

InfoScale Enterprise Across VMware Virtual SAN

VERITAS™

Continuous shared access across VMware Virtual SAN

*Carlos Carrero
Sunny Anthony*

Table of Contents

- 01** Introduction
- 02** VMware Configuration
- 03** InfoScale Enterprise Configuration
- 04** Storage Settings
- 05** Operations

Introduction

- + *vSAN architecture with Veritas Cluster File System*
- + *Flexible Storage Sharing*
- + *Infrastructure*

01

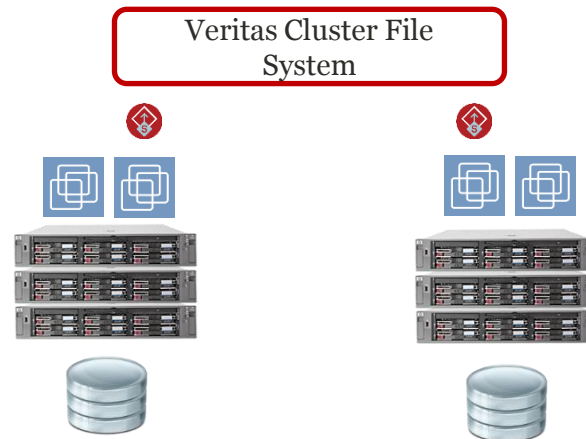
VSAN ARCHITECTURE WITH CLUSTER FILE SYSTEM

VMware Virtual SAN is a shared storage solution for hyperconverged infrastructure that is optimized for vSphere virtual machines. It uses local storage that pools into a datastore where VMs consume that storage.

Veritas InfoScale Storage/Enterprise includes Cluster File System which provides parallel access to the same file system from multiple nodes.

While vSAN just provide a volume that will be consumed by one VM, Cluster File System runs within the VM to actually manage the data and provide concurrent access from any VM.

Veritas Cluster File System can create a shared file system across any VM which is using Virtual SAN, providing fast failover and ETL (Extract-Transform-Load) processing across VMs



FLEXIBLE STORAGE SHARING


Veritas Cluster File System can run inside a VMware Virtual Machine and consume the storage provided in form of VMDK or raw devices. Traditionally, that storage had to be shared across all the VMs so all the VMs had access to the same storage at the same time.

By using Flexible Storage Option, local disk in each of the nodes can be used to create a shared file system. When combining it with vSphere it means that each VM can consume its local storage.

Virtual SAN can also use local storage and still provide resiliency for the VMs within a vSphere cluster. In case of a failure, vSphere will make sure the infrastructure is still operational.

Veritas Cluster File System will guarantee a faster recovery of the application and a continuous access to the data from any node.

While vSphere will guarantee the infrastructure availability, Veritas Cluster File System will accelerate application recovery and will maintain continuous data access across any VM



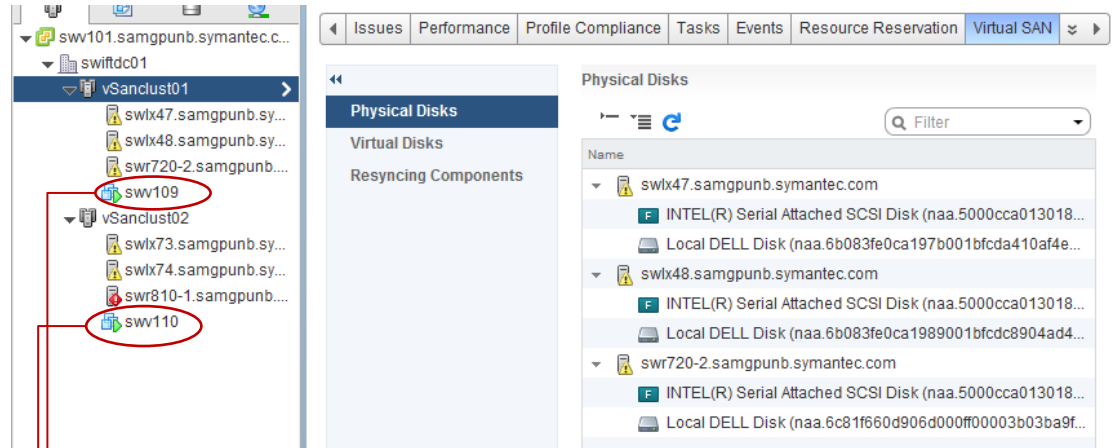
INFRASTRUCTURE

To demo this configuration we are going to use two different and independent Virtual SAN clusters. One VM will be running in each cluster that will be used to create a Veritas Cluster File System (CFS).

Flexible Storage Option for CFS will be used so that one copy of data resides in each of the clusters at all times. Any of the clusters can go down that CFS will keep the other VM running.

Data will be accessed from any of the VMs as they will have a shared file system.

This guide will highlight the key processes but it is not intended to be a complete step by step guide.



Veritas Cluster File System will be created across these two VMs

Each VM will consume local storage from vSAN datastore

VMware Configuration

- + *Enable vSAN*
- + *Switch configuration*
- + *Virtual Machine creation*

02

ENABLE VSAN

We have to enable VSAN in order to consume local storage from each cluster.

