replication IP*' is '4.15.1.130 (data)' and 'Destination Replication IP*' is '4.15.1.120 (data)'. Both have 'CREATE NEW' buttons below them.
- Gateway** and **VLAN ID**: These fields are currently empty.
- Notification**: A green bar at the bottom states 'Replication test passed' with a close button (X).
- Buttons**: A 'TEST REPLICATION SETTINGS' button is located at the bottom left.

Figure 5 - Create Cluster IP

On the next screen (Figure 6) you'll need to add a witness. Use the pull-down menu and select **Add New Witness**. If you're going to use a witness that has already been configured, then you may simply select that from the menu list.

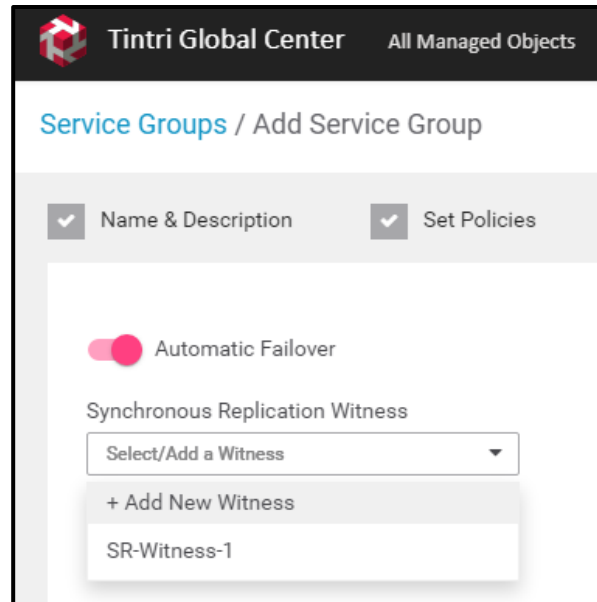


Figure 6 - Turn on Automatic Failover

Continuing with adding a new witness, on the **Configure Witness Server** screen (Figure 7) enter a name, the IP address, port number and authorization key from the witness server. (**Note:** The default port number is 9095)

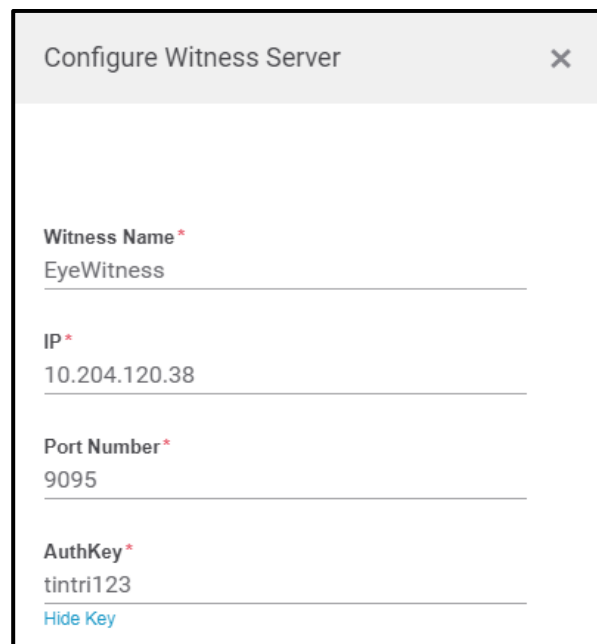


Figure 7 - Add witness

To ensure everything is working correctly click the **Test Witness Configuration** button. The next screen (not shown) determines when out of sync notifications are sent. Set according to your needs however Tintri recommends 1 minute because an out of sync replication is very much business critical. Review the final summary screen and click **Add Service Group**.

After the creating a service group you can manage and monitor it by clicking on **Service Groups** and then selecting the appropriate service group (Figure 8).

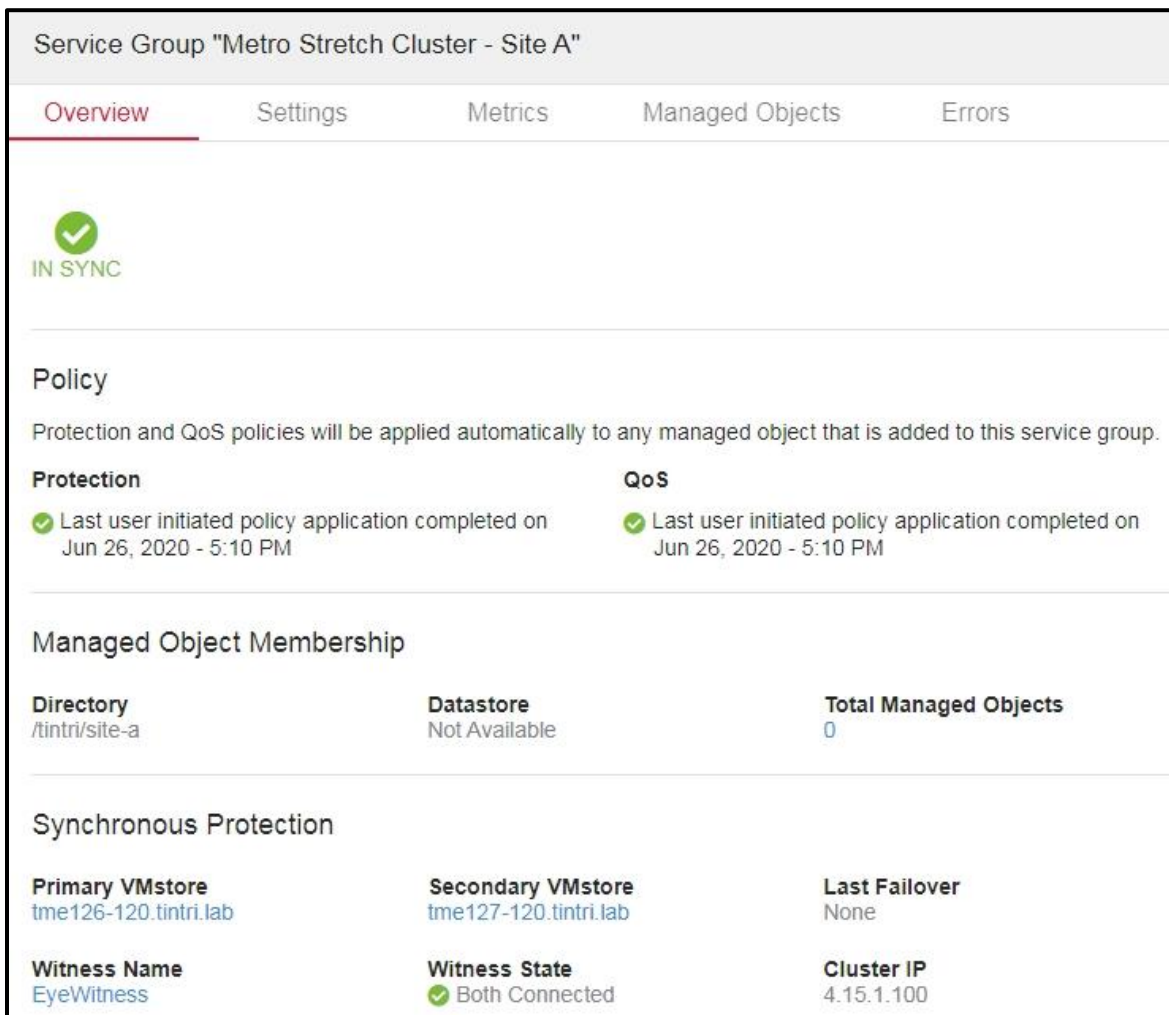
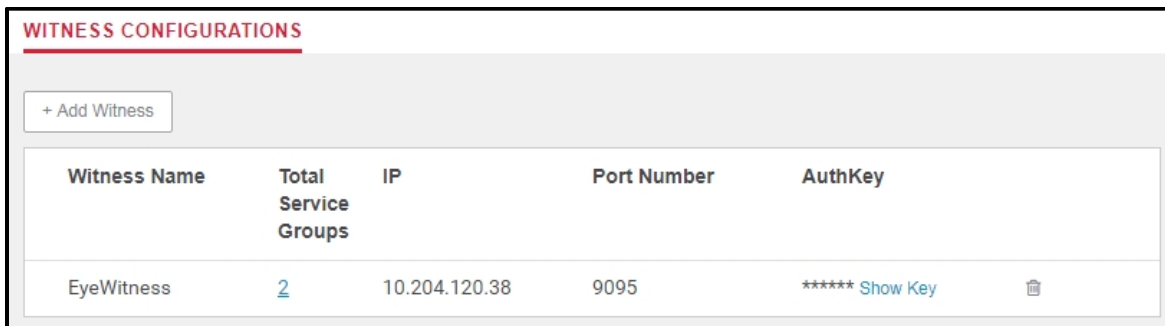


Figure 8 – Service Group Overview

To manage witnesses, click on **Explore, Settings** and then select **Witnesses** (Figure 9).



