

Proof Of Concept Guide for Virtual SAN

Cormac Hogan
VMware Technical Marketing
November 2013

Version 2.0

Virtual SAN (beta refresh) POC Guide

Table Of Contents

Introduction	4
Requirements.....	4
Workflow.....	5
Proof Of Concept Lab Setup	6
Hosts	6
Storage	6
Cluster	6
Datastores	6
Virtual Machines	6
Networking.....	6
Step 1 – Setup the VSAN Network	7
Step 2 – Enable VSAN on the Cluster	9
2.1 Verifying the VSAN datastore.....	13
2.2 Verifying the Storage Provider status	14
2.3 Troubleshooting Installation & Configuration Issues.....	15
Step 3. Scale-Out Storage.....	16
3.1 Create New Disk Group	17
Step 4 – Build VM Storage Policies.....	19
4.1 Deploy VMs	23
4.2 Modifying VM Storage Policies	26
Step 5 – vMotion & Storage vMotion	29
5.1 Storage vMotion from NFS to vsanDatastore	29
5.2 vMotion from host with local storage to host without local storage.....	31
Step 6 - vSphere HA & VSAN Interoperability	32
6.1 Check base-sles object layout	32
6.2 Enable HA on the cluster.....	33
6.3 Host Failure – No running VMs	35
6.4 Host Failure – Running VMs	37
Appendix A: Troubleshooting Virtual SAN	39
A1. Troubleshooting VSAN using the ESXCLI.....	39

Virtual SAN (beta refresh) POC Guide

A2. Troubleshooting VSAN using the RVC.....	40
A3. Troubleshooting VSAN using the VSAN Observer.....	41

Beta Refresh

Virtual SAN (beta refresh) POC Guide

Introduction

This document is the proof of concept (POC) guide for VMware Virtual SAN (VSAN), new scale-out hybrid storage from VMware. The document steps through a number of different test scenarios that a customer can use to evaluate the Virtual SAN (VSAN) product. By completing these tests, a customer will experience the full breath of capabilities that VSAN has to offer.

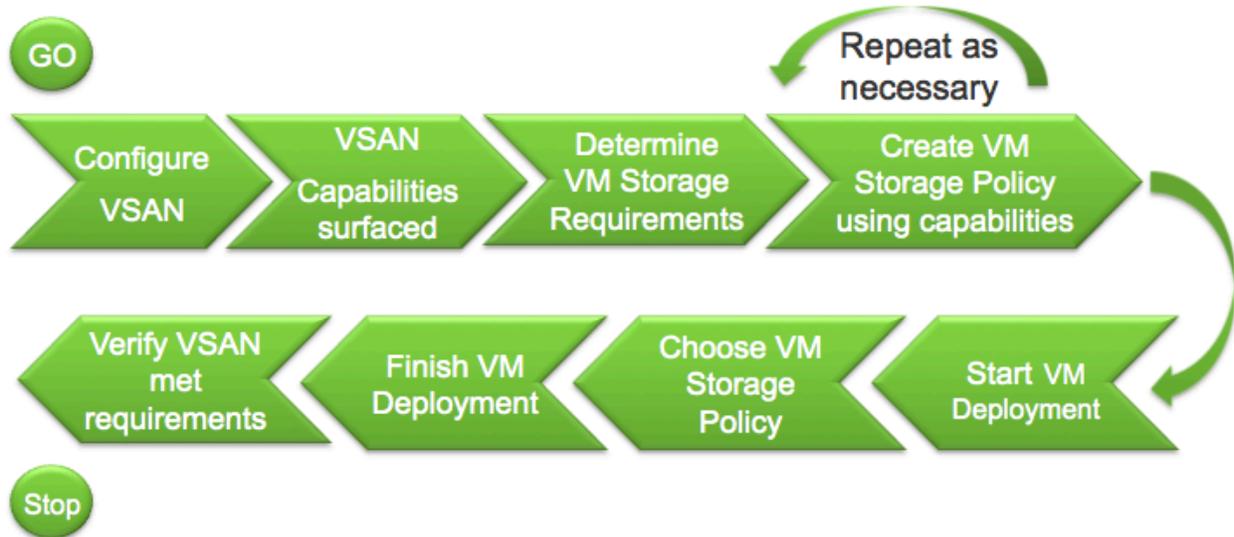
Requirements

Please refer to the pre-requisites checklist which is part of the Virtual SAN Proof Of Concept (POC) Kit. This contains the list of requirements for a successful POC.

Virtual SAN (beta refresh) POC Guide

Workflow

This guide will take customers through a number of workflows, such as configuring the VSAN cluster, the creation of various VM Storage Policies and the deployment of VM with those policies. The workflows can be summarized in the following chart.



Test cases are written in the following sequence:

1. Setup the VSAN network
2. Create the VSAN Cluster
3. Verify operation of VSAN
4. Identifying install & configure issues
5. Create VM Storage Policies
6. Deploy VM
7. Verify VM Storage Policy requirements met by VSAN
8. Check availability on errors
9. Experiment with troubleshooting tools

Virtual SAN (beta refresh) POC Guide

Proof Of Concept Lab Setup

The following is a description of the lab setup used for creating this guide. Detailed information about the setup is found in the POC pre-requisite checklist that is part of the POC Kit.

Hosts

The POC requires a total of 5 ESXi hosts, all running vSphere 5.5.

Storage

Four of the hosts contain empty local disks (both HDD & SSD) that can be consumed by VSAN. There is one empty SSD and two empty HDDs per host.

Cluster

Four of the hosts (esx-01a, esx-02a, esx-03a & esx-04a) are already in a cluster called Cluster Site A.

One of the hosts (esx-04a) does not contain any local storage.

A fifth host (esx-05a) is a stand-alone host that is not part of the cluster.

Datstores

In this guide, one NFS datastore, called **ds-site-a-nfs01**, is present and available to all hosts in the POC.

Virtual Machines

A single VM called **base-sles** resides on this NFS datastore. This is not necessary for a successful POC, but having a VM template available in your environment will speed up the POC. Optionally, you may have to deploy a Virtual Machine from an ISO image to complete the exercises.

Networking

In this setup, there is a distributed switch configured. A distributed switch is not necessary to implement a successful POC; standard virtual switches are fully supported in a VSAN configuration. However a distributed switch configuration allows us to use Network I/O Control to provide Quality Of Service (QoS) on the VSAN traffic.

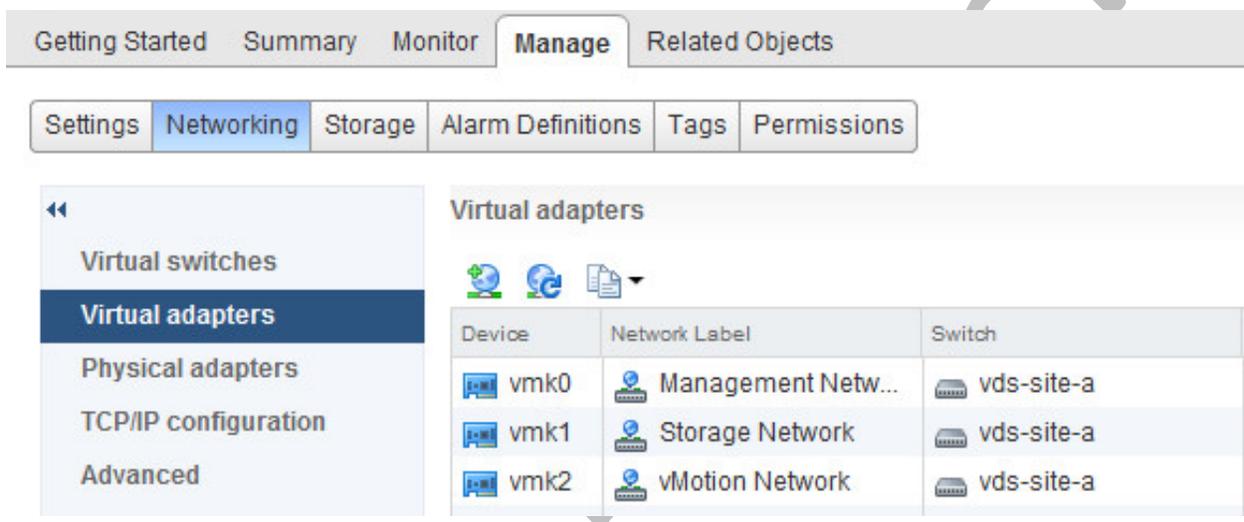
Virtual SAN (beta refresh) POC Guide

Step 1 – Setup the VSAN Network

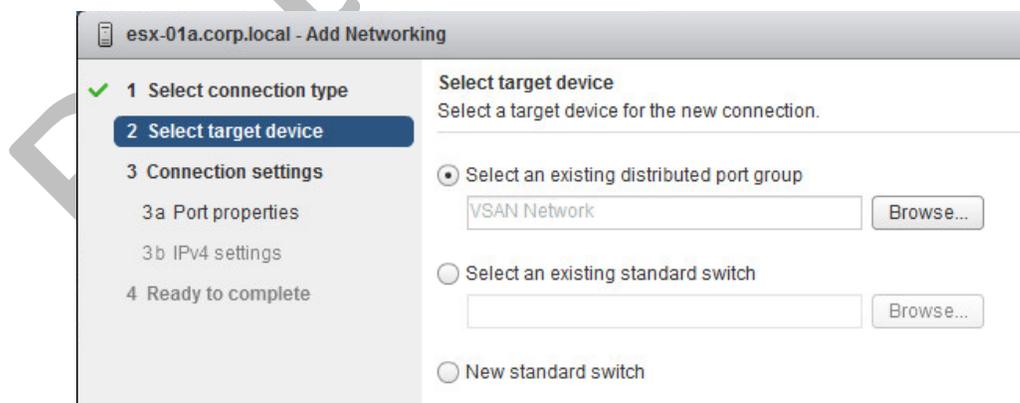
The first step is to set up communication between each of the ESXi hosts in the cluster. In this POC, there are already 3 VMkernel ports configured on each ESXi host. These are:

1. Management Network
2. vMotion Network
3. Storage Network

Navigate from the Home view to vCenter > Hosts & Clusters. Select the first of the four ESXi hosts and then select the Manage tab. Select **Networking** and then the **Virtual Adapters** view. It should look identical to the following:



We must now add the VSAN network. Setting up the VSAN network involves steps identical to setting up any other VMkernel network. Click on the icon to add a new virtual adapter (VMkernel Network Adapter), then select a distributed port group called VSAN network that has already been created:



Virtual SAN (beta refresh) POC Guide

At the point where you select the Port properties, select **Virtual SAN traffic**.

esx-01a.corp.local - Add Networking

- ✓ 1 Select connection type
- ✓ 2 Select target device
- 3 Connection settings
 - 3a Port properties**
 - 3b IPv4 settings
- 4 Ready to complete

Port properties
Specify VMkernel port settings.

VMkernel port settings

Network label: VSAN Network(vds-site-a)

IP settings: IPv4

TCP/IP stack: Default

Available services

Enable services:

- vMotion traffic
- Fault Tolerance logging
- Management traffic
- Virtual SAN traffic

All other settings may be left at the **default** (Obtain IPv4 settings automatically). The VSAN interfaces will be assigned their IP addresses via DHCP. This must be repeated for the four ESXi hosts in the cluster (esx-01a thru to esx-04a). We will deal with ESXi host esx-05a separately in part 6 of the lab.

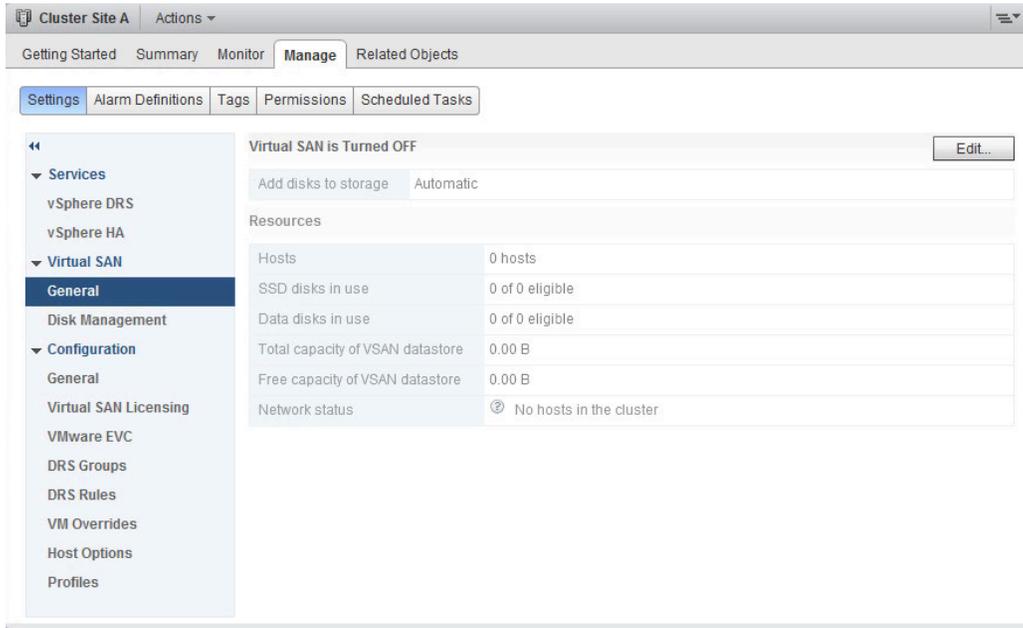
Key Takeaway: VSAN requires a network connection between hosts for communication and I/O purposes. VSAN requires multicast, but **only** in L2.