All the disks will be selected and one datastore will be created on top. Each of the clusters will have a different datastore. From each one, VMDK files will be created and attached to each of the VMs.



vSanclust01 - Claim Disks for Virtual SAN Use

When claiming disks for Virtual SAN, the following rules apply:

1. Hosts that contribute to the Virtual SAN cluster must contribute at least 1 flash disk.
 2. The number of HDD disks that a host contributes must be greater than or equal to the number of flash disks it contributes.
- More flash disks means better performance, and more HDD disks mean more datastore capacity.



Select all eligible disks ⌵ ⌵ Show: **Eligible hosts** ⌵ 🔍 Filter

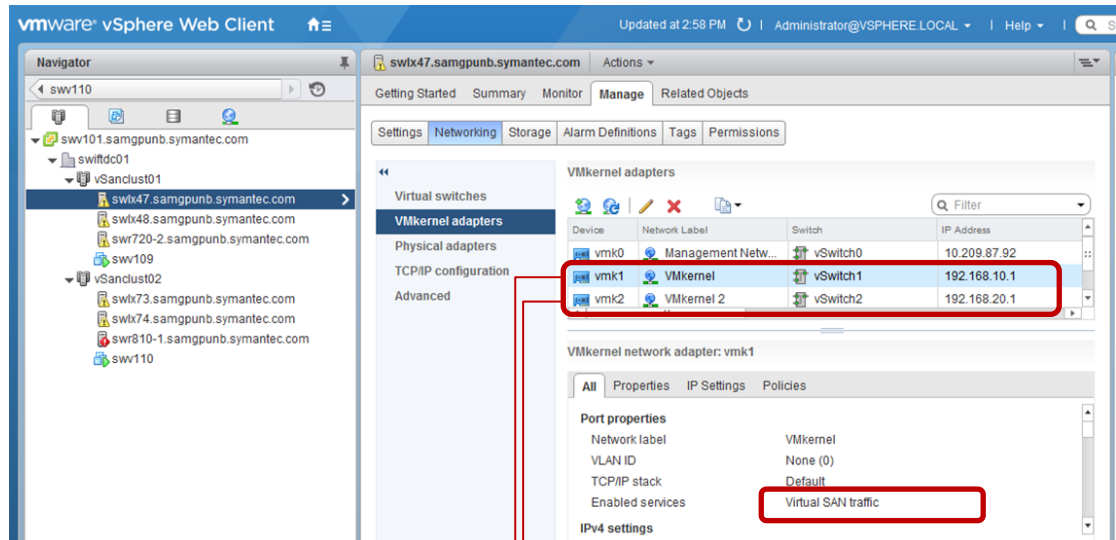
Name	Drive Type	Capacity	Transport Type
<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> swlx47.samgpunb.symantec.com			
<input checked="" type="checkbox"/> Local DELL Disk (naa.6b083fe0ca197b001bfcda410af4eb3f)	HDD	1.09 TB	Parallel SCSI
<input checked="" type="checkbox"/> INTEL(R) Serial Attached SCSI Disk (naa.5000cca013018be0)	Flash	186.31 GB	Block Adapter
<input checked="" type="checkbox"/> INTEL(R) Serial Attached SCSI Disk (naa.5000cca013018be4)	Flash	186.31 GB	Block Adapter
<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> swlx48.samgpunb.symantec.com			
<input checked="" type="checkbox"/> INTEL(R) Serial Attached SCSI Disk (naa.5000cca013018c28)	Flash	186.31 GB	Block Adapter
<input checked="" type="checkbox"/> Local DELL Disk (naa.6b083fe0ca1989001bfcddc8904ad4ee2)	HDD	1.09 TB	Parallel SCSI
<input checked="" type="checkbox"/> INTEL(R) Serial Attached SCSI Disk (naa.5000cca013018c2c)	Flash	186.31 GB	Block Adapter
<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> swr720-2.samgpunb.symantec.com			
<input checked="" type="checkbox"/> Local DELL Disk (naa.6c81f660d906d000ff00003b03ba9fb6)	HDD	278.88 GB	Parallel SCSI
<input checked="" type="checkbox"/> INTEL(R) Serial Attached SCSI Disk (naa.5000cca013018a70)	Flash	186.31 GB	Block Adapter
<input checked="" type="checkbox"/> INTEL(R) Serial Attached SCSI Disk (naa.5000cca013018a74)	Flash	186.31 GB	Block Adapter
<input checked="" type="checkbox"/> Local DELL Disk (naa.6c81f660d906d0001a715aa509f60233)	HDD	278.88 GB	Parallel SCSI
<input checked="" type="checkbox"/> Local DELL Disk (naa.6c81f660d90c1f001a7159ab07d40ff7)	HDD	278.88 GB	Parallel SCSI

VSAN NETWORK CONFIGURATION

In order to Virtual SAN to be properly configured we need to add one VMKernel adapter on each ESX hosts within the cluster.

Select the host, then click on the Networking tab and select VMKernel adapters. Click on the Add host networking icon and on the connection type select VMKernel network adapter.

Same operation needs to be performed in order to enable vMotion. The difference will be the type of traffic to be used within the interface.