Witness Name	Total Service Groups	IP	Port Number	AuthKey
EyeWitness	2	10.204.120.38	9095	***** Show Key

Figure 9 - Witness Configurations

vSphere Configuration

vSphere has several mechanisms for maintaining the uptime and accessibility of VMs. Using a combination of VMware vSphere best practices and failure testing we have compiled a list of best practices that give the most resilient environment. They are discussed in detail in the following pages.

vSphere HA

From the VMware side of this configuration vSphere HA is the fundamental component to enabling a metro cluster solution. It's a stable and robust solution with what can seem like an endless set of choices for ensuring a stable environment during failures. Each of the sections, leveraging VMware's guidance, discusses how each component should be configured in the context of using Tintri's VMstore and synchronous replication with automatic failover.

Admission Control

Ensuring that either site can handle a complete site failure requires the use of vSphere HA admission control. It reserves enough CPU and memory to ensure VMs can run on the remaining hosts in case of host failure. To that end a policy setting of 50 percent for both memory and CPU is recommended. You'll need to override the default calculated quantity as shown in Figure 10.

VMware's recommendation to use a percentage as opposed to a specific number of hosts definitely gives the most flexibility and simplicity of management.

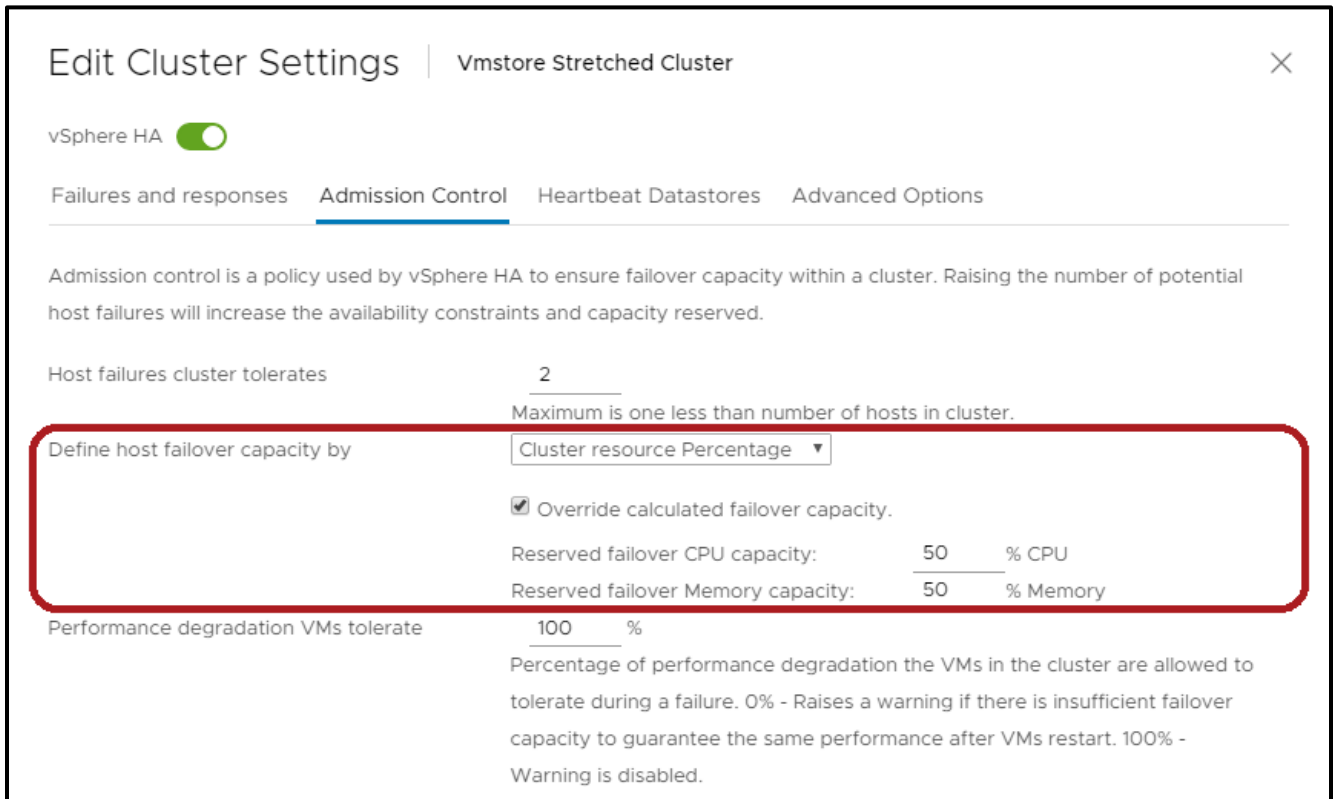


Figure 10 - vSphere HA Admission Control Settings

Host Status and Isolation

To be able to initiate failure remedies vSphere HA has a number of tools to determine a host's availability and its access to storage. There are two methods for which vSphere HA determines this status:

1. *Network heartbeat*
2. *Datastore heartbeat*

The sequence is shown in Figure 11 and are from two perspectives, the HA master (light blue) and a slave host (maroon). The HA master tracks the status of other host members in the cluster through a network heartbeat (A) then datastore heartbeat (B) and finally an ICMP ping to the host management address (C).

The host determines its status by looking at whether or not it's been isolated. It does that by first checking its heartbeat with the master and other hosts in the cluster (1). If there's no network heartbeat it'll ping an isolation address (2). Isolation addresses are configurable from within vSphere and it's recommended to have at least one at each site.

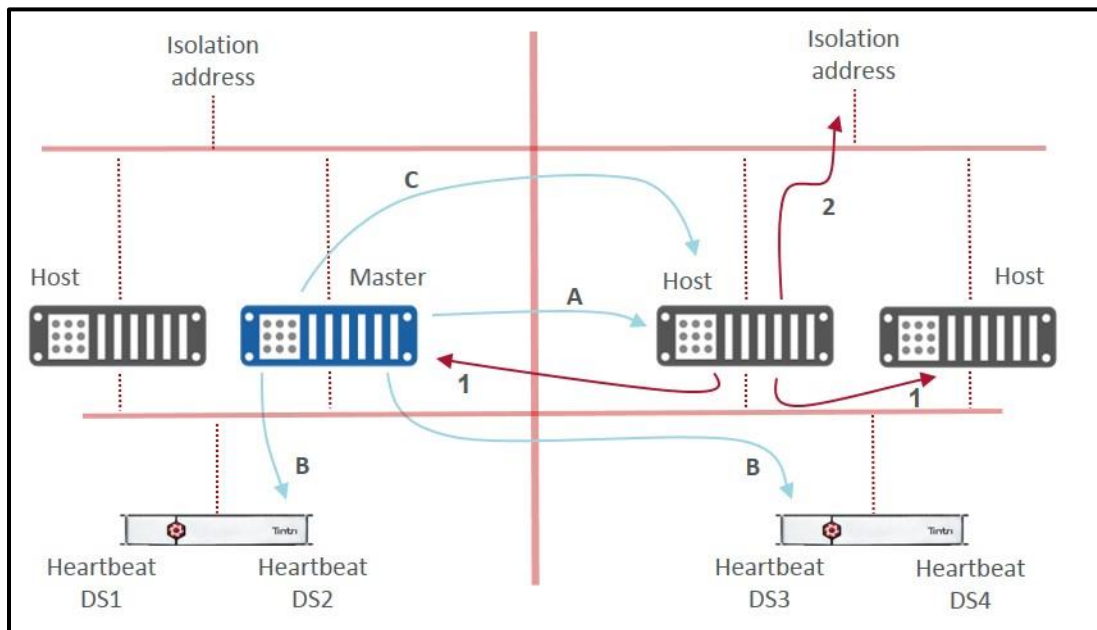


Figure 11 – Master and Host Status Determination

Note: The detail in Figure 11 is minimized to avoid distraction. However, the cluster hosts can and will check both the isolation and heartbeat addresses in the opposite site to ensure its specific isolation state.

Configuring an isolation address entails adding a parameter **das.Isolationaddress** under Advanced Options as shown in Figure 12. It's recommended that you have at least 2, one at each site in the cluster to account for complete network failure at any one site.