Virtual SAN (beta refresh) POC Guide

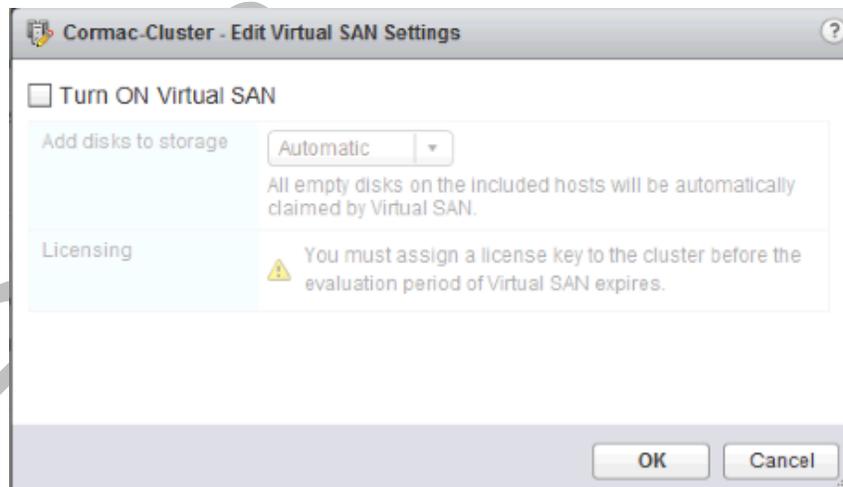
Step 2 – Enable VSAN on the Cluster

At this point, the networking is in place on all nodes. The four of your ESXi hosts are already in a cluster. We can now go ahead and enable the VSAN cluster, and point out how easy it is to setup.

Select the cluster object, then navigate to Manage tab > Settings > Virtual SAN > General:

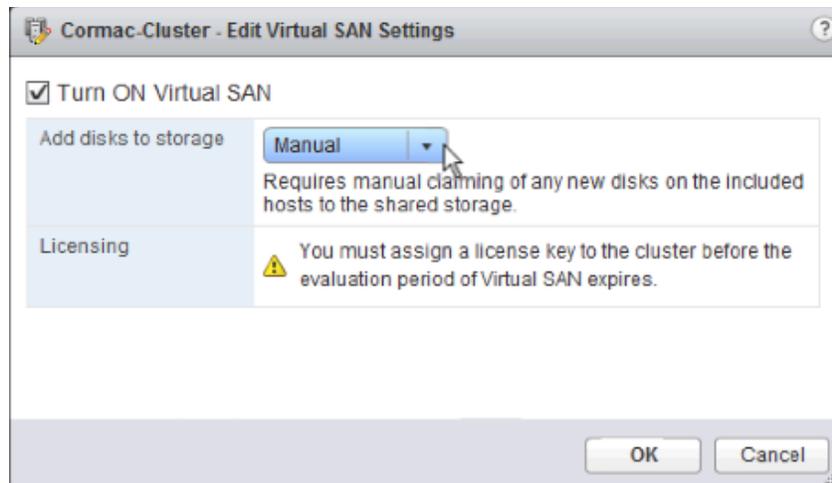


Next then click on the **Edit** button located to the right of the window – **Do not click OK until told to do so:**



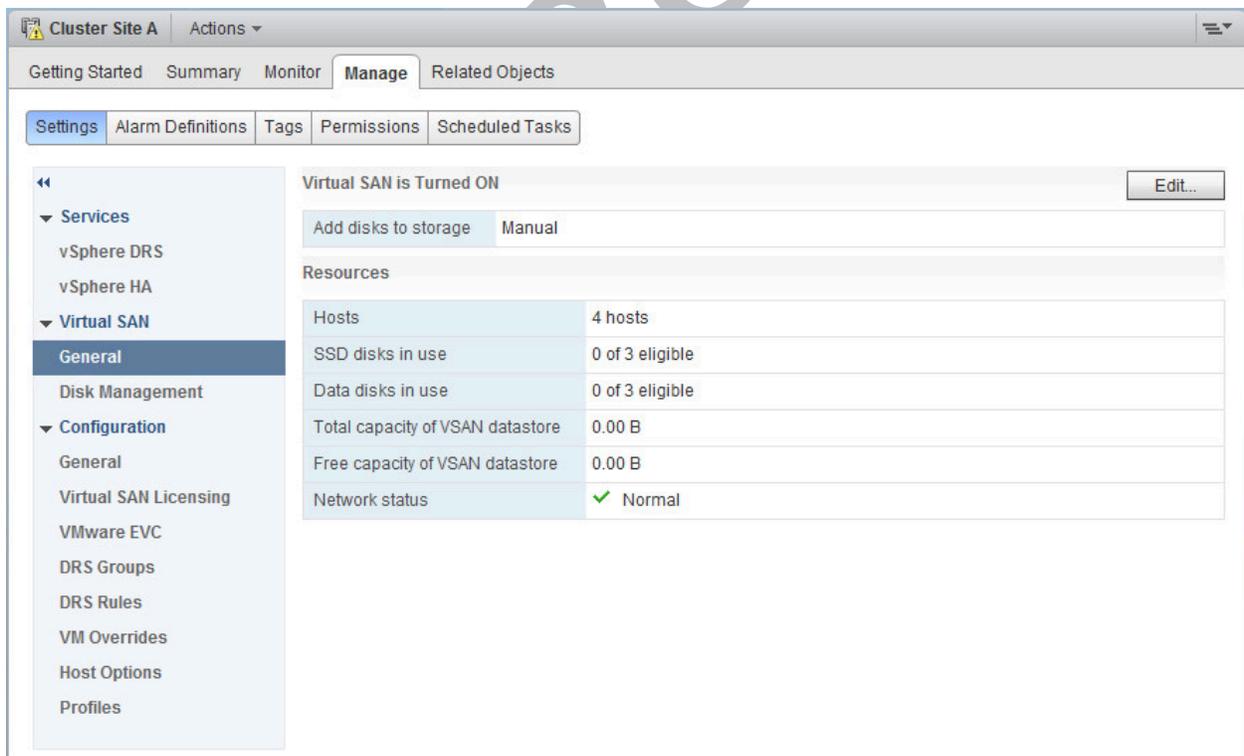
Virtual SAN (beta refresh) POC Guide

By default, VSAN will automatically **add new disks** to the VSAN cluster (Disk claim policy). Ensure that the setting 'Add disks to storage' is set to **manual**, as we manually add disks to the cluster.



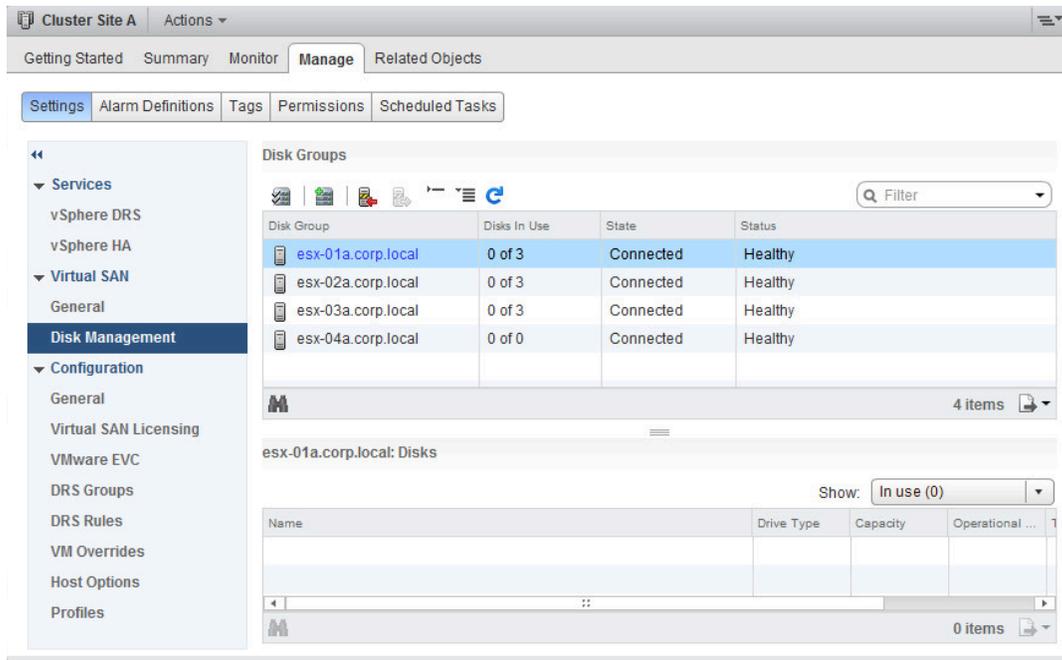
Remember that there is a requirement to have at least 1 SSD drive per host. In production environments, SSD will make up at least 10% of all storage, the remaining 90% being HDD of course. In this lab, there is a single SSD & a two HDD on three of the four ESXi hosts in the cluster. Click **OK**.

After clicking OK, you may have to refresh the vSphere web client for these changes to appear.

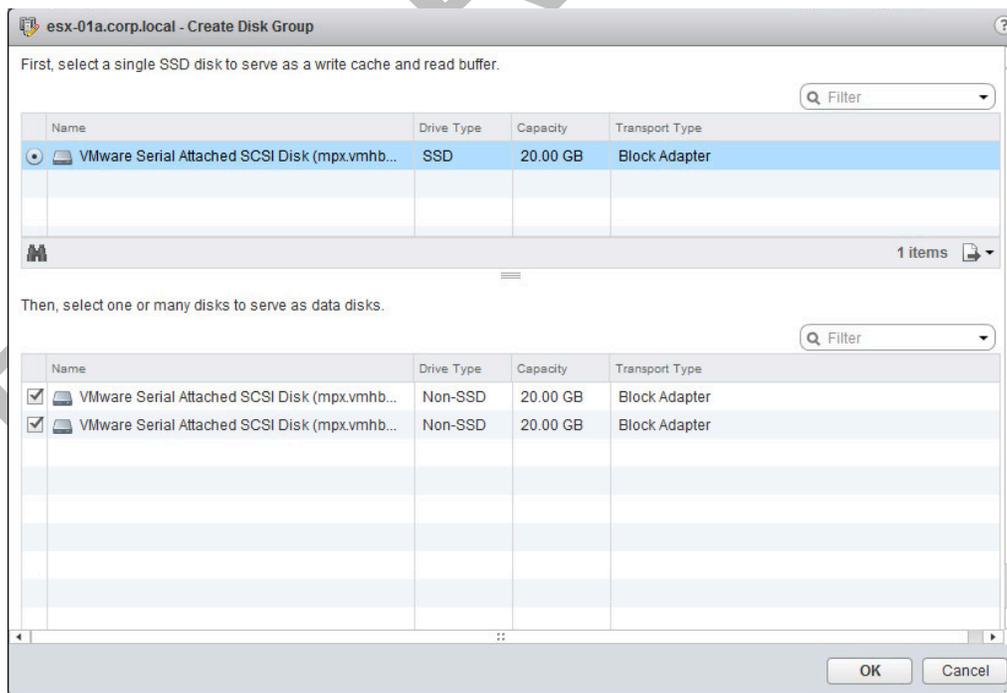


Virtual SAN (beta refresh) POC Guide

Go to the VSAN **Disk Management** view. Here, you will see that there are 4 hosts in the cluster, but no Disk Groups are created. This is because **Manual** mode was selected when the cluster was created, so you have to create the Disk Groups manually.



Select host esx-01a and click on the icon (with the green plus symbol) to create a new Disk Group.



Virtual SAN (beta refresh) POC Guide

Each disk group created may contain only one SSD. The SSD is used for write cache/read buffer and the HDDs are used for data disks/capacity. The SSD & both the HDDs on these hosts have been selected to be part of the VSAN (although a capacity ratio of 1:10 for SSD:HDD is what is deemed a best practice). Repeat this operation for all three hosts that have storage (esx-01a, esx-02a & esx-03a). Host esx-04a does not have any local storage. When completed, this cluster will have three disk groups.

Disk Group	Disks In Use	State	Status
esx-01a.corp.local	3 of 3	Connected	Healthy
Disk group (000...)	3		
esx-02a.corp.local	3 of 3	Connected	Healthy
Disk group (000...)	3		
esx-03a.corp.local	3 of 3	Connected	Healthy
Disk group (000...)	3		
esx-04a.corp.local	0 of 0	Connected	Healthy

Name	Drive Type	Capacity	Operational ...
VMware Serial Attached SCSI Disk (mpx.vmhba2:C0:T0:L0)	SSD	20.00 GB	Mounted
VMware Serial Attached SCSI Disk (mpx.vmhba2:C0:T1:L0)	Non-SSD	20.00 GB	Mounted

At this point, all four ESXi hosts in the cluster should be able to successfully see the VSAN datastore, labeled **vsanDatastore**.