This VMKernel network will be used for vMotion traffic

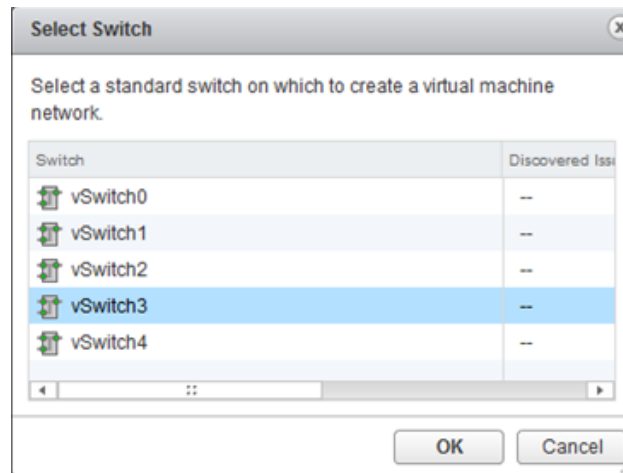
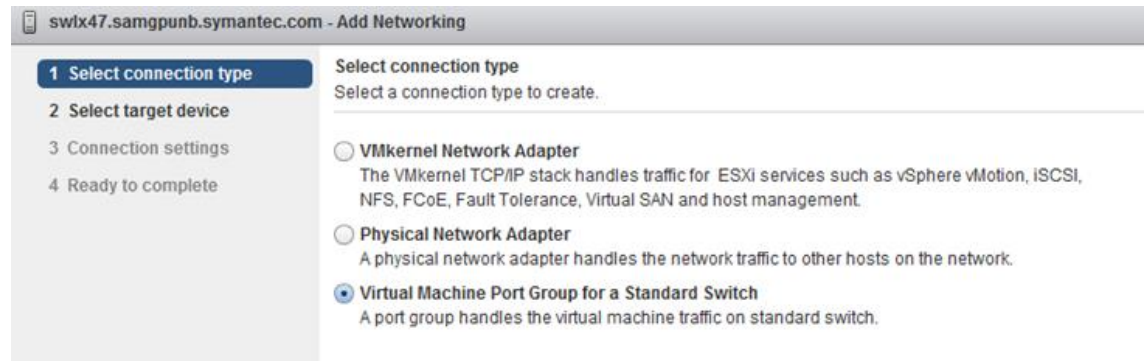
This VMKernel network will be used for Virtual SAN traffic

SWITCH CONFIGURATION FOR CLUSTER FILE SYSTEM TRAFFIC

The VMs need to have at least two private links to provide heartbeat information and share I/O between them.

We will create two standard switches for the Cluster File System traffic.

This network will be later used during InfoScale Enterprise configuration.



VIRTUAL MACHINE CREATION

Select storage
Select the datastore in which to store the configuration and disk files

VM Storage Policy: **Virtual SAN Default Storage Policy**

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)	2.43 TB	0.00 B	2.43 TB	vsan	
Incompatible					
datastore1 (5)	1.08 TB	978.00 MB	1.08 TB	VMFS	
datastore1 (6)	1.08 TB	978.00 MB	1.08 TB	VMFS	
datastore1 (9)	271.25 GB	974.00 MB	270.30 GB	VMFS	

Each VM will use the vSAN Datastore

Customize hardware

Configure the virtual machine hardware

Virtual Hardware | VM Options | SDRS Rules

- *CPU: 16
- *Memory: 16 GB
- New Hard disk: 60 GB
- New SCSI controller: LSI Logic SAS
 - SCSI Bus Sharing: None
 - Change Type: LSI Logic SAS
- *New Network: VM Network Connect...
- New CD/DVD Drive: Client Device Connect...
- New Floppy drive: Client Device Connect...
- Video card: Specify custom settings
- VMCI device
- New SATA Controller
- Other Devices
- New Network: PrivSW01 Connect...
- *New Network: PrivSW02 Connect...

New device: Network

When creating the VM, the two private network links will also be attached

InfoScale Enterprise Configuration

+ *Installing InfoScale Enterprise 7.0*

03

INSTALLING INFOSCALE ENTERPRISE 7.0

The install package can be downloaded in any of the servers and run from there.

For this exercise InfoScale Enterprise and SFCFSHA will be chosen as it contains both Cluster File System and all the High Availability agents that will provide a fast failover configuration.

If ssh has not been enabled between systems, the installer will automatically set it up.

- 1) Veritas InfoScale Foundation
- 2) Veritas InfoScale Availability
- 3) Veritas InfoScale Storage
- 4) Veritas InfoScale Enterprise**

Select a product to perform pre-installation check for: [1-4,q] 4

- 1) Cluster Server (VCS)
- 2) Storage Foundation (SF)
- 3) Storage Foundation and High Availability (SEHA)
- 4) Storage Foundation Cluster File System HA (SFCFSHA)**
- 5) Storage Foundation for Oracle RAC (SF Oracle RAC)

Select a component to perform pre-installation check for: [1-5,q] 4

Enter the superuser password for system swv110:

- 1) Setup ssh between the systems
- 2) Setup rsh between the systems
- b) Back to previous menu

Select the communication method [1-2,b,q,?] (1) 1

INSTALLING INFOSCALE ENTERPRISE 7.0

The installer will auto-detect if any required package is missing. We recommend yum to be configured, so the installer can automatically install any required package.