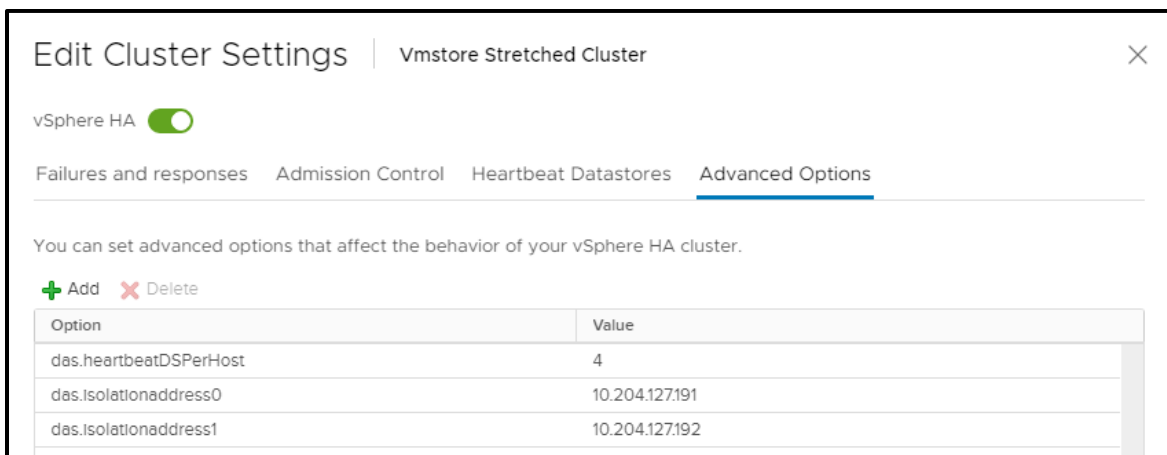


Figure 12 - Advanced Options

VMware suggests using 4 heartbeat datastores and to use more than the default of 2 you need to add a parameter **das.heartbeatDSPerHost** as shown in Figure 12. Once that's set the heartbeat datastores are selected under the same titled tab as shown in Figure 13. Two of the datastores in the figure are the service group datastores mounted through the cluster IPs. The other 2 are data IPs that simply mount the root folder

from each array. This is the most efficient way to achieve the 4 heartbeat datastores without using up the synchronous replication service groups which has a maximum of 16.

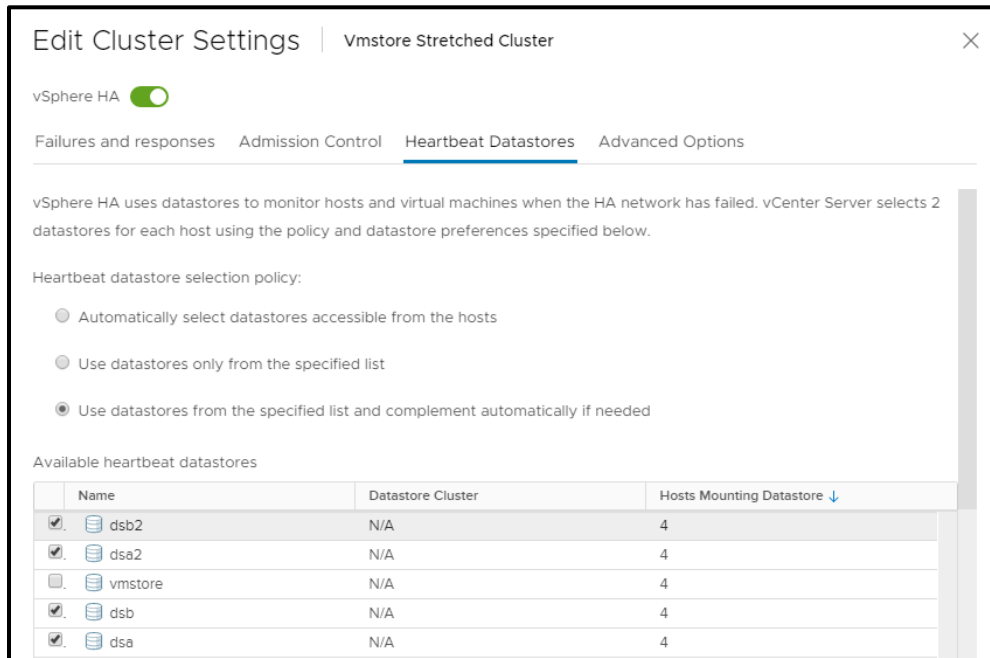


Figure 13 - Heartbeat Datastores

VMware advises to leave the VMs powered on when a host isolation occurs reasoning that with today's redundant environments, they are very rare. However, when one does occur it could prove quite disruptive

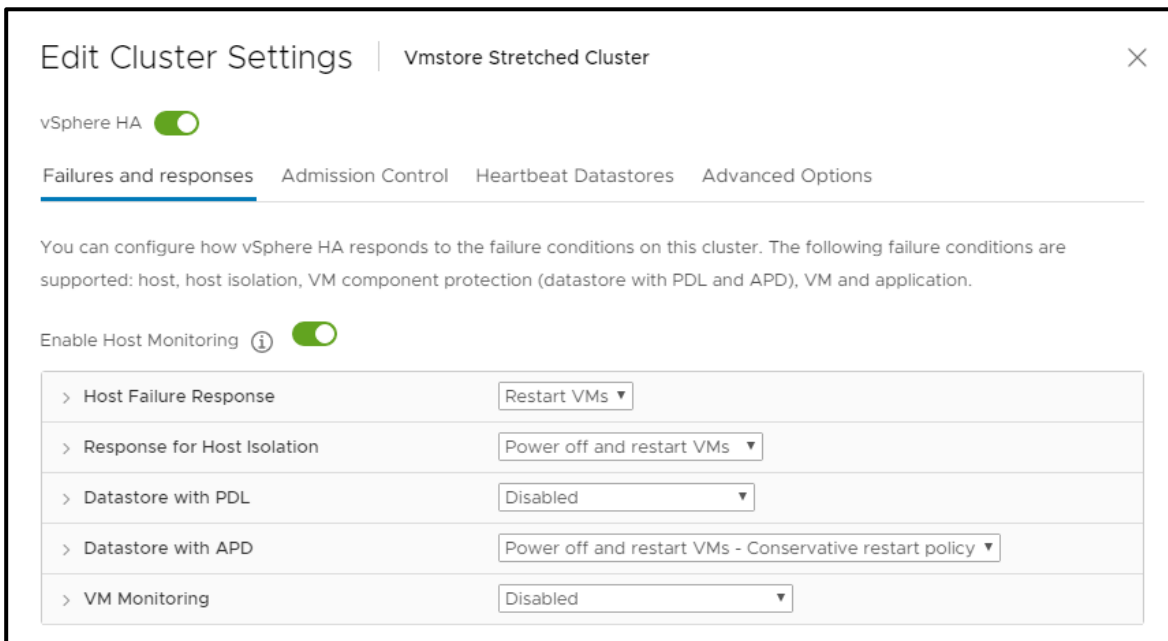


Figure 14 - Host Monitoring Settings

depending on how the infrastructure is configured. Internally the ESXi host may still have access to the VM but externally, there might not be access to that VM. Therefore, Tintri recommends **Power off and restart VMs** as shown in Figure 14.

Permanent Device Loss

Since we're looking at the NFS offering of VMstore in this paper the Permanent Device Loss (PDL) does not apply because it's for block devices which send SCSI sense codes about the status of each path. Therefore, this feature can be turned off as shown in Figure 14.

All Paths Down Scenario

Since PDL is not applicable for monitoring and triggering VM restarts we'll need to take advantage of the All Paths Down feature. The sequence for determining when corrective action is taken is shown in Figure 15.

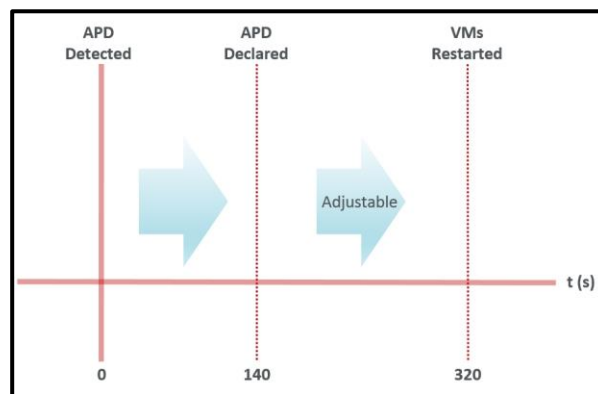


Figure 15 - All Paths Down Sequence

After an all paths down is detected vSphere waits 140s to officially declare a full APD. Once the APD is declared a new 180s (default but configurable) timer starts counting. At the end of 180s if no recovery has happened then the defined corrective action is taken. If at any point after the APD (at 140s) and before the corrective action (at 320s) a recovery occurs, then the **Response Recovery** is performed – either Restart VMs or Disabled.

The APD of 140s can be an eternity in any environment when being down is not an option. Business requirements should drive the responses to both of these settings. However, a conservative approach is recommended for the APD response - **Power off and restart VMs (Conservative)**. The setting for the **Response Recovery** can vary greatly based on the infrastructure. There may be environments that can handle intermittent errors or transient perturbations. Testing or further exploration may be needed to satisfy business requirements. (Both settings shown in Figure 16).