Key Takeaway: The steps to setup VSAN are very simple, and are akin to the steps to setup vSphere HA and DRS clusters.

Virtual SAN (beta refresh) POC Guide

2.1 Verifying the VSAN datastore

At this point, we will verify that the VSAN datastore has been successfully created and that its capacity correctly reflects the total local storage capacity from each of the ESXi hosts. Once the network is created, physical storage is added to disk groups and the VSAN cluster created, a single VSAN datastore is built. Navigate to the Storage view and check the status of the vsanDatastore.

The screenshot displays the vSAN Datastore Summary view. At the top, there's a navigation bar with 'Getting Started', 'Summary', 'Monitor', 'Manage', and 'Related Objects'. Below this, a storage bar shows 'STORAGE' with a progress indicator, 'FREE: 119 GB', 'USED: 0 B', and 'CAPACITY: 119 GB'. A 'Refresh' button is located to the right. The main content area is divided into two panels: 'Details' and 'Tags'. The 'Details' panel contains the following information:

Details	
Location	ds://vmfs/volumes/vsan:5276cfaa875c4885-70b63b8ddc6aa40f/
Type	vsan
Hosts	4
Virtual machines	0

The 'Tags' panel is currently empty, displaying 'This list is empty.' and has 'Assign...' and 'Remove...' buttons at the bottom.

The capacity is an aggregate of the HDDs taken from each of the 3 ESXi hosts in the cluster. That is $3 \times 2 \times 20\text{GB} = 120\text{GB}$ (less some VSAN Datastore overheads). The 3 x SSD (1 on each of the 3 ESXi hosts with storage) are not considered when the capacity calculation is made.

Key Takeaway: VSAN uses the SSDs for read cache and write buffering. HDDs are used for capacity.

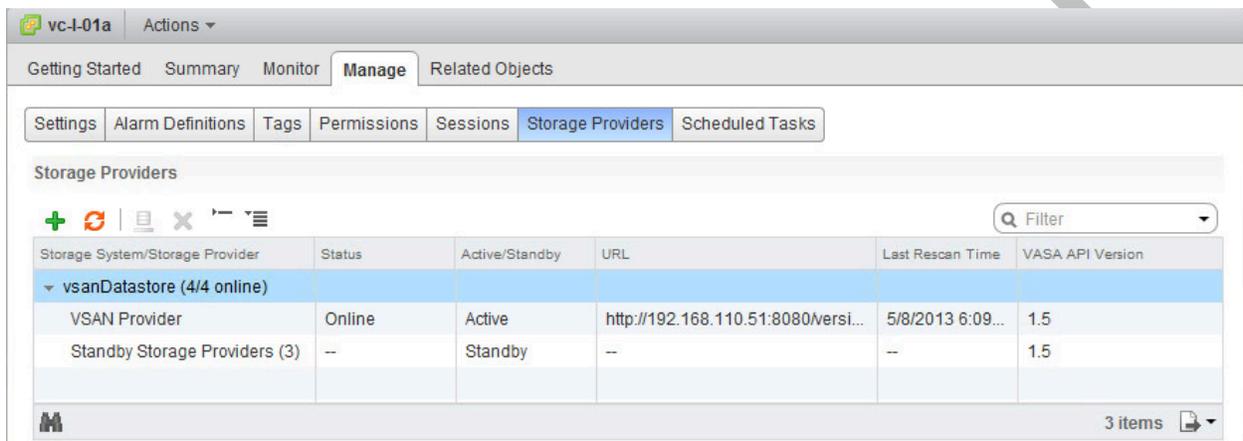
Virtual SAN (beta refresh) POC Guide

2.2 Verifying the Storage Provider status

To learn about the capabilities of VSAN, and to communicate between vCenter and the storage layer, a storage provider needs to be configured. Each ESXi host has a storage provider.

When a VSAN cluster is formed, the storage providers will be registered automatically with SMS, the Storage Management Service (by vCenter). However, it is best to verify that the storage providers on one of the ESXi hosts has successfully registered and is active, and that the other storage providers from the remaining ESXi hosts in the cluster are registered and are in standby mode.

Navigate to the vCenter server > Manage tab > Storage Providers to check the status.



In this four-node cluster, one of the VSAN providers is online and active, while the other three are in Standby. Each ESXi host participating in the VSAN cluster will have a provider, but only one needs to be active to provide VSAN datastore capability information.

Should the active provider fail for some reason, one of the standby storage providers will take over.

Key Takeaway: The storage provider, which surfaces up the VSAN capabilities to vCenter, is highly available.

Virtual SAN (beta refresh) POC Guide

2.3 Troubleshooting Installation & Configuration Issues

A number of issues may occur during the installation and configuration of Virtual SAN. This section will look at some of the common issues encountered.

- **Automatic mode does not claim disks**
 - **Issue:** Disks must be deemed local to be automatically claimed. If ESXi determines that a disk could be shared, it does not report it as "local". Many SAS controllers could (in theory) allow disks to be shared. VSAN currently does not support shared disks and reports these disks as not local.
 - **Solution:** Manually create disk groups with disks that are not "local".
- **Disk group creation fails silently**
 - **Issue:** You manually build a disk group, select SSD & HDDs, and although the operation reports success, the disk group is not created.
 - **Solution:** This is the behavior observed in beta when VSAN is not licensed. License VSAN.
- **Network status reports 'Misconfiguration detected'**
 - **Issue:** VSAN cluster fails to form due to hosts in the cluster being unable to communicate and the fact that the VSAN Traffic port group is configured on each of the hosts and the hosts can successfully ping each other.
 - **Solution:** Ensure ports on network switch allow multicast, ideally using a multicast router. Multicast is only used for relatively infrequent metadata operations, like object creation, change in object status (e.g. after a failure) as well as publication of stats like significant change of free disk space. The publication of stats is throttled to only update on significant changes and so is also infrequent. In steady state multicast traffic will be very low, but it may spike in events of network hiccups or disk failures. Those spikes, while higher than steady state, are still much lower than the TCP traffic from sustained storage IO.

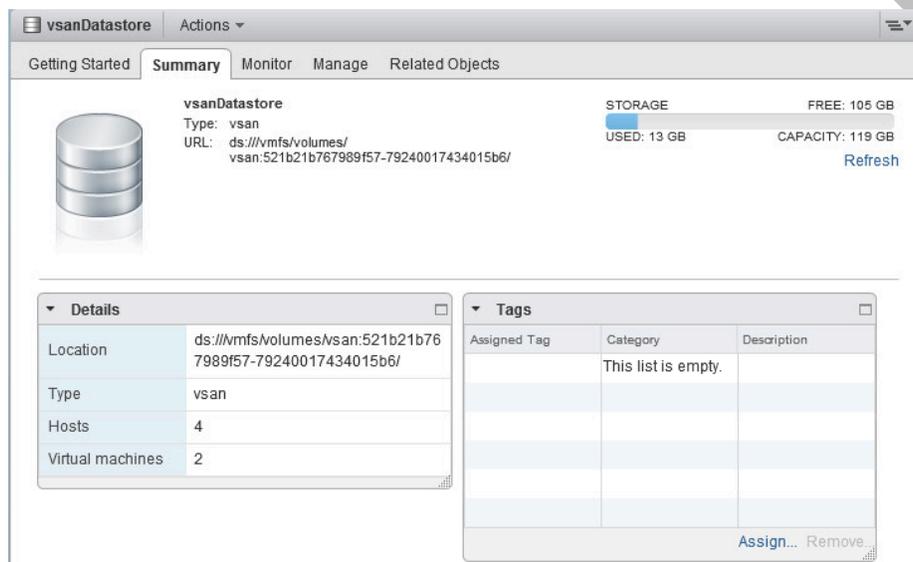
Virtual SAN (beta refresh) POC Guide

Step 3. Scale-Out Storage

In this section of the lab document, we are going to look at the ability to add another ESXi host with storage to the VSAN cluster and observe the scale-out capabilities of the product.

At this point, we have four ESXi hosts in the cluster, although only three are contributing local storage to the VSAN datastore. The VM, **base-sles**, should currently reside on a host that contributes local storage, esx-01a, esx-02a or esx-03a.

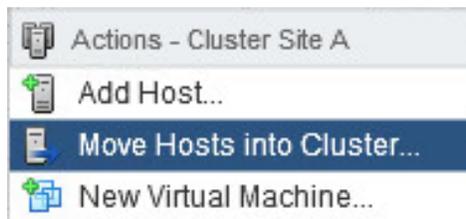
Let's check the status of the vsanDatastore. Navigate to the vsanDatastore summary tab, and it should look like this:



You have now reached the conclusion that you would like to add more compute and storage to the VSAN cluster, which involves adding a new ESXi host that contains additional disks.

There is a fifth ESXi host (esxi-05a) in your inventory that has not yet been added to the cluster. We will do that now and examine how the vsanDatastore seamlessly grows to include this new capacity.

Navigate to the cluster object in the inventory, right click and select the action 'Move hosts into cluster'.



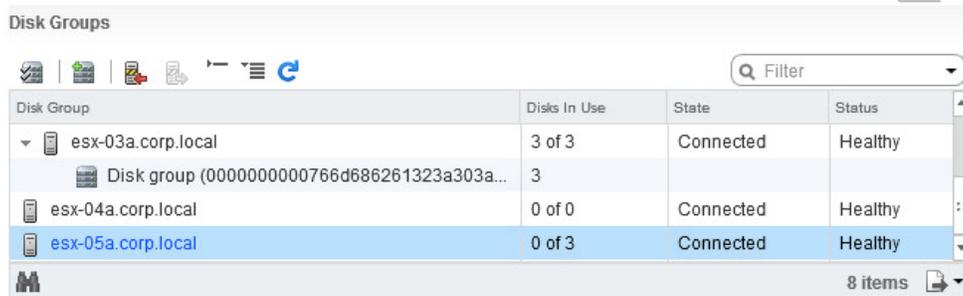
From the list of available hosts (you should only see esx-05a), select the host and click OK.

Virtual SAN (beta refresh) POC Guide

The next step is to add a VSAN network to this host. As per the procedure outlined in step 1 of this lab guide, create a VSAN VMkernel network adapter to this host using the distributed port group called VSAN Network. Make sure you select the Virtual SAN traffic service, and let DHCP provide the IP settings.

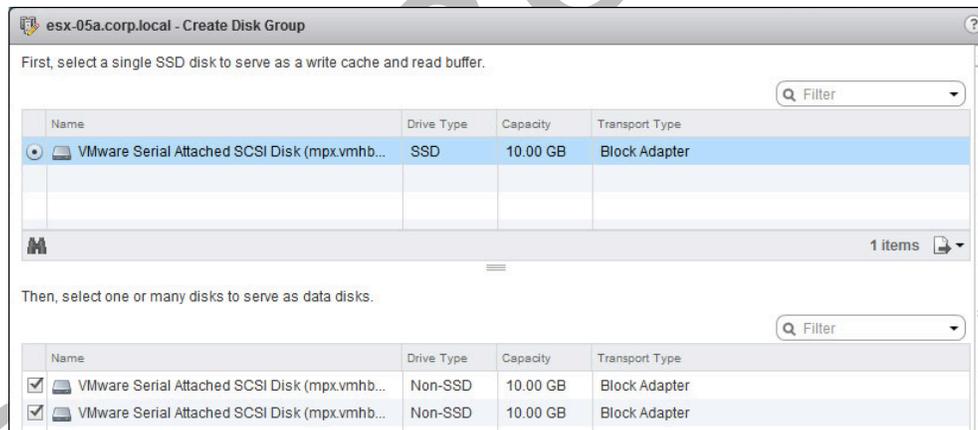
3.1 Create New Disk Group

When that is completed, since the cluster was setup in Manual mode, we need to create a new disk group using the disks on host esx-05a. Select Cluster > Manage > Settings > Disk Management. Select the host just added. You will notice that 0 of 3 disks are in use.



Disk Group	Disks In Use	State	Status
esx-03a.corp.local	3 of 3	Connected	Healthy
Disk group (0000000000766d686261323a303a...)	3		
esx-04a.corp.local	0 of 0	Connected	Healthy
esx-05a.corp.local	0 of 3	Connected	Healthy

Now create a new disk group, and add the disks (one SSD and two HDDs) to the disk group. These disks are all 10GB in size. Click OK and wait for the new disk group to be created.



esx-05a.corp.local - Create Disk Group

First, select a single SSD disk to serve as a write cache and read buffer.

Name	Drive Type	Capacity	Transport Type
VMware Serial Attached SCSI Disk (mpx.vmhb...)	SSD	10.00 GB	Block Adapter

Then, select one or many disks to serve as data disks.

Name	Drive Type	Capacity	Transport Type
<input checked="" type="checkbox"/> VMware Serial Attached SCSI Disk (mpx.vmhb...)	Non-SSD	10.00 GB	Block Adapter
<input checked="" type="checkbox"/> VMware Serial Attached SCSI Disk (mpx.vmhb...)	Non-SSD	10.00 GB	Block Adapter

When it is created, revisit the vsanDatastore summary view and check if the size has increased with the addition of the new host & disks. You should observe that the capacity of the datastore has seamlessly increased from 118GB to 138GB with the addition of two x 10GB HDDs. (remember that SSDs do not contribute towards capacity).