The installer provides some guidance about how to install OS rpms using native methods, like yum, or how to manually install the required OS rpms.

- 1) **Install the missing required OS rpms with yum, if yum is configured on the systems**
- 2) Install the missing required OS rpms manually, (detailed steps are provided)
- 3) Do not install the missing required OS rpms

How would you like to install the missing required OS rpms? [1-3,q,?] (1)

Keyless option enables the software utilization during 60 days. Systems need to be registered with VOM after that period of time.

- 1) Enter a valid license key
- 2) Enable keyless licensing and complete system licensing later

How would you like to license the systems? [1-2,q] (2) 2

- 1) Veritas InfoScale Foundation
- 2) Veritas InfoScale Availability
- 3) Veritas InfoScale Storage
- 4) Veritas InfoScale Enterprise
- b) Back to previous menu

Which product would you like to register? [1-4,b,q] (4) 4

Registering keyless key ENTERPRISE on Veritas InfoScale Enterprise
Successfully registered ENTERPRISE keyless key on **swv109**
Successfully registered ENTERPRISE keyless key on **swv110**

INSTALLING INFOSCALE ENTERPRISE 7.0

The two networks that were previously created with the VM will be used as private links for the Veritas Cluster File System

```
Veritas InfoScale Enterprise 7.0 Configure Program
swv109 swv110
```

- 1) Configure the heartbeat links using LLT over Ethernet
- 2) Configure the heartbeat links using LLT over UDP
- 3) Configure the heartbeat links using LLT over RDMA
- 4) Automatically detect configuration for LLT over Ethernet
- b) Back to previous menu

```
How would you like to configure heartbeat links? [1-4,b,q,?] (4) 1
```

Networks are automatically detected and configured

```
Checking media speed for eno33559296 on swv109..... 10000Mb/s
Checking media speed for eno50338560 on swv109 ..... 10000Mb/s
Checking media speed for eno33559296 on swv110 ..... 10000Mb/s
Checking media speed for eno50338560 on swv110 ..... 10000Mb/s
```

Storage Settings

+ *Attach VMDKs*

+ *Disk.EnableUUID setting*

+ *Disk initialization*

+ *Disk Group creation*

+ *Volume attributes*

+ *File System attributes and configuration*

+ *The usage of the multi-writer flag*

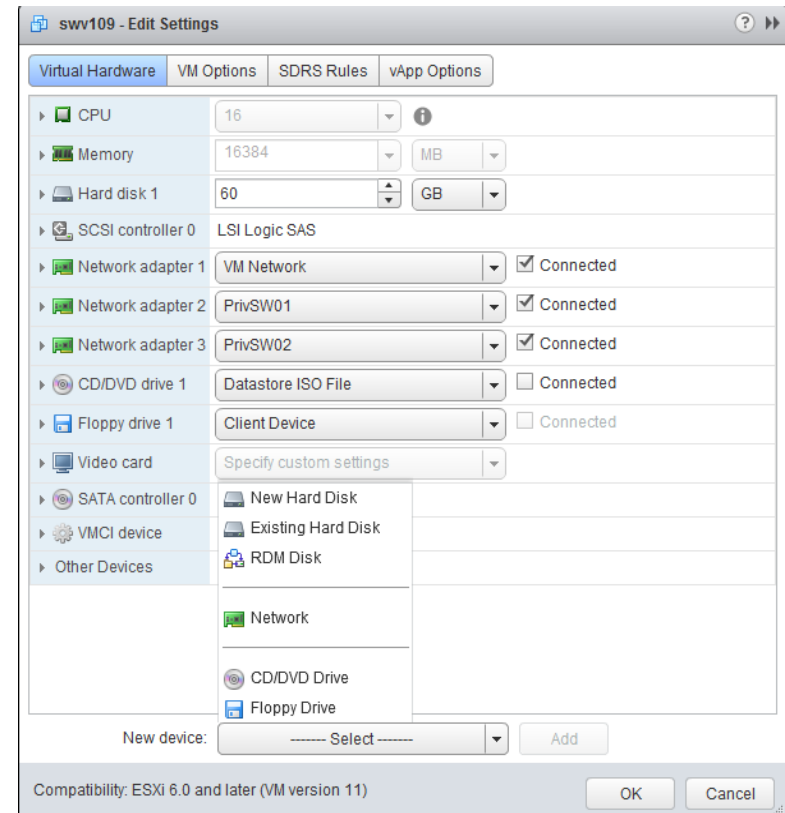
04

ATTACH VMDKS

We are going to attach three new disks of 16GB each one to the VM. Those disks will reside on the vSAN data store.

Select Edit Setting for the VM and select New Hard Disk, using the vSAN data store.

▶ Hard disk 1	60	GB
▼ Hard disk 2	16	GB
Maximum Size	2.40 TB	
Virtual SAN storage consumption	32.00 GB disk size on datastore 8.00 MB used storage space 0.00 B reserved flash space	
VM storage policy	Virtual SAN Default Storage Policy	
Type	As defined in the VM storage policy	
Disk File	[vsanDatastore (1)] 49c80b56-018b-ebb-d-b586-b82a72d64377/swv109_1.vmdk	



DISK.ENABLEUUID SETTING

There is one important step when configuring Veritas Cluster File System on VMware and the storage is provided by VMDK files.