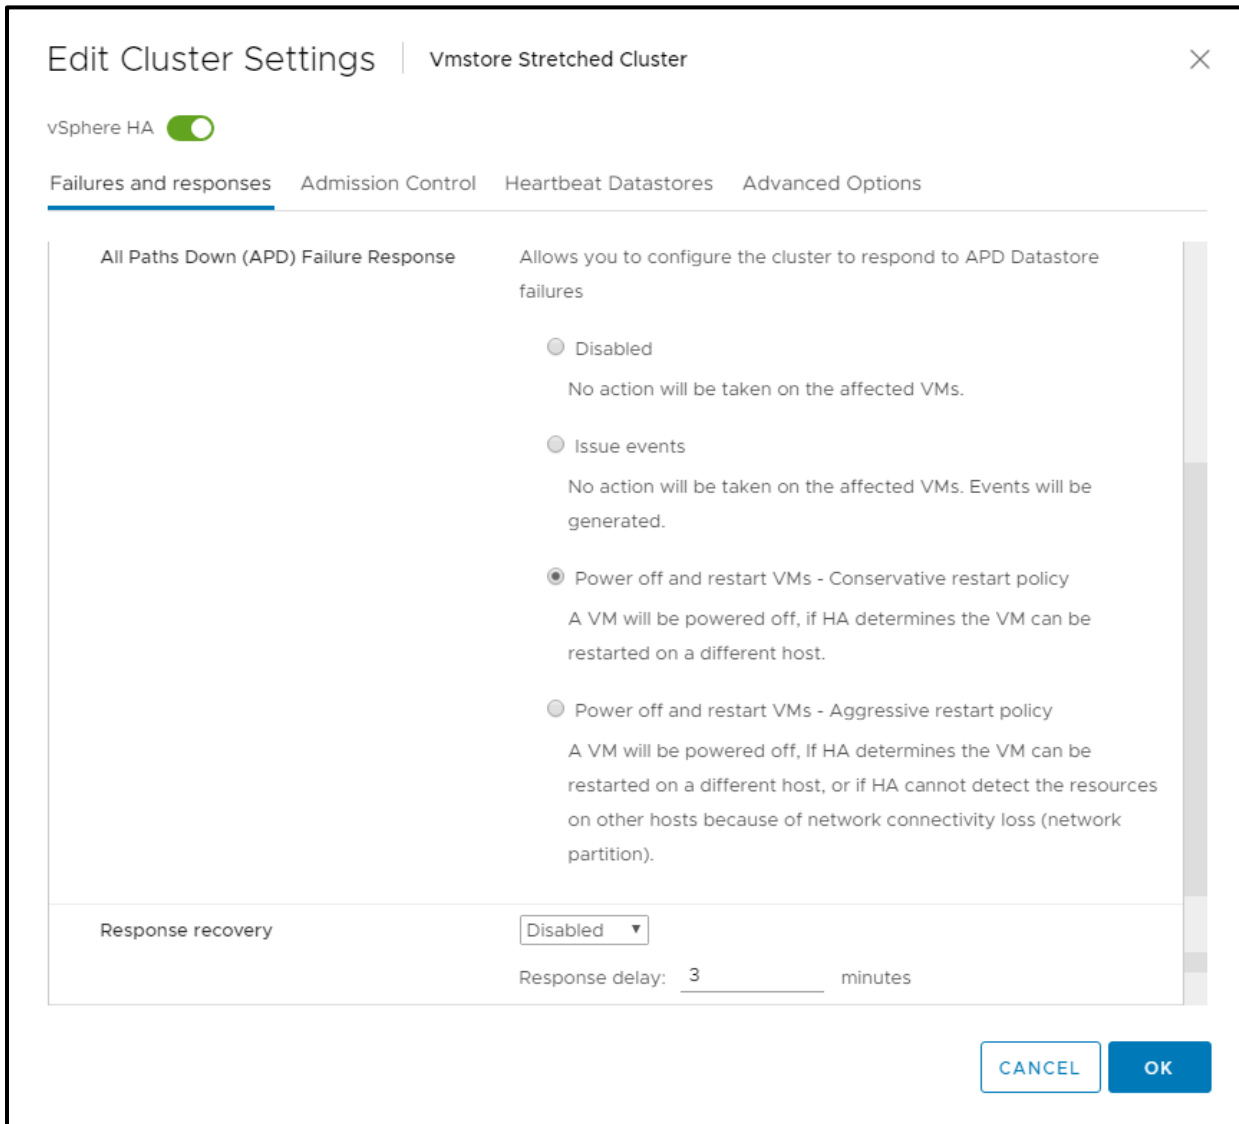


Figure 16 - All Paths Down Settings

vSphere DRS

vSphere DRS is used to distribute load within the cluster. When an ESXi host becomes overly taxed DRS will migrate VMs across the remaining hosts in the cluster. Since each site has its own identity, locality, and set of resources we'll use vSphere DRS affinity rules to place a specific site allegiance on each VM. Configuring these components can be tricky and we'll cover them in detail as well as discuss how each behave during failures in the next section.

Host Group Creation

The first step is to create site A and Site B host groups that house that site's ESXi hosts. To create a host group, click on the ESXi cluster and then in the right window click on the Configure tab and finally click on **VM/Host Groups** (Figure 17).

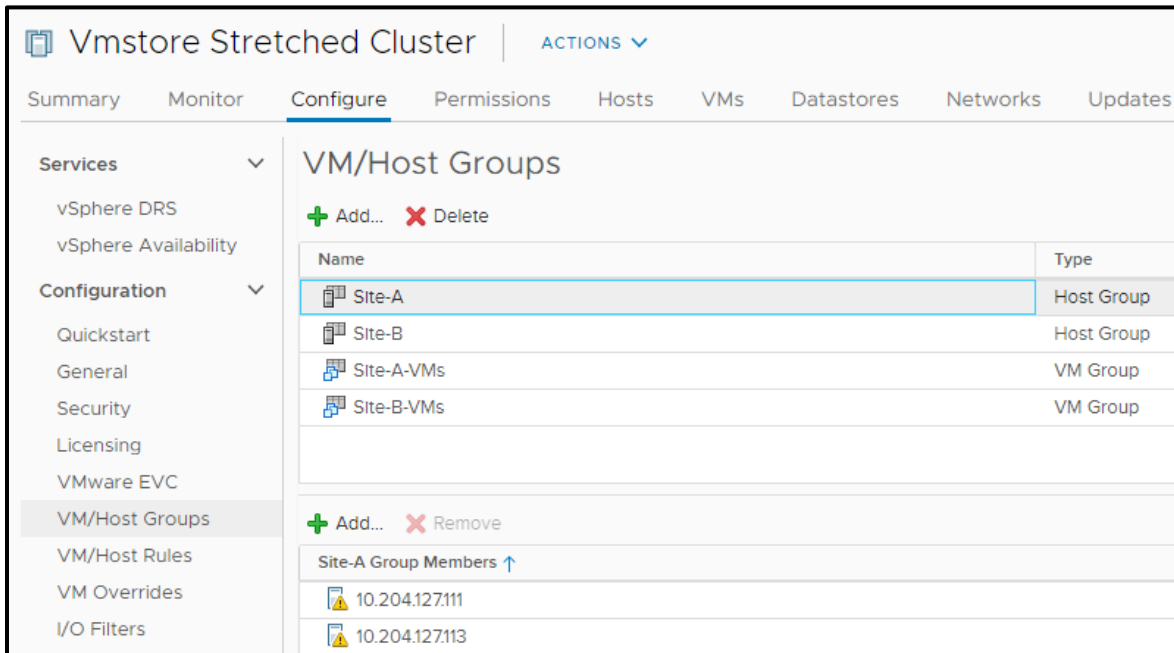


Figure 17 - Host Group Creation

The rest of the configuration is straight forward. The end result is 2 host groups with 2 ESXi hosts each for Site A and Site B respectively as shown in Figure 17.

VM Group Creation

Now let's create the VM Groups placing site specific VMs in each group. Following the same steps as in the host group creation this time ensure the group type is set to **VM Group** and select the appropriate site's VMs to add to each group (Figure 18).




Figure 18 - VM Group Creation

VM Host Rule Creation

Now that the host and VM groups are created let's tie them all together with VM/Host rules. These rules are at the heart of keeping VMs up and running. There are 4 choices for placing VMs where 2 are inclusive and 2 are exclusive criteria. We'll ignore the exclusive rules since we're trying to assign VMs to a specific host group. Looking at the inclusive rules we have *should run* and *must run*. If we were to choose *must run*, then it would limit the ability of VMs to migrate to the other site thereby creating a gap with certain site failures. So, we'll select *should run* so that when there is a host available in that host group DRS should run the VMs there. Now there will be certain scenarios where host saturation can occur, and vSphere DRS will ignore the should rules but they are rare.

To do this click on **VM/Host Rules** under the **Configure** tab of the cluster. Click on **Add** and enter a name of the rule and select the type – **Virtual Machines to Hosts** (Figure 19). For each site select the VM group and host group and give it **Should run on the hosts in group**.