Virtual SAN (beta refresh) POC Guide

The screenshot displays the 'vsanDatastore' management interface. At the top, there are tabs for 'Getting Started', 'Summary', 'Monitor', 'Manage', and 'Related Objects'. The 'Summary' tab is active, showing a storage usage bar with 'USED: 13 GB' and 'FREE: 125 GB'. Below this, a 'Details' pane lists properties: Location (ds://mfs/volumes/vsan:521b21b767989f57-79240017434015b6/), Type (vsan), Hosts (5), and Virtual machines (2). A 'Tags' pane is also visible, which is currently empty.

Assigned Tag	Category	Description
This list is empty.		

As you can clearly see, increasing storage and compute in VSAN is relatively simple. Note that if VSAN were setup in automatic mode, the steps to create a disk group would not be necessary.

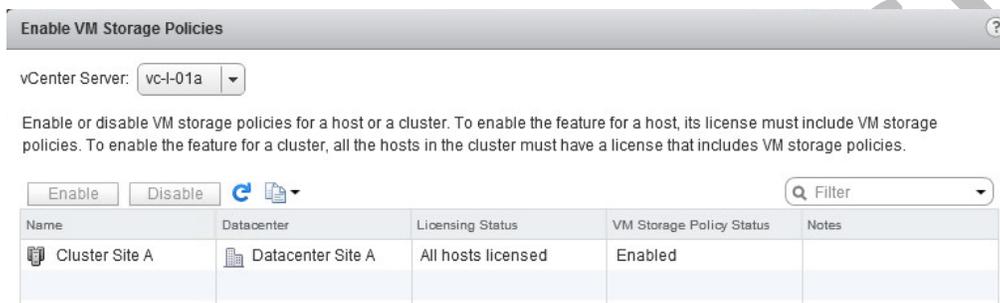
Key Takeaway: Scaling out storage and compute in a VSAN cluster is as simple as adding a new host to the cluster.

Virtual SAN (beta refresh) POC Guide

Step 4 – Build VM Storage Policies

Once the storage provider is added, **capabilities** made available from VSAN will be visible in the VM Storage Policies which can be found in Home > Rules & Profiles. VM Storage Policies are similar in some respects to the vSphere 5.0 & 5.1 Profile Driven Storage feature. There are two icons in this view representing ‘Create VM Storage Policies’ & ‘Enable VM Storage Policies’ respectively. The first step is to enable VM Storage Policies. *It is envisioned that Storage Policies will be automatically enabled on the VSAN cluster in the final product, but for now it must still be enabled manually.*

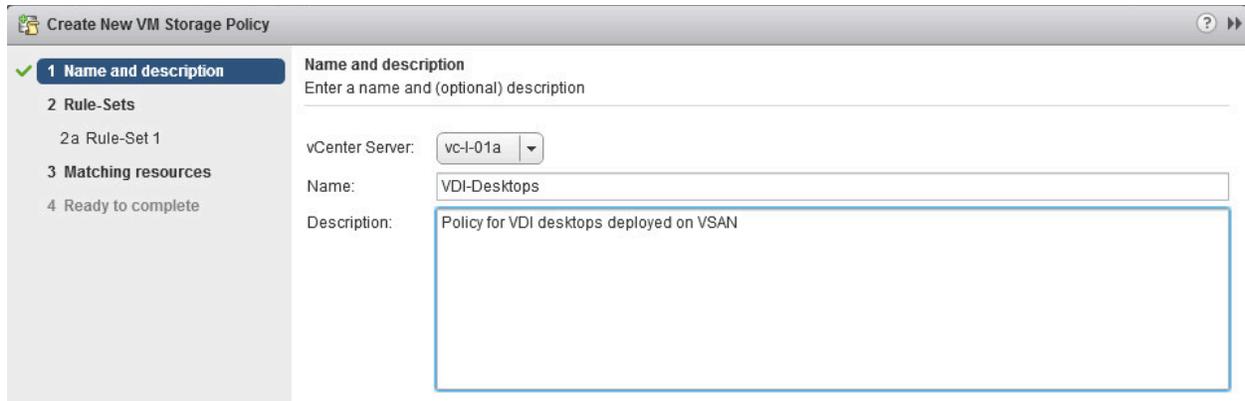
Click on the icon with the check mark to ‘Enable VM Storage Policies per compute resource’, and then click **Enable**:



Once enabled, you may now close the window. The capabilities of the **vsanDatastore** should now be visible during VM Storage Policy creation. By using a subset of the **capabilities**, a vSphere admin will be able to create a storage policy for their VM to guarantee Quality of Service (QoS). Click on the icon with the plus sign representing ‘Create New VM Storage Policy’ to begin.

Virtual SAN (beta refresh) POC Guide

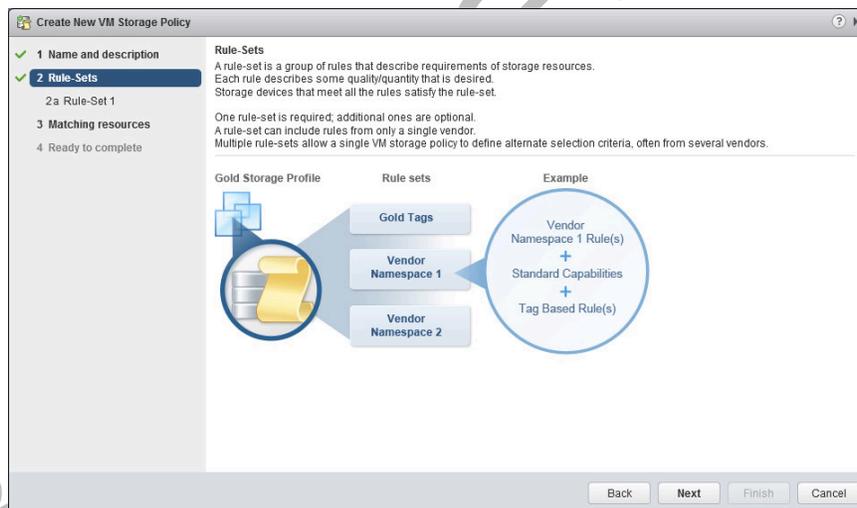
The first step is to give the VM Storage Policy a name. I will call it my **VDI-Desktops** profile for the purposes of this example:



The screenshot shows the 'Create New VM Storage Policy' wizard at the 'Name and description' step. The left sidebar shows the progress: 1 Name and description (checked), 2 Rule-Sets, 2a Rule-Set 1, 3 Matching resources, and 4 Ready to complete. The main area is titled 'Name and description' and contains the following fields:

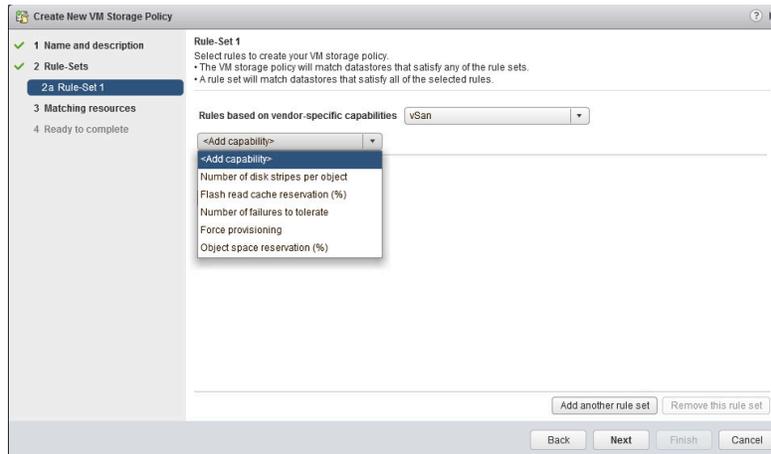
- vCenter Server: A dropdown menu with 'vc-l-01a' selected.
- Name: A text box containing 'VDI-Desktops'.
- Description: A text box containing 'Policy for VDI desktops deployed on VSAN'.

Next we get a description of rule sets. Rule-sets are a way of using storage from different vendors, e.g. for example you can have a single “bronze” policy which contains two separate rule-sets, one of which is a VSAN Rule-Set and the other may be a 3rd party storage vendor Rule-Set. When “bronze” is chosen as the VM Storage Policy at VM deployment time, both VSAN and the 3rd party storage are checked to see if they match the requirements in the policy.

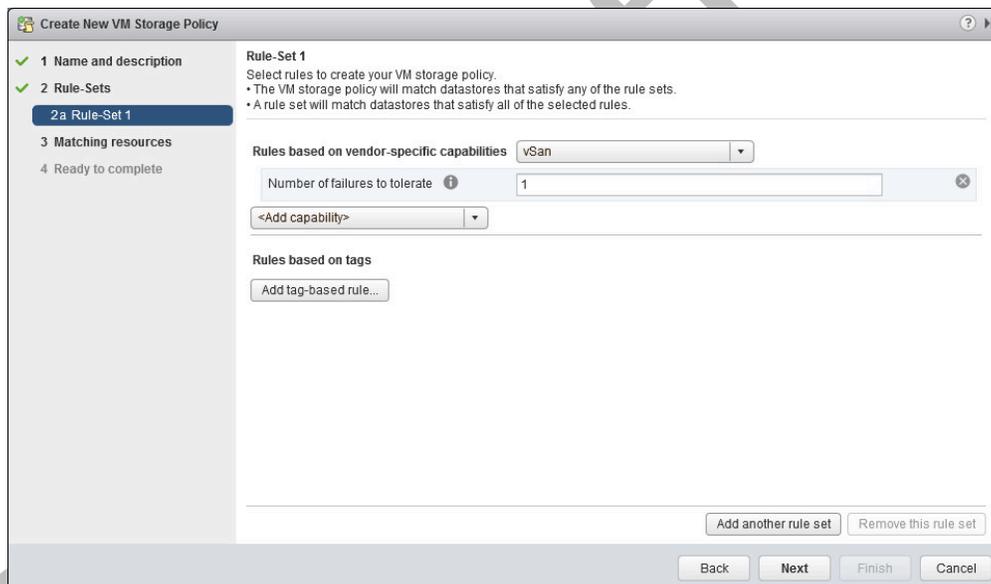


The next step is to select a subset of all of the **vsanDatastore** capabilities. Refer to official documentation for a full description of the capabilities. To begin you need to select the vendor, in this case it is called **vSan**.

Virtual SAN (beta refresh) POC Guide



The next step is to add the capabilities required for the virtual machines that you wish to deploy in your environment. In this particular example, I wish to specify an availability requirement. In this case, I want the VMs which have this policy associated with them to tolerate at least one failure (host, network or disk).



Virtual SAN (beta refresh) POC Guide

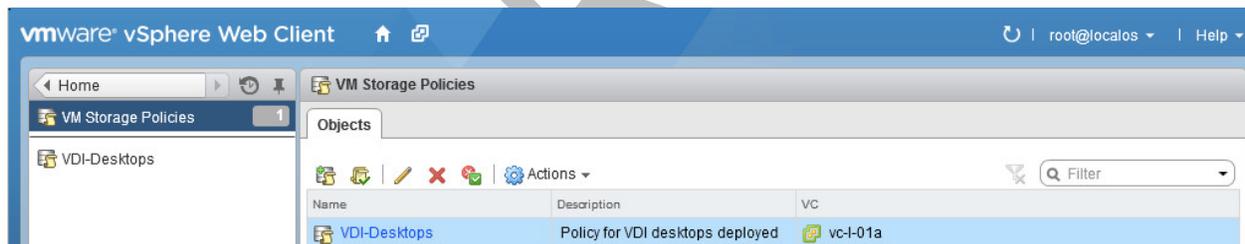
The nice thing about this is immediately I can tell whether or not any datastores are capable of understanding the requirement in the **matching resources** window. As you can see, my vsanDatastore is capable of understanding these requirements that I have placed in the VM Storage Policy:



Note that this is no guarantee that the datastore can meet the requirements in the VM Storage Policy. It simply means that the requirements in the VM Storage Policy can be understood by the datastores which show up in the matching resources.

This is where we start to define the requirements for our VMs and the applications running in the VMs. Now, we simply tell the storage layer what our requirements are by selecting the appropriate VM Storage Policy during VM deployment, and the storage layer takes care of deploying the VM in such a way that it meets those requirements.