VMware offers a setting that provides an unique UUID (identifier) for each VMDK. This is important so the disks will provide a consistent identifier upon any reboot.

Once a VM is created that is going to be used with CFS, this setting will be enabled.

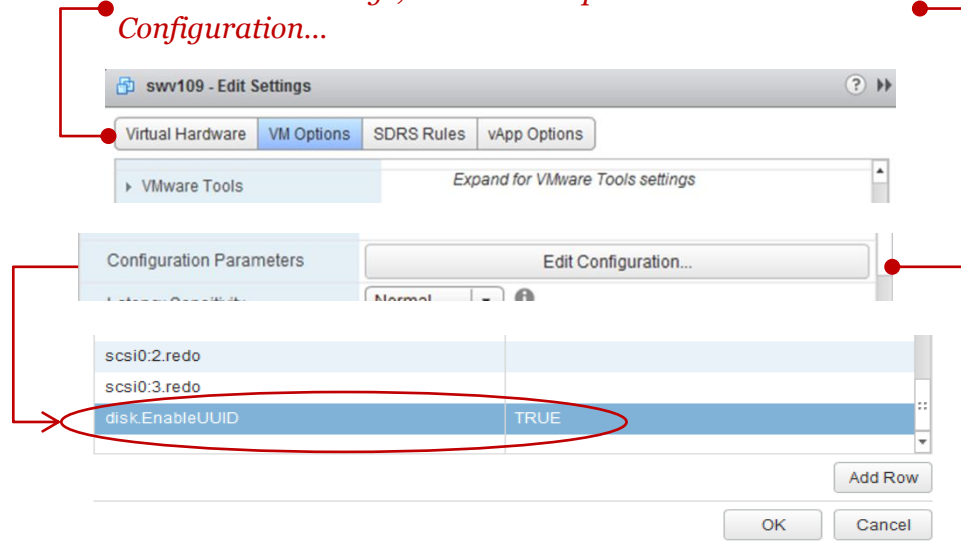
If the setting has not been previously enabled, the disks appear as sdb, sdc, etc.. This is not correct and means that the Veritas ASL (Array Support Library) is not recognizing the disks properly.

Once `disk.EnableUUID` has been set at TRUE, the disks are properly recognized within the VM with the vmdk* prefix on the disk name.

```
swv109.samgpunb.symantec.com:/root> vxdisk list
DEVICE      TYPE      DISK      GROUP      STATUS
sda         auto:none -          -          online invalid
sdb         auto:none -          -          online invalid
sdc         auto:none -          -          online invalid
sdd         auto:none -          -          online invalid
```

● *The disks are not properly recognized*

● *Click on Edit Settings, choose VM Options and click on Edit Configuration...*



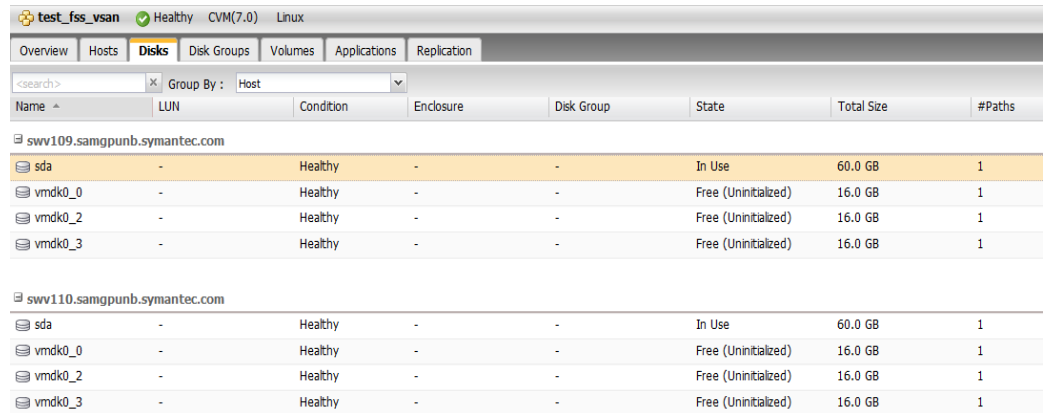
```
vmdk0_0     auto:none -          -          online invalid
vmdk0_2     auto:none -          -          online invalid
vmdk0_3     auto:none -          -          online invalid
```