The screenshot shows a dialog box titled "Create VM/Host Rule" for a "Vmstore Stretched Cluster". The "Name" field contains "Site-A-VM-rules" and there is a checked checkbox for "Enable rule.". The "Type" dropdown is set to "Virtual Machines to Hosts". The "Description" is "Virtual machines that are members of the Cluster VM Group Site-A-VMs should run on host group Site-A.". The "VM Group" dropdown is set to "Site-A-VMs". The "Host Group" dropdown is set to "Site-A". The "Should run on hosts in group" dropdown is also set to "Should run on hosts in group".

Figure 19 - VM/Host Rule Creation

When completed you'll have two sets of rules as shown in Figure 20

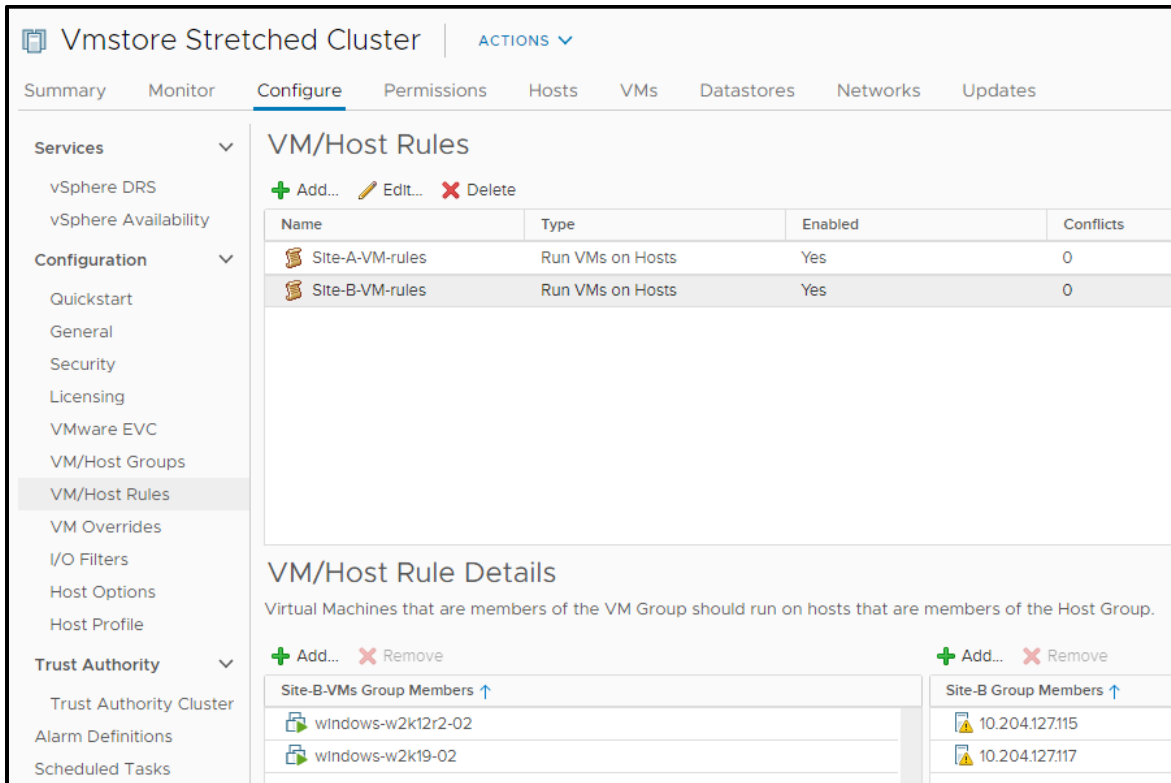


Figure 20 - VM/Host Rules

vSphere Storage DRS

In this solution vSphere Storage DRS is not used as there is no need for aggregation of datastores. The simplicity of the VMstore product is that you're presenting a single datastore to the vSphere cluster. This eliminates VM disk migrations between multiple datastores thereby keeping the resources free for normal I/O and synchronous replication.

To that end VMstore scale-out is not compatible with synchronous replication and the datastore path that the service group presents.

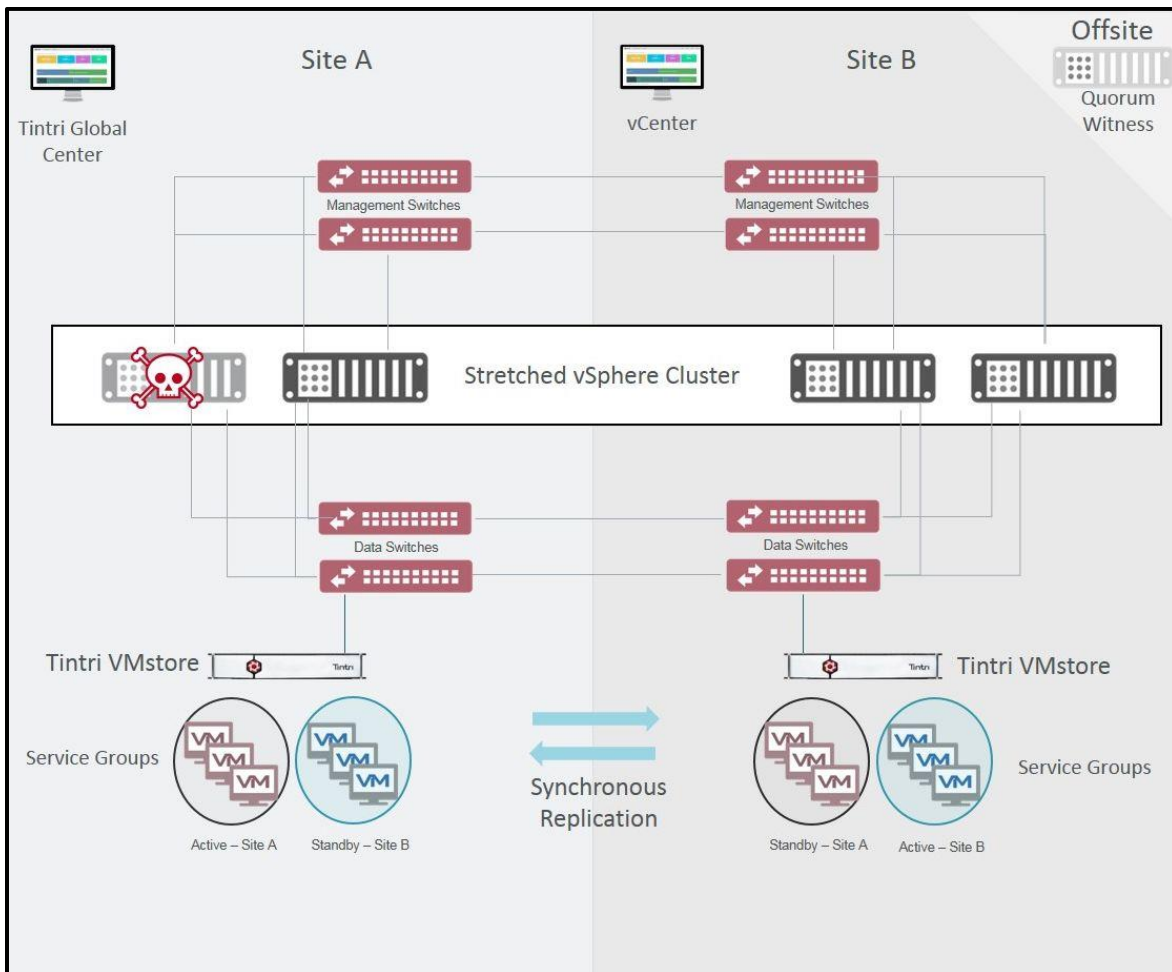
Failure Scenarios

This section covers only the major failure types and are addressed from a quality assurance perspective where we look at expected behavior and the actual result. Then we discuss how the settings achieved the expected behavior.

Single-Host Failure in Site A

Scenario: In this failure a single ESXi host fails leaving the VMs inaccessible

Expected: The VMs should be restarted in Site A per site affinity rules. VMstore cluster is unaffected.