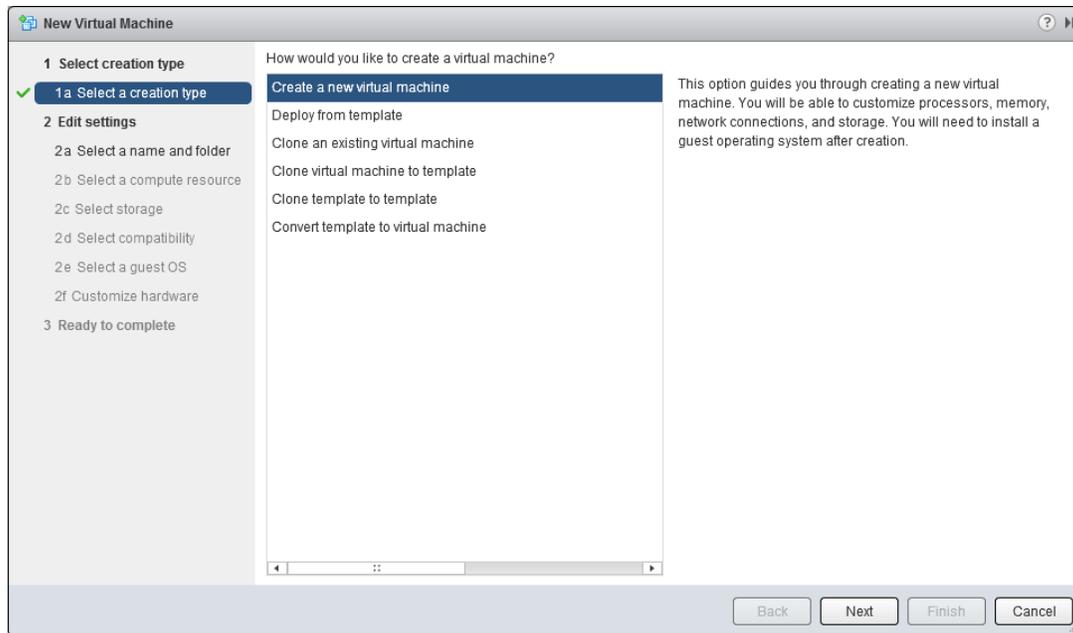
Complete the creation of the VM Storage Policy. This new 'policy' should now appear in the list of VM Storage Policies.



Virtual SAN (beta refresh) POC Guide

4.1 Deploy VMs

Create a virtual machine, which uses the **VDI-Desktops** profile created earlier.



Since DRS is not enabled, you will have to choose a host for this VM. Choose host esx-01a, esx-02a or esx-03a. Do not use esx-04a at this time. When it comes to selecting storage, you can now specify a VM Storage Policy (in this case **VDI-Desktops**). This will show that **vsanDatastore** is **Compatible** as a storage device, meaning once again that it understands the requirements placed in the storage policy. *It does not mean that the vsanDatastore will implicitly be able to accommodate the requirements – just that it understands them. This is an important point to understand about Virtual SAN.*

VM Storage Policy: **VDI-Desktops**

The following datastores are accessible from the destination resource that you selected. Select the destination datastore for the virtual machine configuration files and all of the virtual disks.

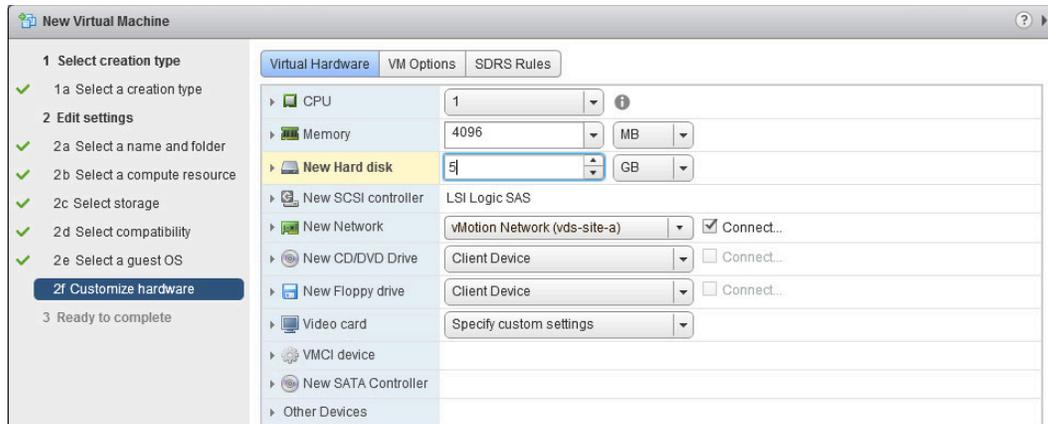
Name	Capacity	Provisioned	Free	Type	Storage DRS
Compatible					
 vsanDatastore (1)	138.00 GB	52.46 GB	130.41 GB	vsan	
Incompatible					
 ds-site-a-nfs01	14.27 GB	2.95 GB	12.74 GB	NFS	

Continue with the creation of this Virtual Machine, selecting the defaults for the remaining steps, including compatibility with ESXi 5.5 and later and Windows 2008 R2 (64-bit) as the Guest OS.

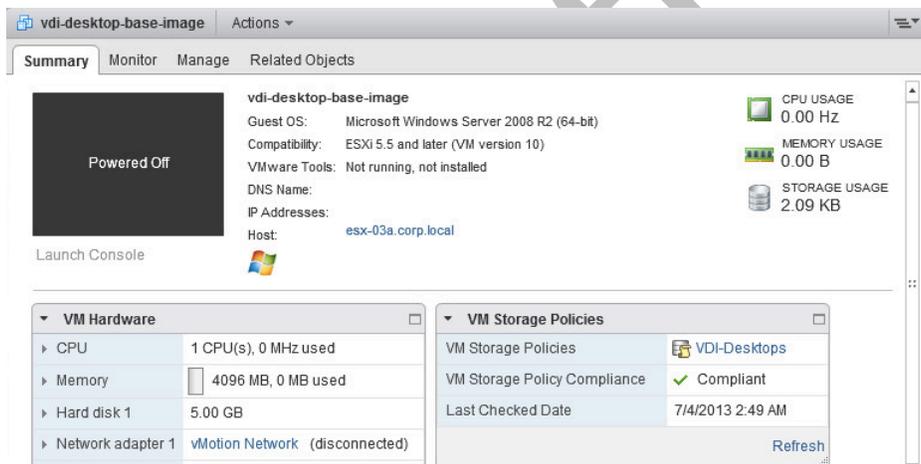
When you get to the **2f. Customize hardware** step, in the Virtual Hardware tab, expand the New Hard Disk virtual hardware and you will see VM Storage Policy set to **VDI-Desktops**. **Reduce the Hard Disk Size**

Virtual SAN (beta refresh) POC Guide

to 5GB in order for it to be replicated across hosts (the default size is 40GB – we want to reduce this as this is a small lab environment)

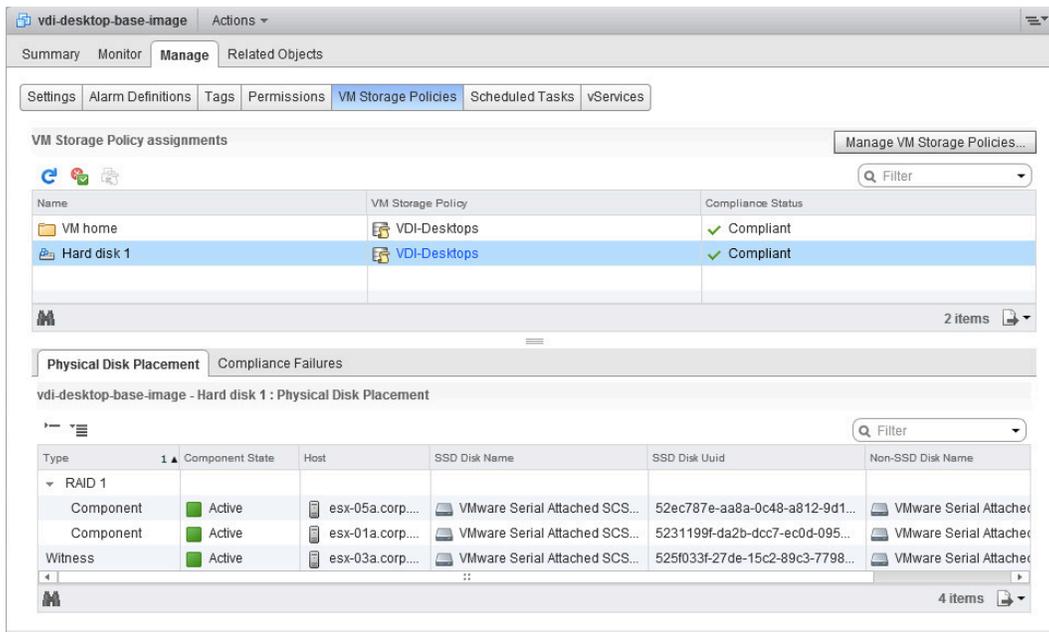


Complete the wizard. When the VM is created, look at its **Summary** tab and check the compliance state in the VM Storage Policies window. It should say Compliant with a green check mark.



As a final step, you might be interested in seeing how your virtual machine's objects have been placed on the **vsanDatastore**. To view the placement, select your Virtual Machine > Manage tab > VM Storage Policies. If you select one of your objects, the Physical Disk Placement will show you on which host the components of your objects reside, as shown in the example below.

Virtual SAN (beta refresh) POC Guide



The RAID 1 indicates that the VMDK has a replica. This is to tolerate a failure, the value that was set to 1 in the policy. So we can continue to run if there is a single failure in the cluster. The witness is there to act as a tiebreaker. If one host fails, and one component is lost, then this witness allows a quorum of storage objects to still reside in the cluster.

Notice that all three components are on different hosts for this exact reason. At this point, we have successfully deployed a virtual machine with a level of availability that can be used as the base image for our VDI desktops.

Examining the layout of the object above, we can see that a RAID1 configuration has been put in place by VSAN, placing each replica on different hosts. This means that in the event of a host, disk or network failure on one of the hosts, the virtual machine will still be available.

If the host on which the VM does not reside fails, then no action is required. If the host on which the VM resides fails, then vSphere HA can be used to automatically bring the VM online on one of the remaining hosts in the cluster.

We will examine this interoperability with vSphere HA in a later module.

Key Takeaway: Policies enable Software Driven Storage. For the first time, administrators can communicate their storage requirements on a per VM basis to the storage layer from vCenter.

Virtual SAN (beta refresh) POC Guide

4.2 Modifying VM Storage Policies

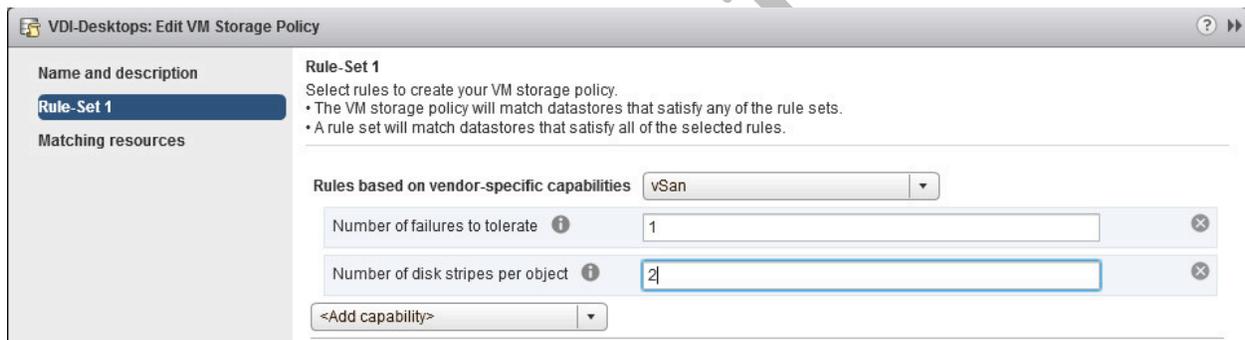
Scenario: customer notices that the VM deployed with the VDI-Desktop policy is getting a 90% read cache hit rate. This implies that 10% of reads need to be serviced from HDD. At peak time, this VM is doing 3000 IOPS. Therefore, there are 300 reads that need to be serviced from HDD. The specifications on the HDDs imply that each disk can do 150 IOPS, meaning that a single disk cannot service these additional 300 IOPS. To meet the I/O requirements of the VM implies that a stripe width of two disks should be implemented.

Edit Profile

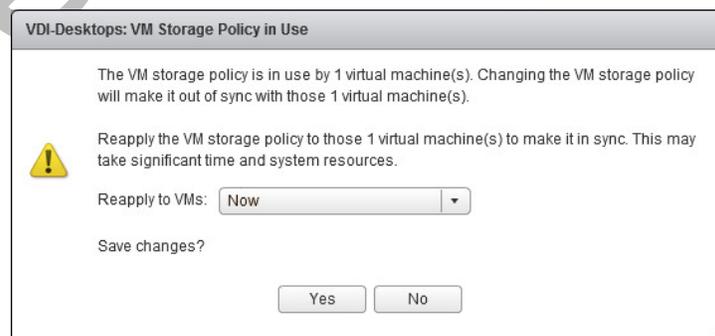
The first step is to edit the VDI-Desktops profile created earlier and add a stripe width requirement to the policy. Navigate back to Rules & Profiles, select VM Storage Policy, select the VDI-Desktop policy and click on 'edit'.

Add Stripe Width Capability

In the Rule-Set1, add a new capability called 'Number of disk stripes per object' and set the value to 2. This is the number of disks that the stripe will span.



Click OK. You will observe a popup which states that the policy is already in use. We will need to synchronize the virtual machine with the policy after saving the changes. Change the 'Reapply to VMs' to Now and Click Yes.