DISK INITIALIZATION

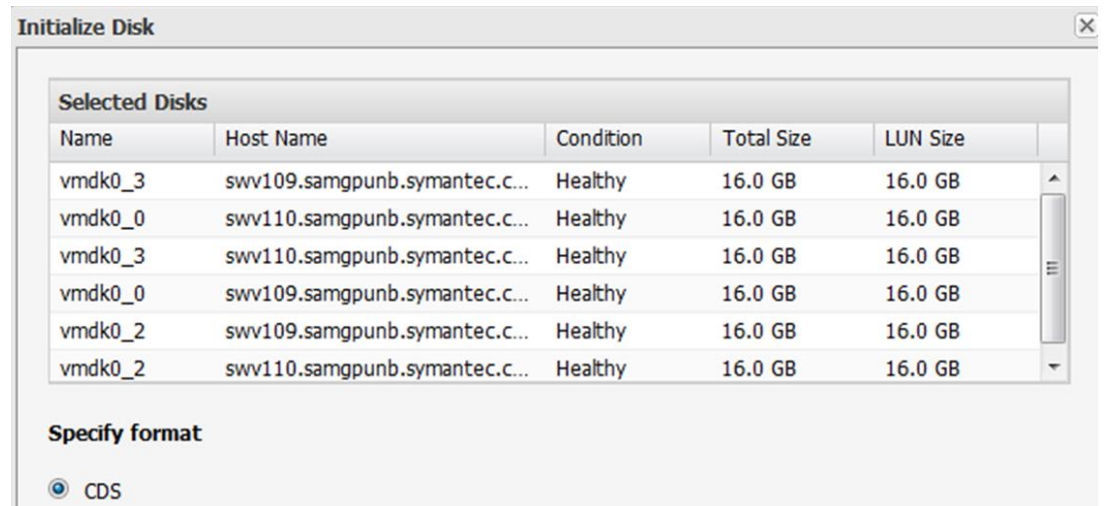
Once the disks are properly configured with the VMs, a file system will be created. Note that each VM is treating each disks as local, but we are going to create a global namespace across the two VMs.

Our first step is to export those disks so they are visible across the cluster. All these tasks can be performed either using CLI or the VOM GUI.

All the 6 disks can be selected and initialized with default CDS format.



Name	LUN	Condition	Enclosure	Disk Group	State	Total Size	#Paths
swv109.samgpunb.symantec.com							
sda	-	Healthy	-	-	In Use	60.0 GB	1
vmdk0_0	-	Healthy	-	-	Free (Uninitialized)	16.0 GB	1
vmdk0_2	-	Healthy	-	-	Free (Uninitialized)	16.0 GB	1
vmdk0_3	-	Healthy	-	-	Free (Uninitialized)	16.0 GB	1
swv110.samgpunb.symantec.com							
sda	-	Healthy	-	-	In Use	60.0 GB	1
vmdk0_0	-	Healthy	-	-	Free (Uninitialized)	16.0 GB	1
vmdk0_2	-	Healthy	-	-	Free (Uninitialized)	16.0 GB	1
vmdk0_3	-	Healthy	-	-	Free (Uninitialized)	16.0 GB	1



Initialize Disk

Selected Disks

Name	Host Name	Condition	Total Size	LUN Size
vmdk0_3	swv109.samgpunb.symantec.c...	Healthy	16.0 GB	16.0 GB
vmdk0_0	swv110.samgpunb.symantec.c...	Healthy	16.0 GB	16.0 GB
vmdk0_3	swv110.samgpunb.symantec.c...	Healthy	16.0 GB	16.0 GB
vmdk0_0	swv109.samgpunb.symantec.c...	Healthy	16.0 GB	16.0 GB
vmdk0_2	swv109.samgpunb.symantec.c...	Healthy	16.0 GB	16.0 GB
vmdk0_2	swv110.samgpunb.symantec.c...	Healthy	16.0 GB	16.0 GB

Specify format

CDS

DISK INITIALIZATION

Once the disks are initialized, they can be exported so they will be visible across all the servers.

The screenshot displays a storage management interface for two servers: **swv109.samgpunb.symantec.com** and **swv110.samgpunb.symantec.com**. Each server has a list of disks: **sda**, **vmdk0_0**, **vmdk0_2**, and **vmdk0_3**. The **vmdk0** disks are highlighted in orange. A context menu is open over the **vmdk0_0** disk of the second server, listing various actions. The **Export** option is highlighted with a red box. A red arrow points from the text below to this **Export** option.

Server	Disk	State	Health
swv109.samgpunb.symantec.com	sda	-	Healthy
	vmdk0_0	-	Healthy
	vmdk0_2	-	-
	vmdk0_3	-	-
swv110.samgpunb.symantec.com	sda	-	-
	vmdk0_0	-	-
	vmdk0_2	-	-
	vmdk0_3	-	-

- Create Disk Group
- Add To Disk Group
- Remove From Disk Group
- Online
- Offline
- Disconnect
- Reclaim Thin Storage
- Resize
- Rename
- Evacuate
- Replace
- Recover
- Initialize
- Export**
- Unexport

Export the disks so they are visible across the nodes

DISK GROUP CREATION

Select the 6 disks and click on Create Disk Group. Enter any name and select all the disk available.