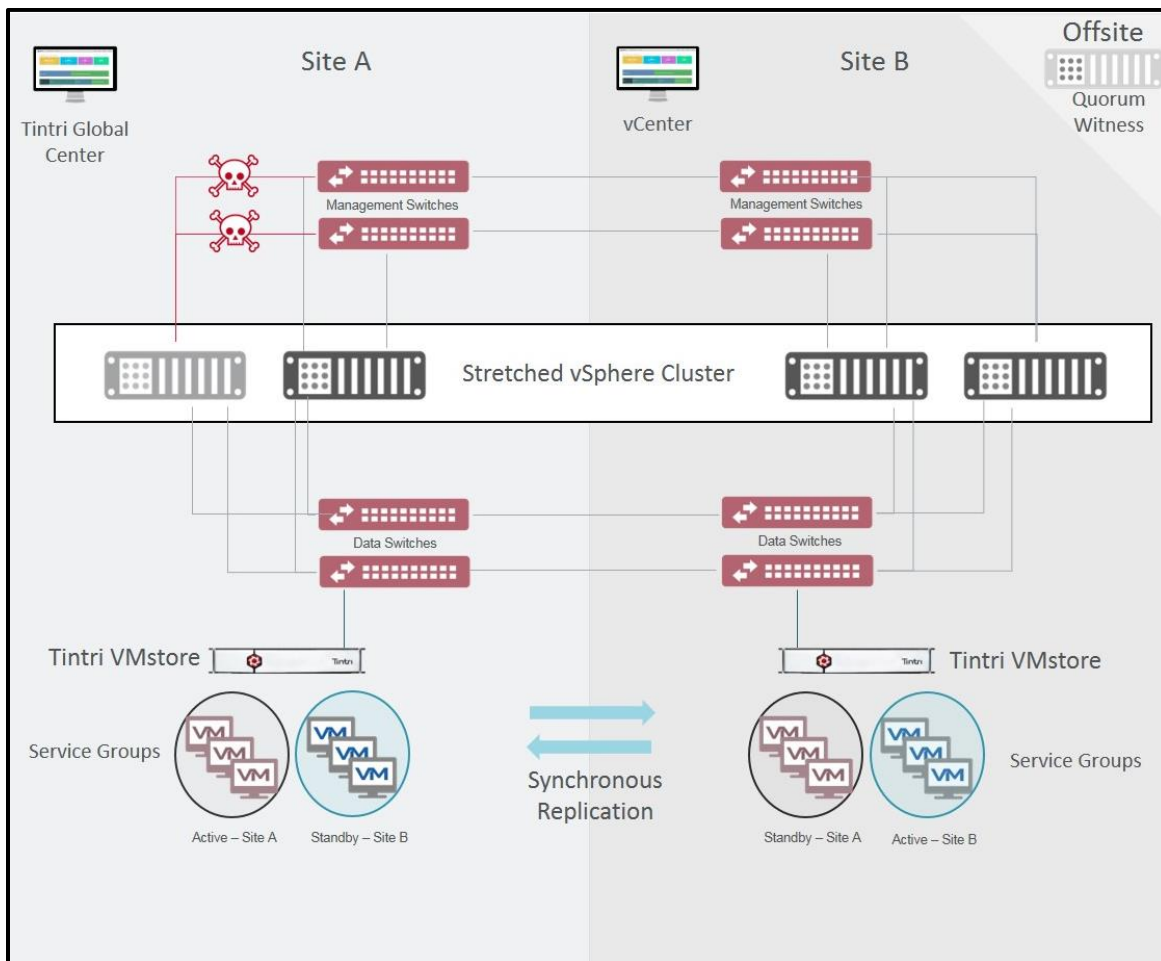


Result/Discussion: The master node went through its sequence of checks to determine the host's status. It lost network heartbeat, datastore heartbeat and finally did not receive any reply from an ICMP ping from the management IP address. At which point it declared the host dead and restarted all the VMs that were running at that host in that site. The VMstore synchronous replication IP is not disrupted in this scenario so both site's datastores are unaffected.

Single-Host Network Isolation in Site A

Scenario: ESXi host loses management connectivity and becomes isolated from the rest of the cluster

Expected: ESXi host realizes it's isolated and performs isolation response. The VMstore cluster is unaffected.

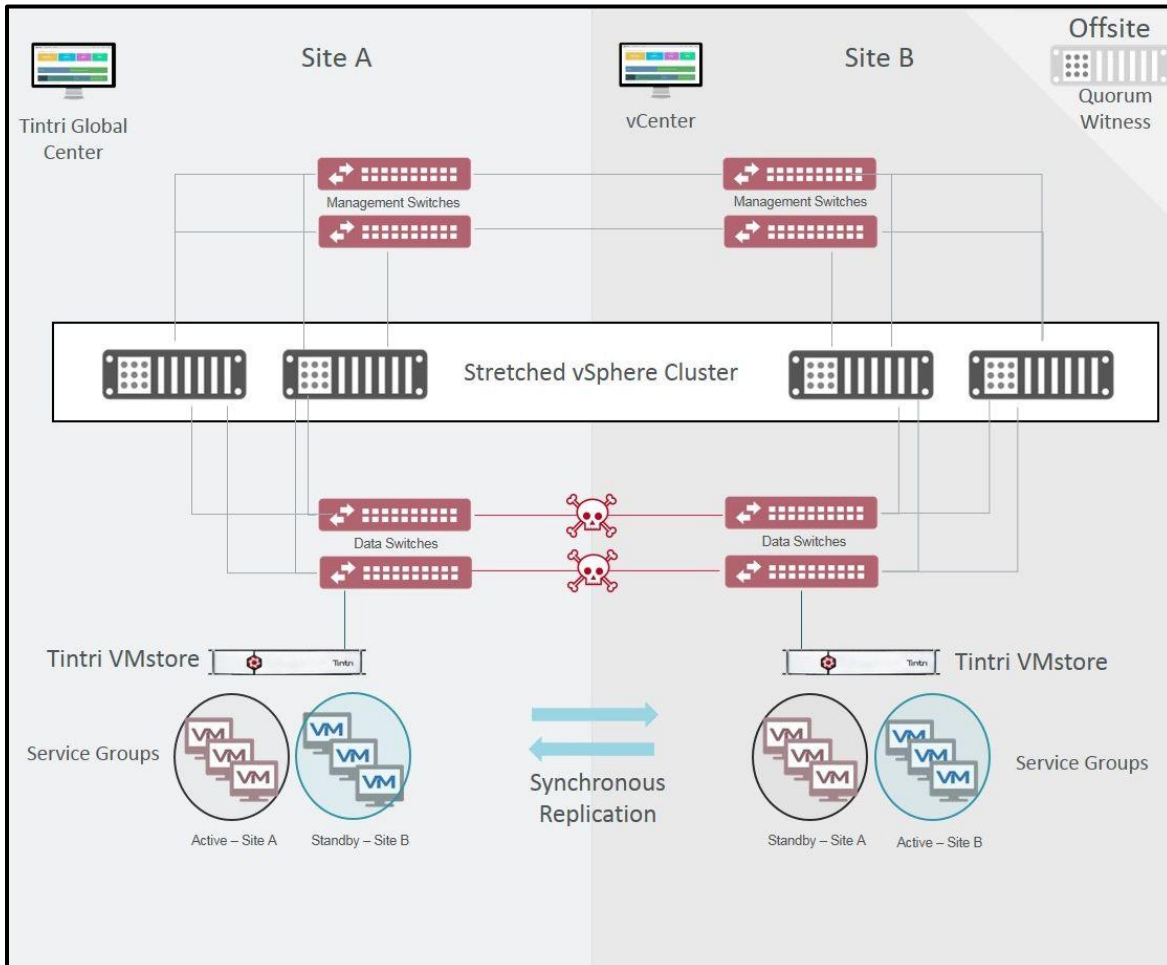


Result/Discussion: This one can be a little confusing based on the language you hear. The term isolation and partition get tossed around as being one and the same. In an isolation scenario, a single host is isolated from the rest of the hosts in the cluster. The vSphere HA master loses network heartbeat with the host but discovers that the datastore heartbeats continue. With host still alive the HA master waits and monitors the VMs. Meanwhile the isolated host determines its state by noticing network heartbeat loss and being unable to ping the isolation addresses. After the APD expires the isolation response is executed. As VMware discusses, an isolated event is very rare in today's redundantly designed networks. However, if the event does happen then it seems prudent to **Power off and restart VMs**. If the network is designed where the management network of the VMs is separate from the data access you could potentially leave the APD response as **Disabled**. Since none of the VMstores nor the witness were affected no action was taken.

Storage Partition

Scenario: Storage network between the data centers is severed.

Expected: Tintri synchronous replication determines that both Site A and B arrays are fine, but replication is no longer possible. VMs do not detect any downtime and continue operating on their assigned hosts.