Virtual SAN (beta refresh) POC Guide

Resync Virtual Machine with Policy Changes

Staying on the VDI-Desktops policy, click on the Monitor tab. In the VMs & Virtual Disks view, you will see that the Compliance Status is 'Out of Date'.

The screenshot shows the 'VDI-Desktops' interface with the 'Monitor' tab selected. Under 'VMs and Virtual Disks', a table lists resources with their compliance status. The 'VM home' folder and 'Hard disk 1' are both marked as 'Out of Date' with a red 'X' icon. The last checked time for both is '5/9/2013 3:18 AM'. A toolbar above the table includes a reapply policy icon (3rd from left).

Name	Compliance Status	Last Checked
vd-desktop-base-image		
VM home	Out of Date	5/9/2013 3:18 AM
Hard disk 1	Out of Date	5/9/2013 3:18 AM

Click on the reapply policy icon (3rd from left) to reapply policy to all out of date entities. Answer yes to the popup. The compliance state should now change once the updated policy is applied.

The screenshot shows the same 'VDI-Desktops' interface after the policy has been reapplied. The 'VM home' folder and 'Hard disk 1' are now marked as 'Compliant' with a green checkmark icon. The last checked time for both is '5/9/2013 7:54 AM'.

Name	Compliance Status	Last Checked
vd-desktop-base-image		
VM home	Compliant	5/9/2013 7:54 AM
Hard disk 1	Compliant	5/9/2013 7:54 AM

As a final step, we will now re-examine the layout of the storage object to see if the request to create a stripe width of 2 has been implemented. Return to the Virtual Machine view > Manage > VM Storage Policy & select the Hard Disk 1 object:

Virtual SAN (beta refresh) POC Guide

Physical Disk Placement Compliance Failures

vdi-desktop-base-image - Hard disk 1 : Physical Disk Placement

Filter

Type	Component State	Host	SSD Disk Name	SSD Disk Uuid
RAID 1				
RAID 0				
Component	Active	esx-03a.corp...	VMware Serial Attached SCS...	525f033f-27de-15c2-89c3-7
Component	Active	esx-03a.corp...	VMware Serial Attached SCS...	525f033f-27de-15c2-89c3-7
RAID 0				
Component	Active	esx-05a.corp...	VMware Serial Attached SCS...	52ec787e-aa8a-0c48-a812
Component	Active	esx-02a.corp...	VMware Serial Attached SCS...	522c4ada-f695-0e56-1e5c-
Witness	Active	esx-05a.corp...	VMware Serial Attached SCS...	52ec787e-aa8a-0c48-a812
Witness	Active	esx-02a.corp...	VMware Serial Attached SCS...	522c4ada-f695-0e56-1e5c-
Witness	Active	esx-03a.corp...	VMware Serial Attached SCS...	525f033f-27de-15c2-89c3-7

10 items

Now we can see that the disk layout has changed significantly. Because we have requested a stripe width of two, the components that make up the stripe are placed in a RAID-0 configuration. Since we still have our failures to tolerate requirement, these RAID-0s must be mirrored by a RAID-1. Now that we have multiple components distributed across the 3 hosts, additional witnesses are needed in case of a host failure.

We are not going to install a Guest OS in this virtual machine. Instead we will focus our attention on another small virtual machine available in your environment for the remaining tests.

Key Takeaway: As virtual machine storage requirements change, administrators of VSAN can simply update the policy. Compare this to a physical SAN or NAS infrastructure where a new datastore would have to be provisioned to satisfy changing virtual machine I/O requirements.

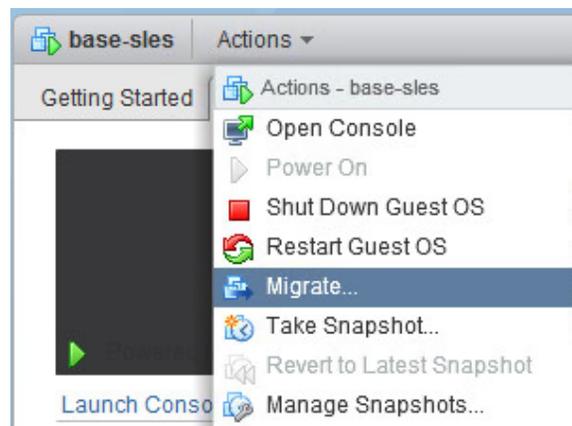
Virtual SAN (beta refresh) POC Guide

Step 5 – vMotion & Storage vMotion

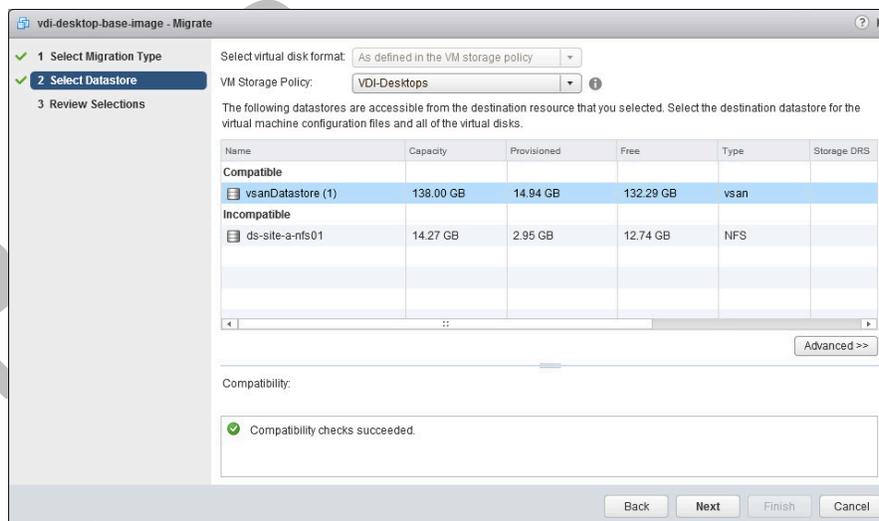
In this step, we will examine the interoperability of VSAN with core vSphere features such as vMotion & Storage vMotion. Power on the virtual machine called **base-sles** which resides on host esx-01a. This is a very small virtual machine, but will be sufficient for the purposes of this lab. Wait until the VMware Tools show as running before continuing. This should only take a moment or two.

5.1 Storage vMotion from NFS to vsanDatastore

This VM currently resides on an NFS datastore called **ds-site-a-nfs01**. We will migrate this virtual machine to the **vsanDatastore**. With the virtual machine **base-sles** selected in the inventory, from the actions list, choose the option to **migrate**.



Choose the option to **'Change datastore'**. At the Select Datastore window, change the VM Storage Profile to **VDI-Desktops**. This will show the vsanDatastore as **Compatible**.

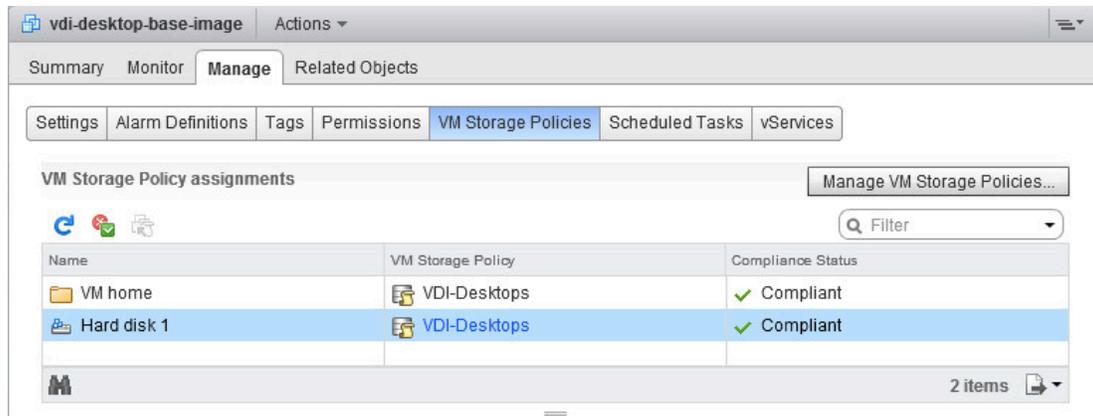


Finish the migration process and wait for the VM to migrate (a few minutes). This demonstrates that you can migrate from traditional datastore formats such as NFS & VMFS to the new vsanDatastore format.

Virtual SAN (beta refresh) POC Guide

Once the VM has been successfully migrated to the vsanDatastore, examine the layout of the VM. It should have the same layout as the vdi-desktop VM that we created earlier, i.e. a mirror and stripe configuration.

Navigate to the Virtual Machine view > Manage > VM Storage Policies. On first observation, you may see the VM home and Hard disk 1 objects state that the compliance state is **not applicable**. Simply click on the check compliance state icon (middle icon), and this should make them compliant.



Now select the Hard Disk 1 object and have a look at the physical disk placement:

Type	Component State	Host	SSD Disk Uuid	Non-SSD Disk Uuid
RAID 1				
RAID 0				
Component	Active	esx-02a.corp.local	52f459ff-ed1b-4d73-382d-c5e8...	52aa3248-0af2-5ee6-e580-1419...
Component	Active	esx-01a.corp.local	527c17b4-4474-bc78-320a-a7c...	52cb9451-092d-1ff6-106c-3ed7a...
RAID 0				
Component	Active	esx-03a.corp.local	52155f2a-2170-ae2f-9861-b94b...	528ef519-7218-22f3-1e3c-c751c...
Component	Active	esx-03a.corp.local	52155f2a-2170-ae2f-9861-b94b...	524dc479-50e7-86f8-6659-4693...
Witness	Active	esx-02a.corp.local	52f459ff-ed1b-4d73-382d-c5e8...	5226ce36-4f6e-530d-d60e-4277...
Witness	Active	esx-01a.corp.local	527c17b4-4474-bc78-320a-a7c...	52016b44-c006-3d04-0385-fa5a...
Witness	Active	esx-02a.corp.local	52f459ff-ed1b-4d73-382d-c5e8...	52aa3248-0af2-5ee6-e580-1419...

You should be able to see the stripe width of 2 (RAID-0) and the replication/mirror of each (RAID-1). And of course, we also have the tie-breaker witness disks.

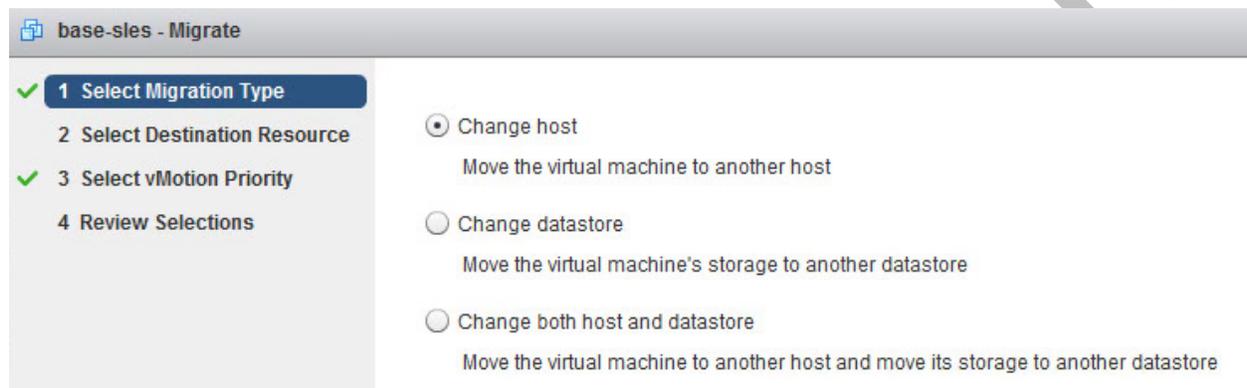
Virtual SAN (beta refresh) POC Guide

5.2 vMotion from host with local storage to host without local storage

Now we will show how hosts which are in the VSAN cluster, but do not have any local storage, can still use the vsanDatastore to run VMs.

At this point, the virtual machine **base-sles** resides on the vsanDatastore. The VM is currently on a host that contributes local storage to the vsanDatastore (esx-01a.corp.local). We will now move this to a host (esx-04a.corp.local) that does not have any local storage.

Once again select the **base-sles** virtual machine from the inventory. From the Actions drop down menu, once again select **Migrate**. This time we choose the option to 'Change host'.



At step 3 where a host selection needs to be made, select host esx04a.corp.local. Complete the migration by selecting all of the defaults from the remaining steps in the migration wizard.

When the migration has completed, you will see how hosts that do not contribute any local storage to the vsanDatastore can still run virtual machines. This means that VSAN can be scaled out on a compute basis.

To complete this section, migrate the VM back to a host that has local storage making up the VSAN datastore, e.g. esx-01a, esx-02a or esx-03a. Leave the VM residing on the vsanDatastore.

Key Takeaway: VMs can be migrated between a vsanDatastore, traditional VMFS and NFS datastores.