Create Disk Group - Disk Group Options

Disk Group Name :

Enable cross-platform data sharing

Shared (Applicable for clustered host)

Enable Flexible Storage Sharing

Coordinator (Coordinator disk group is used by VCS)

Description :

Create Disk Group - Disk Selection

Disks matching the below filter criteria [Edit](#)

Enclosure is: **Any** Vendor is: **Any** LUN Type is: **Any**

No of Active Paths: **Any**

Additional Criteria: **Free VxVM disks on the host: swv109.samgpunb.symantec.com**

<input checked="" type="checkbox"/>	Name	Condition	Enclosure	State	Total Size	Free Size	Thin	# Paths
<input checked="" type="checkbox"/>	vmdk0_2	Healthy		Free (Ini...	15.97 GB			1
<input checked="" type="checkbox"/>	vmdk0_3	Healthy		Free (Ini...	15.97 GB			1
<input checked="" type="checkbox"/>	vmdk0_0	Healthy		Free (Ini...	15.97 GB			1
<input checked="" type="checkbox"/>	vmdk0_0_1	Healthy		Free (Ini...	15.97 GB			0
<input checked="" type="checkbox"/>	vmdk0_2_1	Healthy		Free (Ini...	15.97 GB			0
<input checked="" type="checkbox"/>	vmdk0_3_1	Healthy		Free (Ini...	15.97 GB			0

Add all the available disks into the Disk Group

VOLUME ATTRIBUTES

Select Create Volume to get the volume attributes tab

The screenshot shows the 'Create Volume - Volume Attributes' dialog box. The 'Volume Name' is 'vol1' and the 'Size' is '30 GB'. The 'Layout' section has 'Mirror-Stripe' selected. The 'Mirror Info' section shows 'Total Mirrors' as 2 and 'Mirror across' as Host. The 'Stripe Info' section shows 'Columns' as 3, 'Stripe Unit Size' as 128, and 'Stripe across' as None. The 'Hosts' list includes 'swv110' and 'swv109'. The 'Disk selection method' is set to 'Automatic'.

Name and size

Two copies across hosts

Hosts involved

Stripe internally across the 3 available disks

FILE SYSTEM ATTRIBUTES

Next step is to create the file system

Do not create file system

Create file system

File System Type: vxfs

File System options

Block size: 8192 Allocation Unit: Default

Size: 30 GB

Mount Options

Mount point: /vol1

Read only Add to file system table

Honor setuid Mount on boot

SmartIO writeback fsck pass

Mount Type: Local Cluster

Select nodes to mount the file system:

<input checked="" type="checkbox"/>	swv110
<input checked="" type="checkbox"/>	swv109

Specify the mount point in the VM

Block Size

Cluster option so the file system can be mounted in more than one host at a time

What hosts will mount the file system

FINAL FILE SYSTEM CONFIGURATION

Following previous example, as many file systems as needed can be configured across the VMs.

In the screenshot we can see how the volume named vol1 is mounted under /vol1 in both hosts, swv109 and swv110.

The screenshot displays the Veritas VxFS configuration interface for a cluster named 'test_fss_vsan'. The 'Volumes' tab is active, showing a table of configured volumes. Below this, the 'Mount Points' section shows the configuration for two hosts: swv110 and swv109.


Name	Condition	State	Disk Group	Size	Layout	FS Type	Mount Point	FS Size
swv109.samgpunb.symantec.com								
vol1(/vol1)	Healthy	Healthy	cfsdg	30.0 GB	Mirrored-Stripe/RAI...	vxfs	/vol1	30.0 GB
swv110.samgpunb.symantec.com								
vol1(/vol1)	Healthy	Healthy	cfsdg	30.0 GB	Mirrored-Stripe/RAI...	vxfs	/vol1	30.0 GB

Host	Mount Point	Mounted	CFS Role
swv110.samgpunb.symantec.com	/vol1	Yes	Secondary
swv109.samgpunb.symantec.com	/vol1	Yes	Primary

THE USAGE OF THE MULTI-WRITER FLAG

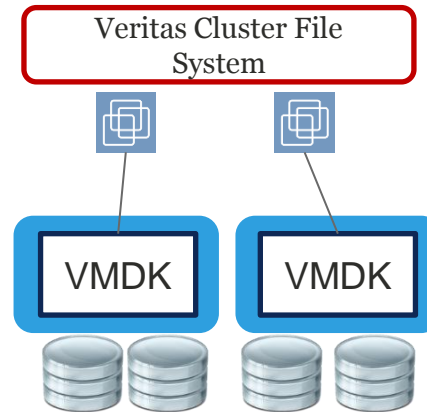
It is common to use the VMware multi-writer flag when Veritas Cluster File System is using shared storage. In that case, two or more VMs need to access to the same VMDK file at the same time. That flag is used to override the default VMware protection. Veritas Cluster File System provides the data consistency so several nodes can access concurrently the same data.

In the approach taken in this architecture there is no need for multi-writer flag. Each VM is consuming local storage and the VMDK file is solely attached to one VMDK. We use Cluster File System Flexible Storage Option to make that VMDK visible to all the other nodes within the cluster and to create redundancy across them.



In the figure at the top, there is a single data store that contains two VMDK files. Each of the VMs need to get access to the VMDK files, so we need the multi-writer flag to be enabled.

The figure below is using two data stores and there is one VMDK file in each one. The VMDK file only needs to be mapped to the local VM. Using Flexible Storage Sharing feature, a file system can be created on top of the two VMDK files without having to use the multi-writer flag.



Flexible Storage Sharing allows the usage of local storage to create a shared storage architecture.

Operations

+ *vMotion*

+ *VM failure*

+ *ESX failure*

05

VMOTION

VMware vMotion can be run in any of the VMs. The advantage is that the VM can be moved to any other physical hosts within the same VMware cluster.