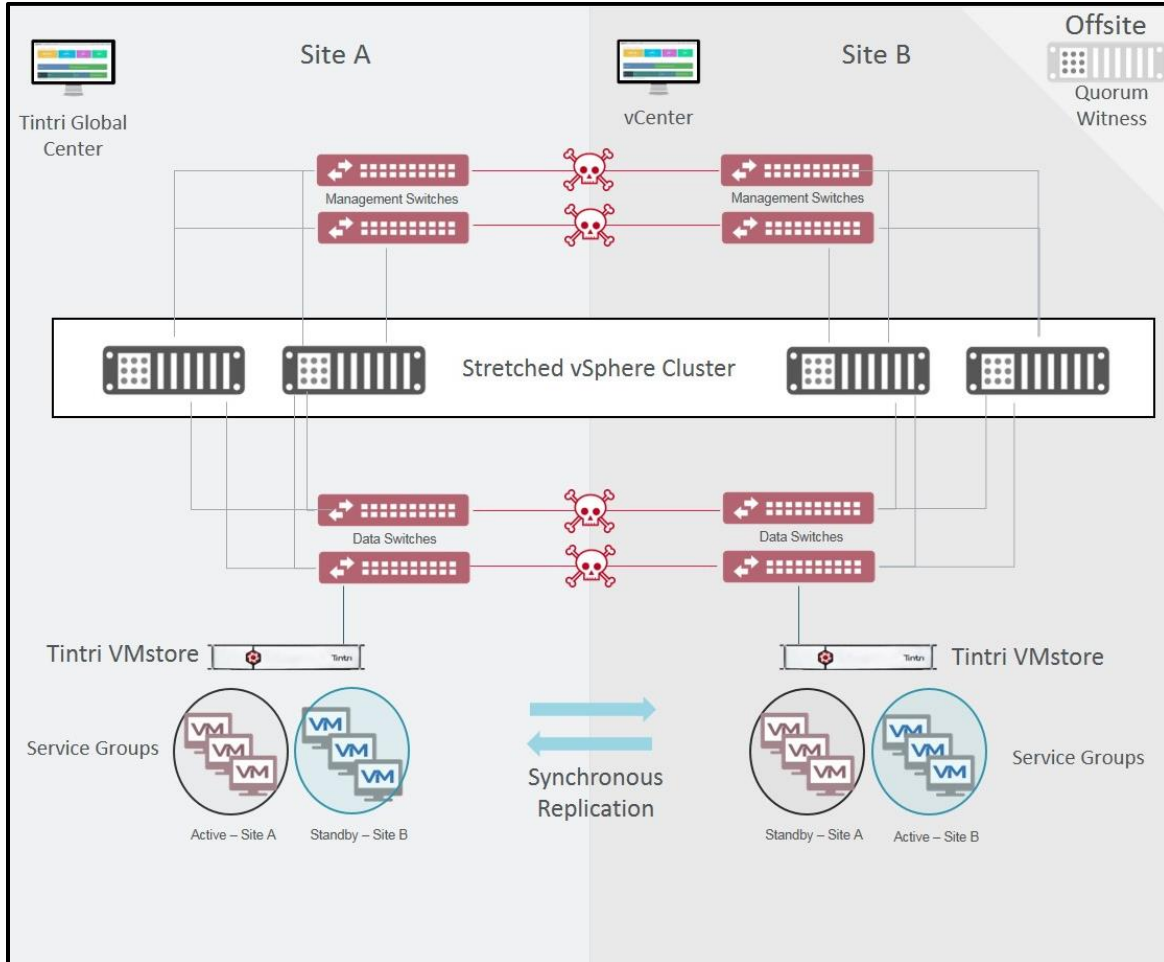


Result: As expected the VMs are not migrated as they were assigned to the appropriate ESXi host which could access the proper VMstore in each site. If, before the failure, a VM had been assigned to a host in the opposite site of its share then an APD will occur. With the APD response set to power off and restart a VM it would be restarted with the proper site affinity. The VMstores determine they are fully functional, but the replication link is down. Upon correcting the error, the replication will reconnect and resync.

Data Center Partition

Scenario: Site A network and storage get completely severed from Site B

Expected: Tintri synchronous replication is blocked. All VMs maintain uptime.

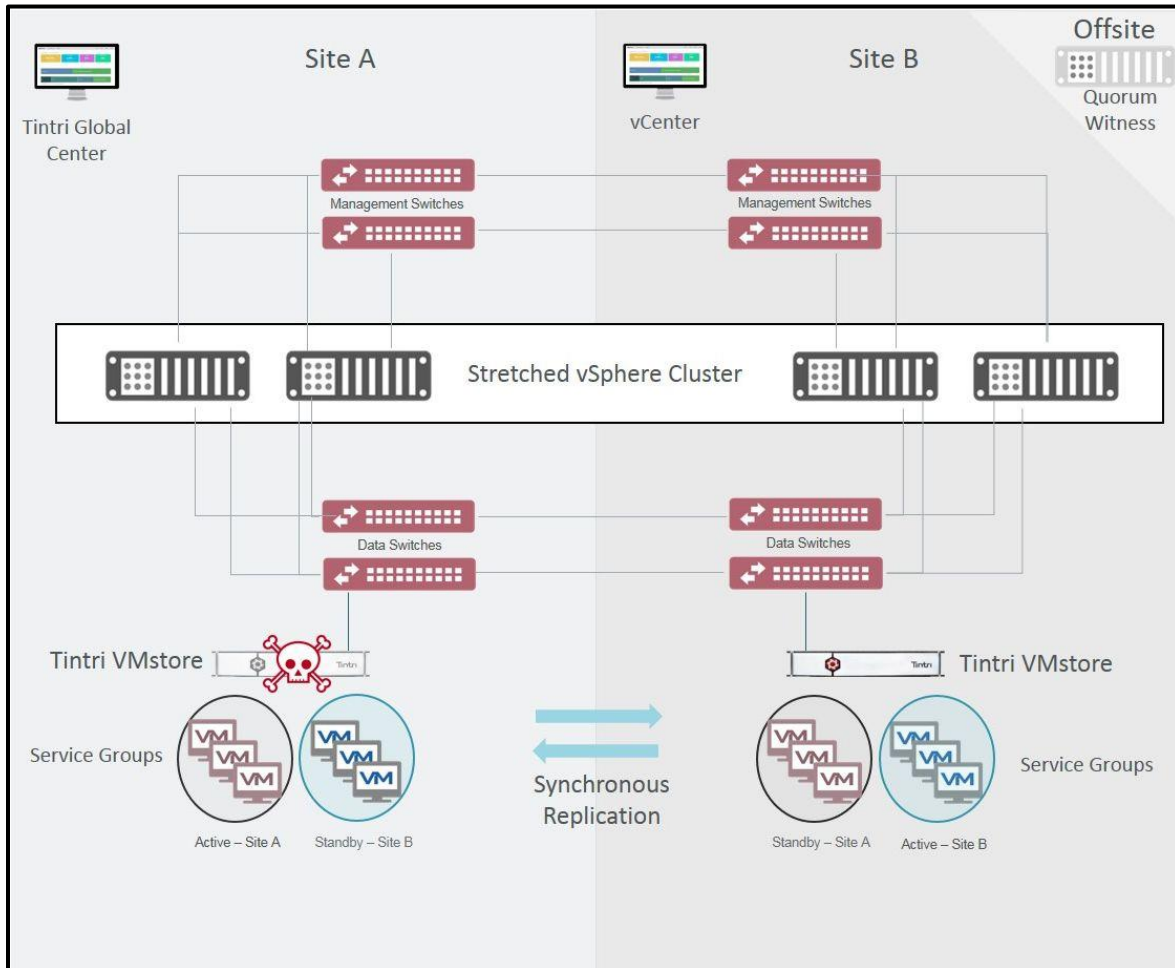


Result/Discussion: Tintri synchronous replication determines both arrays are fine, but replication is now blocked. All VMs maintain uptime because their affinity was set properly and were placed accordingly with access to the VMstore with the VM's VMDK. As in the storage partition failure, if a VM had failed site affinity settings and was running in the opposite data center then there is a chance that after the APD response is triggered a split-brain scenario can occur. vSphere HA will eventually resolve the split brain however avoiding a split brain should be of utmost priority. vSphere DRS triggers every 5min and should be closely monitored to ensure site affinity is aligned properly. VMware has a lot of recommendations on how to best achieve this.

Full Storage Failure in Site A

Scenario: Storage at Site A completely fails leaving access to that VMstore impossible

Expected: Tintri synchronous replication determines Site A array is down and promotes Site B array to primary in the Service Group. VMs do not detect any downtime and continue operating.



Result/Discussion: As expected after a 25s period the remaining VMstore at Site B and the witness declare Site A's VMstore is down and promote Site B to Primary and the copy of the replicated Site A VMs becomes active. The VMs continue I/O without a hiccup. Depending on the array type and the site affinity of the VMs the performance could be slightly lower. Remember Site A VMs are assigned to the hosts in Site A. So, as long as that criteria are maintained the VMs will continue to be serviced by ESXi hosts in Site A even though the physical location of the VMs is in Site B.