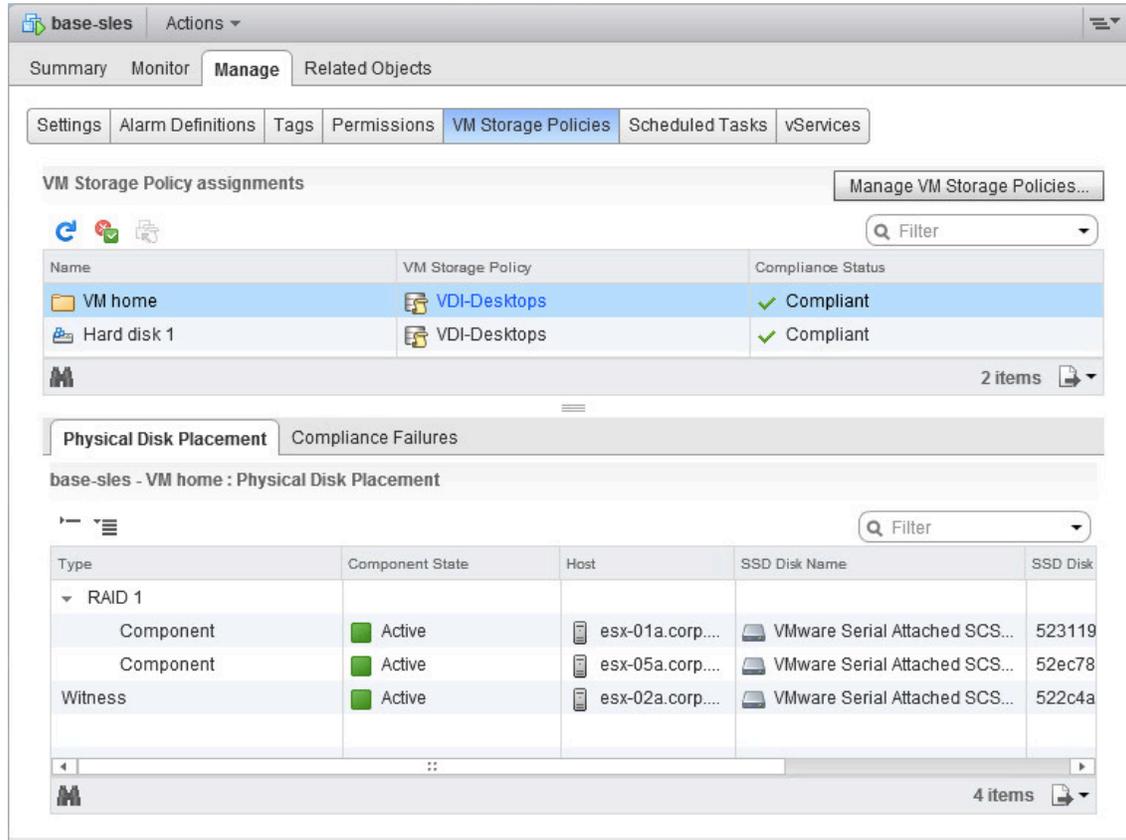
Virtual SAN (beta refresh) POC Guide

Step 6 - vSphere HA & VSAN Interoperability

This final section will provide details on how to evaluate VSAN with vSphere HA.

6.1 Check base-sles object layout

First, let's examine the object layout of the virtual machines. Let's first look at VM Home:



This storage object has 3 components, two of which are replicas making up a RAID-1 mirror. The third is a witness disk that is used for tie-breaking.

The next object is the disk, which we have looked at a number of times already. Just to recap, this has a StripeWidth set to 2, therefore there is a RAID-0 stripe component across two disks. There is no magic here – to mirror an object with a striped width of 2, 4 disks are required. Again, since ComponentFailuresToTolerate is set to 1, there is also a RAID-1 configuration to replicate the stripe. So we have two RAID-0 (stripe) configurations, and a RAID-1 to mirror the stripes. The witnesses are once again used for tie-breaking functionality in the event of failures.

Virtual SAN (beta refresh) POC Guide

Physical Disk Placement | Compliance Failures

base-sles - Hard disk 1 : Physical Disk Placement

Filter

Type	Component State	Host	SSD Disk Uuid	Non-SSD Disk Uuid
RAID 1				
RAID 0				
Component	Active	esx-02a.corp.local	52f459ff-ed1b-4d73-382d-c5e8...	52aa3248-0af2-5ee6-e580-1419...
Component	Active	esx-01a.corp.local	527c17b4-4474-bc78-320a-a7c...	52cb9451-092d-1ff6-106c-3ed7a...
RAID 0				
Component	Active	esx-03a.corp.local	52155f2a-2170-ae2f-9861-b94b...	528ef519-7218-22f3-1e3c-c751c...
Component	Active	esx-03a.corp.local	52155f2a-2170-ae2f-9861-b94b...	524dc479-50e7-86f8-6659-4693...
Witness	Active	esx-02a.corp.local	52f459ff-ed1b-4d73-382d-c5e8...	5226ce36-4f6e-530d-d60e-4277...
Witness	Active	esx-01a.corp.local	527c17b4-4474-bc78-320a-a7c...	52016b44-c006-3d04-0385-fa5a...
Witness	Active	esx-02a.corp.local	52f459ff-ed1b-4d73-382d-c5e8...	52aa3248-0af2-5ee6-e580-1419...

10 items

The next step is to invoke some failures in the cluster to see how this impacts the components that make up our virtual machine storage objects, but also how VSAN & vSphere HA interoperate to enable availability.

6.2 Enable HA on the cluster

Navigate to the Cluster and select the Manage tab > Settings. Select the vSphere HA service. vSphere HA is currently Turned OFF:

Cluster Site A | Actions

Getting Started | Summary | Monitor | **Manage** | Related Objects

Settings | Alarm Definitions | Tags | Permissions | Scheduled Tasks

Services

- vSphere DRS
- vSphere HA**
- Virtual SAN
- General
- Disk Management

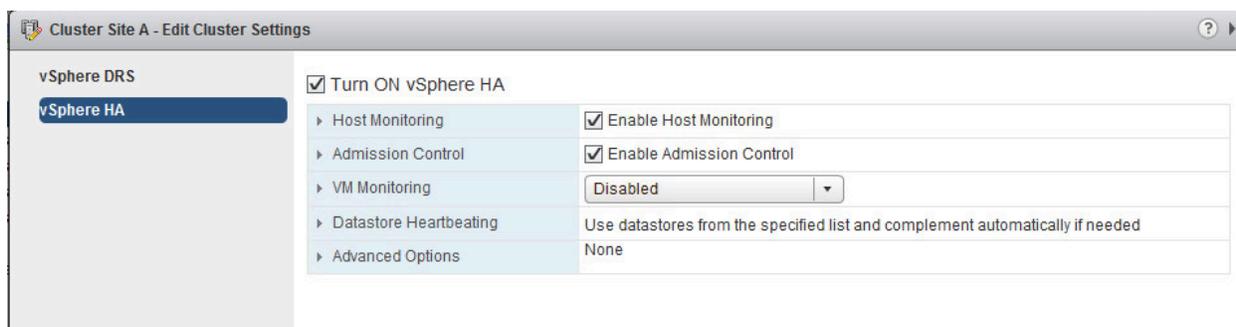
vSphere HA is Turned OFF [Edit...]

Runtime information for vSphere HA is reported under vSphere HA Monitoring

Host Monitoring	Enabled
Admission Control	Enabled
VM Monitoring	Disabled
Datastore Heartbeating	Use datastores from the specified list and complement automatically if needed
Advanced Options	

Click on the Edit button, and click on the checkbox to Turn ON vSphere HA.

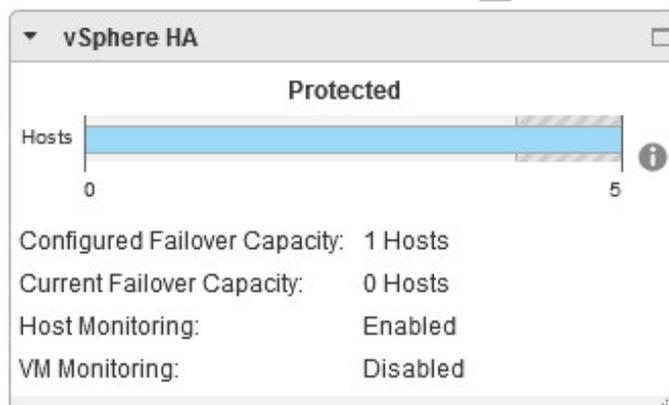
Virtual SAN (beta refresh) POC Guide



By default, the vSphere HA Admission Controls have been set to tolerate a single host failure. You can examine this if you wish by breaking open the Admission Control settings to verify. When satisfied, click on the OK button to enable HA.

After enabling HA, you will see a warning about insufficient resources to satisfy vSphere HA failover level. This is a transient warning and will eventually go away after a few moments, once the HA cluster has finished configuring. You can try refreshing from time to time to remove it.

The cluster summary tab should show a vSphere HA overview as follows:

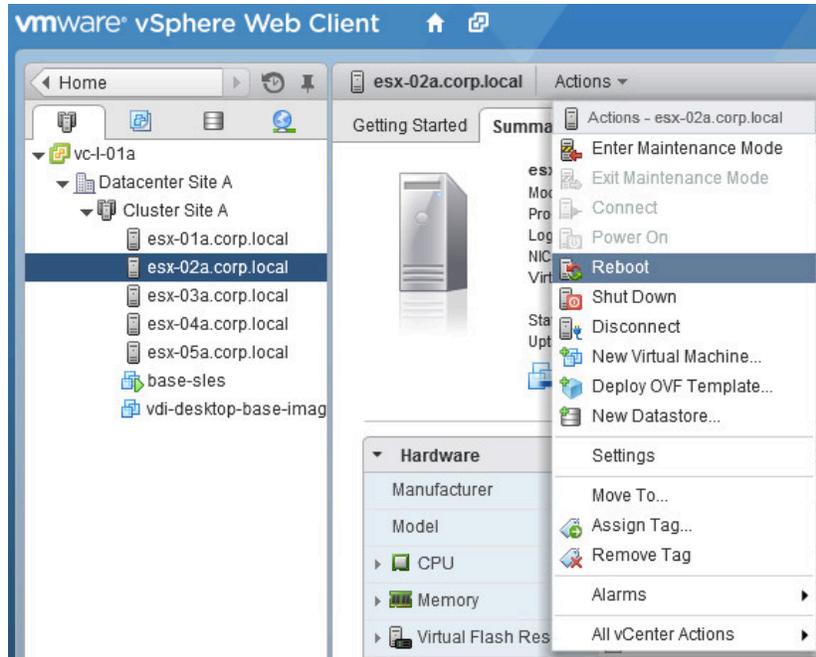


Virtual SAN (beta refresh) POC Guide

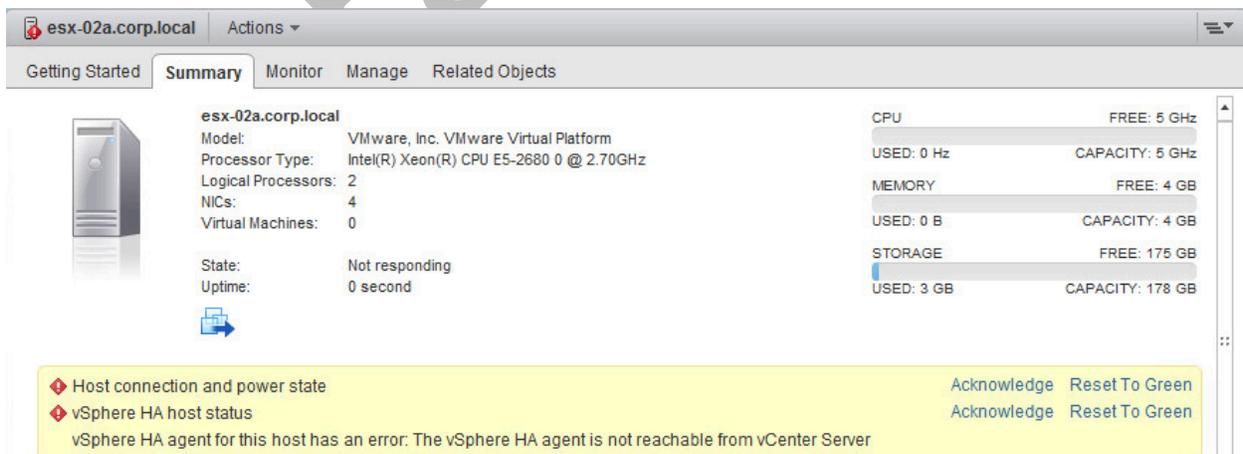
6.3 Host Failure – No running VMs

In this first failure scenario, we will take one of the hosts out of the cluster. This host does not have any running VMs, but we will use it to examine how the VSAN replicas provide continuous availability for the VM, and how the Admission Control setting in vSphere HA and the ComponentFailuresToTolerate are met.

In this step, host esx-02a is rebooted. Select the **Reboot** option from the ESXi host actions:



In a short time, we see warnings and errors related to the fact that vCenter can no longer reach the HA Agent and then we see errors related to host connection and power status.



If we check on other hosts in the cluster, we see VSAN communication issues.