On top, Veritas Cluster File System will continue running without noting that the VM is now running in a different physical hosts.

The image displays three sequential screenshots of the VMware vMotion migration wizard for a VM named 'swv109'.

Step 1: Select the migration type
 The wizard asks to 'Select the migration type' and provides three options:

- Change compute resource only: Migrate the virtual machines to another host or cluster.
- Change storage only: Migrate the virtual machines' storage to a compatible datastore or datastore cluster.
- Change both compute resource and storage: Migrate the virtual machines to a specific host or cluster and their storage to a specific datastore or datastore cluster.
 - Select compute resource first
 - Select storage first

Step 2: Select a compute resource
 The wizard asks to 'Select a compute resource' and provides a table of available resources:

Name	Cluster
swk48.samgpunb.symante...	vSanclust01

Step 3: Select network
 The wizard asks to 'Select network' and provides a table of network options:

Source Network	Used By	Destination Network
PrivSW01	1 VMs / 1 Network adapters	PrivSW01
VM Network	1 VMs / 1 Network adapters	VM Network
PrivSW02	1 VMs / 1 Network adapters	PrivSW02

Recent Tasks
 A table showing the completion of the migration:

Task Name	Target	Status
Relocate virtual machine	swv109	Completed

VM FAILURE

Veritas Cluster File System has the capability to detect a VM outage in almost no time. This clearly enhance the SLA of vSphere. Recovery times for vSphere are longer, so Cluster File System will detect the application, VM or ESX outage and will migrate the service to any other VM.

We simulate a VM failure running the halt command and in the next second (note different timezones) the other node detect the failure and initiate the service recovery, while VM still think that the VM is up.

A failure is simulated in the first node of the Veritas cluster

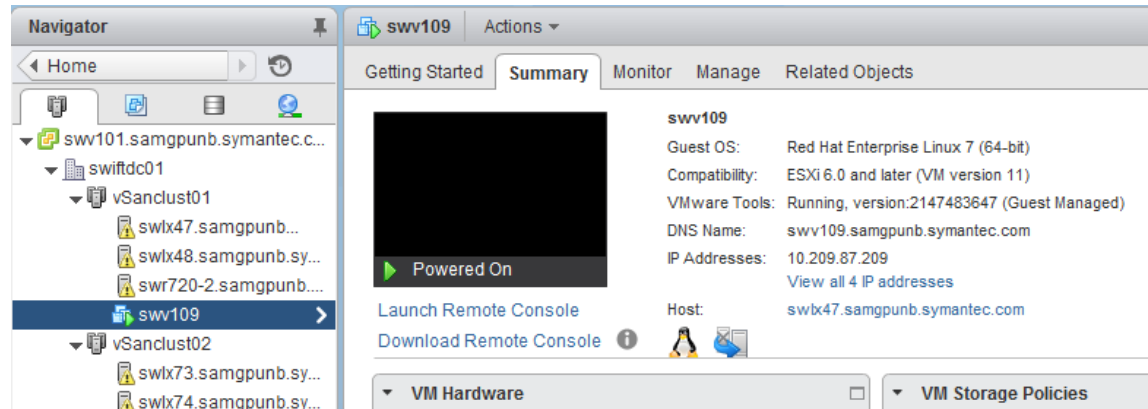
```
swv109.samgpunb.symantec.com:/root>date;halt
Fri Oct 9 18:49:49 BST 2015
```

The second node automatically detects the failure and initiates the app recovery

```
2015/10/09 23:19:50 VCS NOTICE V-16-1-10322 System swv109 (Node '0')
changed state from RUNNING to LEAVING
```

ESX server and VM had different time zones. The difference is just 1 second

vSphere still thinks the VM is up



ESX FAILURE

No matter whether the VM has failed or the ESX server has crashed, Veritas Cluster Server will detect that crash and will make sure the service continues running in another node of the cluster. The file system will still be mounted and data will be accessible from the other nodes.

vSphere HA will detect the ESX failure later and will bring the VM up in another ESX within the vSphere cluster.

By combining Veritas Cluster File System and vSphere HA we get highest protection for our applications. While Veritas takes care of the file system access and application recovery in no time, vSphere guarantees that the failed VM will continue running in another host.

vSphere will restart the VM upon host failure

Cluster monitored failure conditions and responses

Failure	Response	Details
Host failure	Restart VMs	Restart VMs using VM restart priority ordering.
Host Isolation	Disabled	VMs on isolated hosts will remain powered on.
Datastore with Permanent Device Loss	Disabled	Datastore protection for Permanent Device Loss is disabled.
Datastore with All Paths Down	Disabled	Datastore protection for All Paths Down is disabled.
Guest not heartbeating	Reset VMs	VM and application monitoring enabled. VMs will be reset.

Around 8 seconds after the ESX has been powered off, CFS detects the failure

```
2015/10/09 23:49:12 VCS ERROR V-16-1-10079 System swv109 (Node '0') is in Down State - Membership: 0x2
```

Later vSphere HA restart the VM in another ESX server

vSphere HA restarted virtual machine [swv109](#) on host [swlx48.samgpunb.symantec.com](#) in cluster [vSanclust01](#)