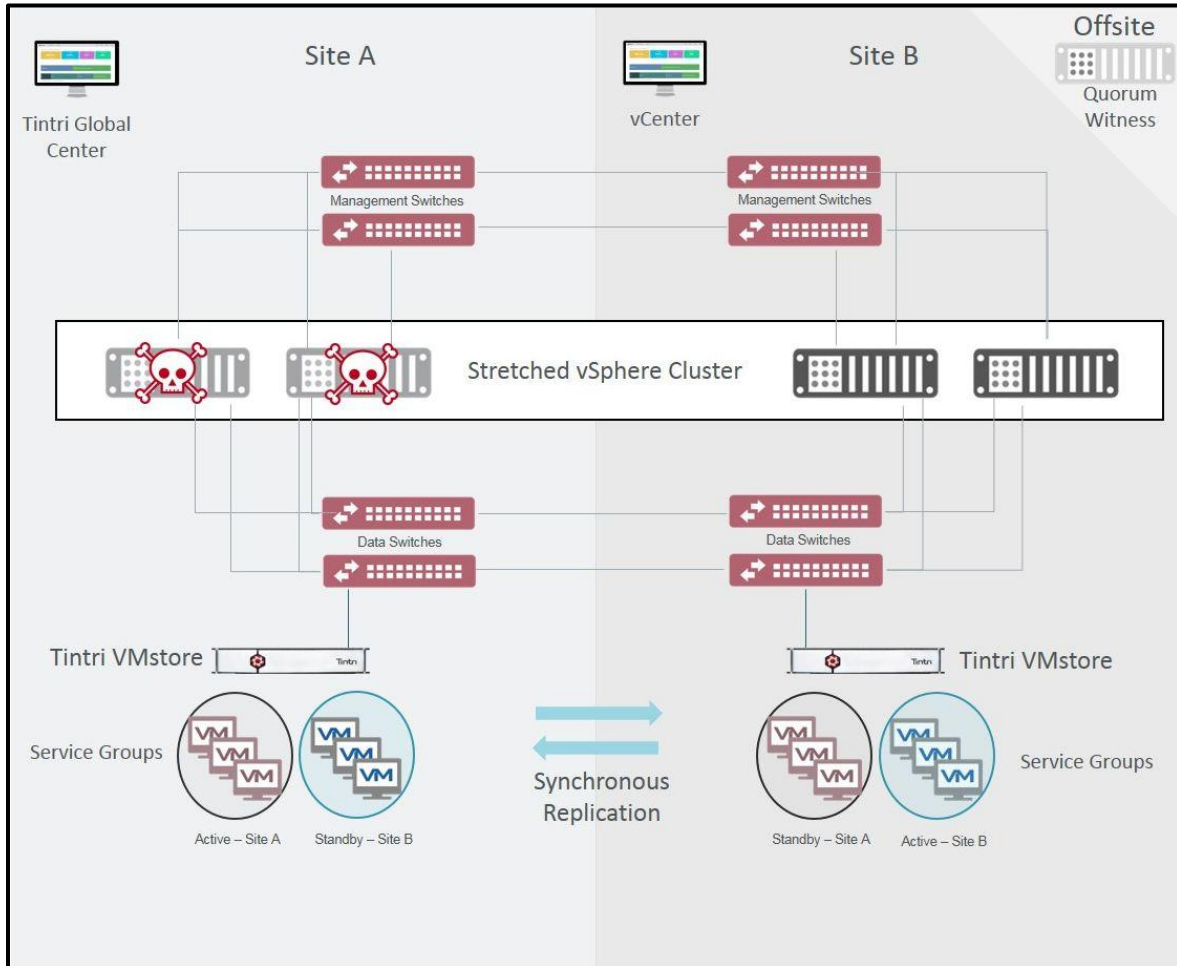
Permanent Device Loss

Since PDL is not supported for NFS shares there is no failure scenario to test.

Full Compute Failure in Site A

Scenario: In this failure all ESXi hosts fail in Site A leaving the VMs managed by those hosts inaccessible.

Expected: The VMs should be restarted in Site B per site affinity rules. VMstore cluster is unaffected.

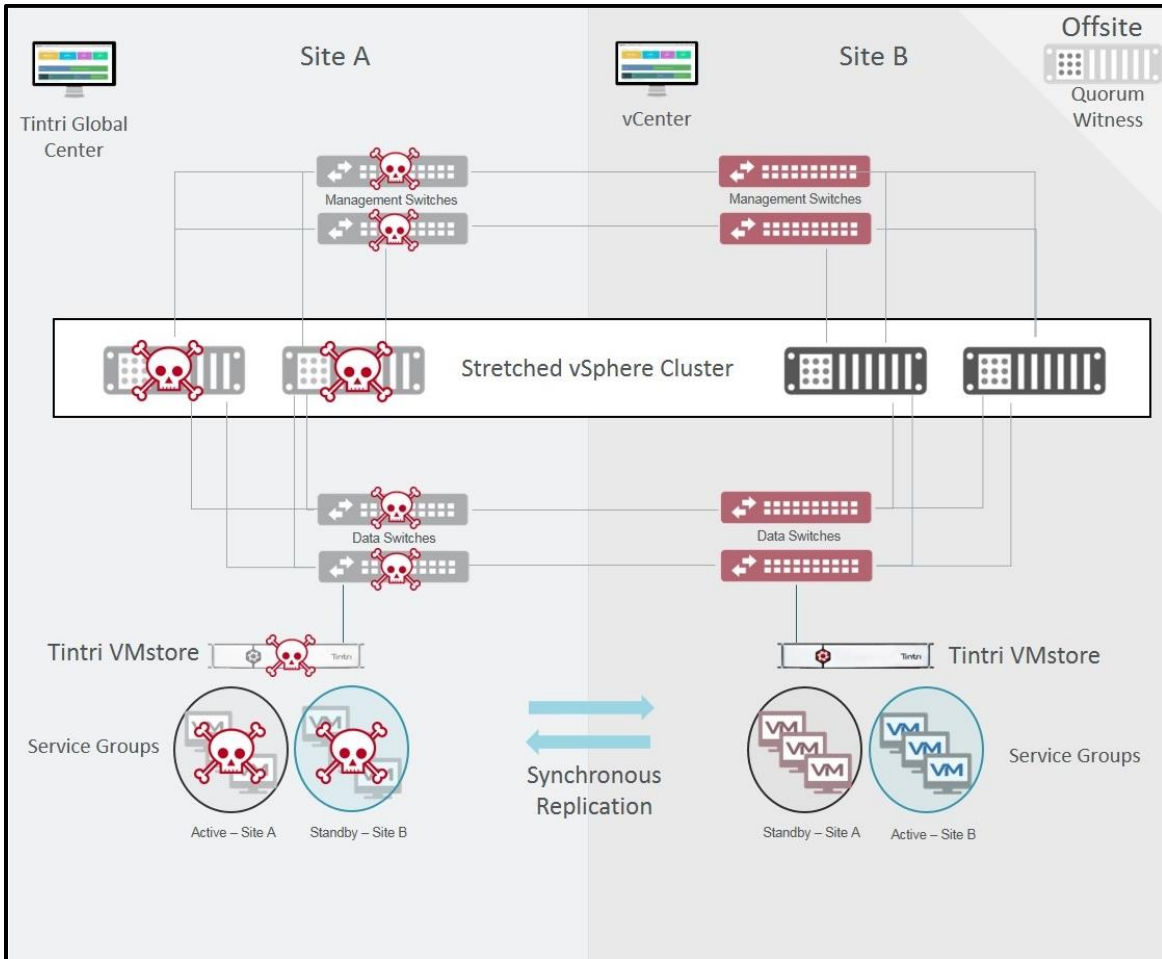


Result/Discussion: Similar to the single host failure, the vSphere HA master (electing one if needed) went through its sequence of checks to determine the hosts' status. It lost network heartbeat, datastore heartbeat and finally did not receive any reply from an ICMP ping from the management network. At this point all VMs that were managed by the failed hosts are restarted in Site B. As in the single host failure, since none of the VMstores nor the witness were affected no action was taken.

Loss of Site A

Scenario: In this failure the entirety of Site A fails. Everything managed or running in Site A is down.

Expected: The VMs should be restarted in Site B per site affinity rules. VMstore migrates cluster IP to Site B.



Result/Discussion: The VMstore at Site B and the witness determine within 25s that Site A's VMstore is down. Site B VMstore is promoted to primary for the Site A's service groups and the replicated datastores become active. In this scenario due to vSphere HA settings and the automatic failover of the VMstore cluster IP, all decisions from vSphere HA are based off of all shares being available and there being no failure of storage detected. In Site B if there was no master one is elected. The master then determines which VMs are not running and restarts them because the hosts have active access the replicated VMs.

Summary

Using a well-known cluster solution but stretching the network and storage across a campus allows end-users a viable option for ensuring their sites and customers have constant access. As can be seen with the various failure scenarios presented in this document and with the many other potential failure combinations this solution can provide a high level of resiliency and automation.

At the core of the host side is VMware's vSphere HA with its robust and highly configurable cluster that allows VMs to easily move between physical hosts all the while having an 'affinity' to a home site of ESXi servers for best performance.

Tying the solution together is VMstore's Synchronous Replication with automatic failover. Synchronous replication is the key in this solution where the data needs to be in sync at both sites to allow for any type of failure and avoid down time. The automatic failover of the VMstore's synchronous replication completes the intelligent infrastructure of this solution.