Virtual SAN (beta refresh) POC Guide

esx-01a.corp.local

Model: VMware, Inc. VMware Virtual Platform
 Processor Type: Intel(R) Xeon(R) CPU E5-2680 0 @ 2.70GHz
 Logical Processors: 2
 NICs: 4
 Virtual Machines: 1

State: Connected
 Uptime: 21 hours

CPU: FREE: 4 GHz
 USED: 1 GHz CAPACITY: 5 GHz

MEMORY: FREE: 3 GB
 USED: 1 GB CAPACITY: 4 GB

STORAGE: FREE: 137 GB
 USED: 2 GB CAPACITY: 139 GB

Host cannot communicate with all other nodes in the VSAN enabled cluster

With one host out of the cluster, and object components that were held on that host are displayed as **Absent** – Object not found. First we will look at the **VM home**:

Physical Disk Placement Compliance Failures

base-sles - VM home : Physical Disk Placement

Type	Component State	Host	SSD Disk Uuid	Non-SSD Disk Uuid
Witness	Active	esx-01a.corp.local	527c17b4-4474-bc78-320a-a7c...	52016b44-c006-3d04-0385-fa5a...
RAID 1				
Component	Active	esx-03a.corp.local	52155f2a-2170-ae2f-9861-b94b...	528ef519-7218-22f3-1e3c-c751c...
Component	Absent	Object not found	Object not found	Object not found

Next we can look at the **Hard Disk**:

Physical Disk Placement Compliance Failures

base-sles - Hard disk 1 : Physical Disk Placement

Type	Component State	Host	SSD Disk Uuid	Non-SSD Disk Uuid
Component	Absent	Object not found	Object not found	Object not found
Witness	Active	esx-01a.corp.local	527c17b4-4474-bc78-320a-a7c...	52016b44-c006-3d04-0385-fa5a...
Component	Absent	Object not found	Object not found	Object not found
RAID 1				
RAID 0				
Component	Absent	Object not found	Object not found	Object not found
Component	Active	esx-01a.corp.local	527c17b4-4474-bc78-320a-a7c...	52cb9451-092d-1ff6-106c-3ed7a...
RAID 0				
Component	Active	esx-03a.corp.local	52155f2a-2170-ae2f-9861-b94b...	528ef519-7218-22f3-1e3c-c751c...
Component	Active	esx-03a.corp.local	52155f2a-2170-ae2f-9861-b94b...	524dc479-50e7-86f8-6659-4693...

10 items

Virtual SAN (beta refresh) POC Guide

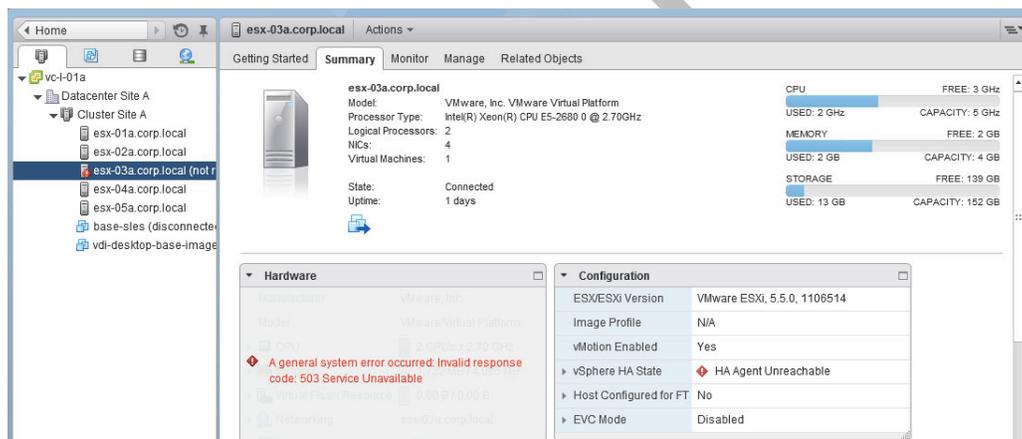
Basically any components on the rebooted host show up as **Absent**. When the host rejoins the cluster, all components are put back in an Active state. A bitmap of blocks that have changed since the component went absent is maintained. The resync process only needs to resync changed blocks. Now we can see one part the availability aspect of VSAN, and how virtual machines continue to run even if components go absent. If the host remains absent for more than 30 minutes, the missing components are rebuilt (reconfigured) on the remaining hosts and disks in the cluster.

6.4 Host Failure - Running VMs

Wait for the host to reboot from the previous test before continuing. Remember that we have only set *ComponentFailuresToTolerate* to 1. In this next example, we will halt the ESXi host (in this example, esx-03a), which contains a running VM **base-sles**. Here we will see interoperability between HA and VSAN.

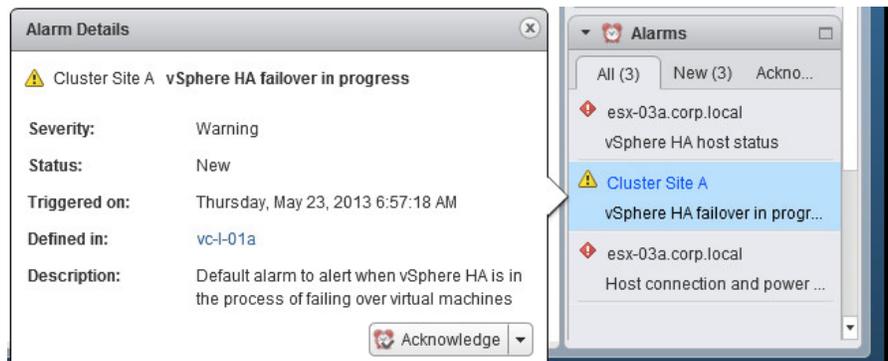
From the Control Center desktop Start button, navigate to All Programs > PuTTY > PuTTY. Select the host esx-04 from the list and launch an SSH session. Login with the credentials *root* and give the *VMware1!* password. Type the command **halt** in the shell.

Once again, you will see vSphere HA detect the error. And as before, vCenter reports on the host connection and power state:

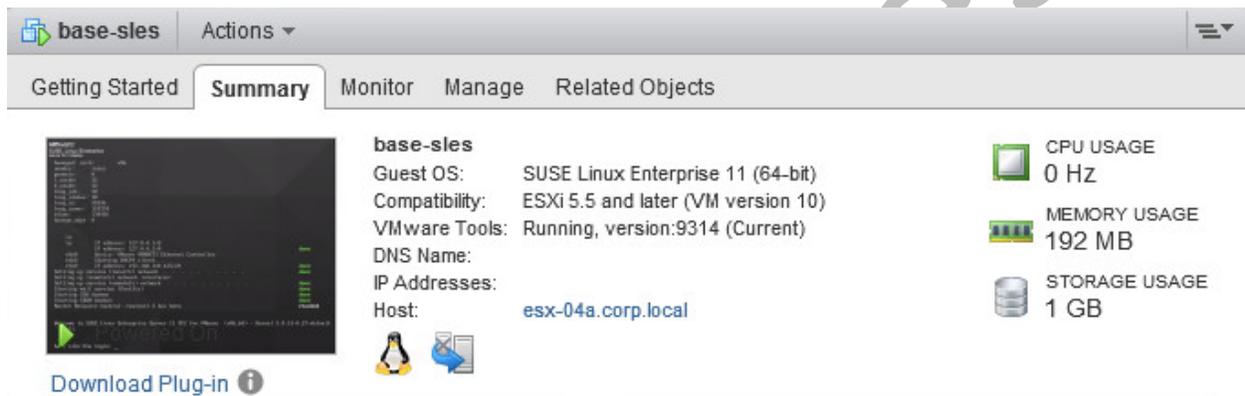


If you go immediately to look at the VM's storage object layout, you might find that you can no longer query the component state since the host that owned the object has gone down. Once the VM has successfully failed over to an alternate host, you will once again be able to query the object layout. You should see vSphere HA kicking in, and failing over the VM from the failed host to another host in the cluster.

Virtual SAN (beta refresh) POC Guide



Finally, check the status of the VM. In my test, the VM is successfully started on host esx-04a (although it may be restarted on a different host in your environment). Note that you should refresh the UI periodically to see these changes occur.



The VM is running, now examining the object layout. It should reveal that not all components are present. However there is still a quorum of objects available, enabling the VM to tolerate a host failure. If the failure persists for longer than 30 minutes, the components will be rebuilt on the remaining disks in the cluster.

Key Takeaway: While the policy can provide highly available virtual machine storage objects, interoperability between VSAN & vSphere HA provides high availability at both the storage AND compute layers .

This Completes the POC.

Virtual SAN (beta refresh) POC Guide

Appendix A: Troubleshooting Virtual SAN

In this appendix, a number of utilities that can help you to troubleshoot VSAN are covered.

A1. Troubleshooting VSAN using the ESXCLI

vSphere 5.5 introduces a new **esxcli** namespace specifically for Virtual SAN.

```
~ # esxcli vsan
Usage: esxcli vsan {cmd} [cmd options]

Available Namespaces:
  datastore      Commands for VSAN datastore configuration
  network        Commands for VSAN host network configuration
  storage        Commands for VSAN physical storage configuration
  cluster        Commands for VSAN host cluster configuration
  maintenancemode Commands for VSAN maintenance mode operation
  policy         Commands for VSAN storage policy configuration
  trace          Commands for VSAN trace configuration

~ # █
```

Refer to the VSAN Troubleshooting guide for further information on how to use the ESXCLI command for troubleshooting VSAN issues.

Virtual SAN (beta refresh) POC Guide

A2. Troubleshooting VSAN using the RVC

The vCenter Appliance now ships with RVC, the Ruby Virtual Console. This tool is an excellent troubleshooting aid for monitoring and troubleshooting issues on a VSAN. The tool is available in vCenter 5.5 and is pre-installed on both the Windows version of vCenter and the appliance version.

To get started with the RVC, from your vCenter command line type **rvc <hostname>** where <hostname> can be an ESXi host or your vCenter server. Navigation of the inventory is done using **ls** and **cd** commands. Each inventory object is associated with a numeric value.

RVC comes with an extensive help subsystem. Type **help vsan** for a list of VSAN supported commands. Type **help vsan.<command>** for help on a specific command:

```
> help vsan
Commands:
enable_vsan_on_cluster: Enable VSAN on a cluster
disable_vsan_on_cluster: Disable VSAN on a cluster
host_consume_disks: Consumes all eligible disks on a host
host_info: Print VSAN info about a host
cluster_info: Print VSAN info about a cluster
disks_info: Print physical disk info about a host
cluster_set_default_policy: Set default policy on a cluster
object_info: Fetch information about a VSAN object
disk_object_info: Fetch information about all VSAN objects on a given physical disk
cmmms_find: CMMDS Find
vm_object_info: Fetch VSAN object information about a VM
disks_stats: Show stats on all disks in VSAN

To see commands in a namespace: help namespace_name
To see detailed help for a command: help namespace_name.command_name
> █
```

Refer to the VSAN Troubleshooting guide for further information on how to use the RVC for troubleshooting VSAN issues. There are many other commands that can assist you in getting additional details around VSAN. It is worthwhile spending some time evaluating the RVC and the various commands that you have at your disposal.

Virtual SAN (beta refresh) POC Guide

A3. Troubleshooting VSAN using the VSAN Observer

VSAN Observer is a performance-monitoring tool that is invoked through the RVC of vCenter. It is web based and provides granular performance statistics on each of the components in a Virtual SAN. It provides live statistics every 60 seconds and allows an administrator to identify bottlenecks anywhere within the VSAN environment.

As well as looking at the VSAN cluster as a whole, performance information can be retrieved on SSD IOPS, Latency and Read Cache Hit Ratio, Magnetic disk performance, VSAN CPU utilization & Virtual Machine VMDK performance.

Data is refreshed every 60 seconds to give the “Average” for that 60 second period.

There are three ways that VSAN Observer can be ran

- Live Monitoring
 - Can be cancelled
 - Automatically terminates after two hours
- Offline Monitoring
 - Data can be replayed on another system through a web browser
- RAW stats bundle
 - For deeper “typically” engineering analysis

From an RVC command window navigate to your VSAN cluster and run the following command:

```
vsan.observer ~/computers/VSAN <options>
```

where the <options> can be:

- Live Monitoring : **--run-webserver --force**
- Offline Monitoring : **--generate-html-bundle**
- RAW Stats Bundle : **--filename**

By default, VSAN Observer will output to port 8010, but this can be changed in the options.

To see the VSAN performance, point a browser to the IP Address of your vCenter server and the port specified when starting the observer tool.

Virtual SAN (beta refresh) POC Guide

VSAN Observer – VSAN client (per host) view:



This is the VSAN Observer landing page and gives a good overview of the VSAN cluster as a whole. This cluster has three ESXi hosts. Essentially what you are looking for is that each of the graphs that are displaying information is 'green'. Green means that everything is operating within its expected parameters. Each of the graphs can be clicked on for a full-sized view.

One of the things we have heard from many of our customers is that they do not like the fact that storage is sometimes a black box. Hopefully, this overview of VSAN Observer will help you to see that VSAN is not a black box from a storage perspective, and that the monitoring and troubleshooting of VSAN is a primary focus of VMware.

Refer to the VSAN Troubleshooting guide for further information on how to use the VSAN Observer for troubleshooting VSAN issues.