



**vSAN datastore support for VMDK shared using multi-writer  
and InfoScale Enterprise**

# vSAN datastore support for VMDK shared using multi-writer and InfoScale Enterprise

## *Table of Contents*

Executive summary .....	3
Overview .....	4
Veritas InfoScale Enterprise™ .....	4
VMware vSAN .....	5
Setup Diagram .....	5
Creating VMDKs on vSAN datastore: .....	6
Creating eager-zeroed disks (PowerCLI) .....	6
Creating eager-zeroed disks (ESXi command line) .....	6
Test cases .....	8
Conclusion .....	9

# vSAN datastore support for VMDK shared using multi-writer and InfoScale Enterprise

## *Executive summary*

This qualification study demonstrates that InfoScale products can run seamlessly on the multi-writer enabled shared VMDKs created on the vSAN datastore.

When vSAN is enabled, a single shared datastore is presented to all the hosts which are part of the vSAN enabled cluster. Typically, all the hosts will contribute performance (SSD) and capacity (magnetic disks or flash) to the vSAN shared datastore. It leverages the local storage from a number of ESXi hosts which are a part of a cluster. A distributed vsan datastore is then created leveraging the local storage from each of the ESXi hosts. This can then be used for VM placement, and also supports a range of core vSphere technologies like vMotion, DRS and vSphere HA. vSAN is completely integrated with the vSphere.

As per the [VMware Knowledge Base](#), the earlier InfoScale Enterprise support was available for the multi-writer enabled shared VMDKs, created on the VMFS datastore. Now, support is added for multi-writer enabled shared VMDKs created on the vSAN datastore.

# vSAN datastore support for VMDK shared using multi-writer and InfoScale Enterprise

## Overview

Infoscale Enterprise offers great features such as Veritas ODM (Oracle Disk Manager) and node arbitration via server-based IO fencing which increases the performance and provides resiliency.

## Veritas InfoScale Enterprise™

Veritas InfoScale Enterprise provides an end-to-end solution for the enterprise storage management. This solution virtualizes heterogeneous storage over the heterogeneous servers into the logical objects.

Veritas InfoScale Enterprise includes the following components and features:

- Veritas Volume Manager virtualizes any storage that is exposed to the system. The customer can then choose a desired resiliency level with the software RAID of the Volume Manager. Volume Manager can even replicate data across remote sites with the Veritas Volume Replication (VVR) feature. Volume Manager features such as Dynamic Multi-Pathing (DMP), Snapshots, Fast Mirror Resync (FMR) and SmartMove migration provide further resiliency and faster data recovery.
- Veritas Cluster File System (CFS) is a highly consistent and resilient file system. CFS provides global namespace across multiple servers. An application can concurrently access data from any of the nodes in a cluster. CFS tuning can improve the application caching and data-flow to the storage. With the checkpoints and fsck features, data and metadata can be quickly recovered. Further, Veritas File System has demonstrated linear scalability of application performance for a range of common workloads. Thus, it guarantees scale-out compute power for an application.
- Veritas Cluster Server (VCS) maintains cluster membership information at all times and enables building of clusters with a large number of servers. VCS enables seamless failover of an application from one server to another server. Customers can count on business continuity and avoid downtime of the services during failure or maintenance of the servers.
- I/O fencing is a core component of VCS, focused on properly handling a cluster partition event due to the loss of cluster communication. I/O fencing consists of two distinct components, Membership Arbitration and Data Protection. These two components together can deliver maximum data integrity in a cluster environment. Membership arbitration is necessary to ensure, that when the cluster members are unable to communicate over the cluster heartbeat network, only a single subcluster is active.
- In some environments, SCSI3-PR support is not available. There is also a data protection aspect of Non-SCSI3 based fencing. This is implemented through the use of judicious timing. When a fencing race occurs, a minimal time gap is put in place before attempting to bring the application online on the subcluster that wins the race. This is to ensure that there is enough time for the losing node to panic and reboot. In environments that do not support SCSI3-PR, Veritas recommends deployment of Non-SCSI3 based fencing. Non-SCSI3 based fencing can be configured [using CPS servers](#) which are outside of its client InfoScale clusters.

# vSAN datastore support for VMDK shared using multi-writer and InfoScale Enterprise

## VMware vSAN

VMware vSAN provides a unique, software-defined approach to hyperconvergence, leveraging the hypervisor (VMware vSphere) to deliver compute, storage and management in a tightly integrated software stack.

Following tests have been carried out as a sanity qualification on the setup given below.

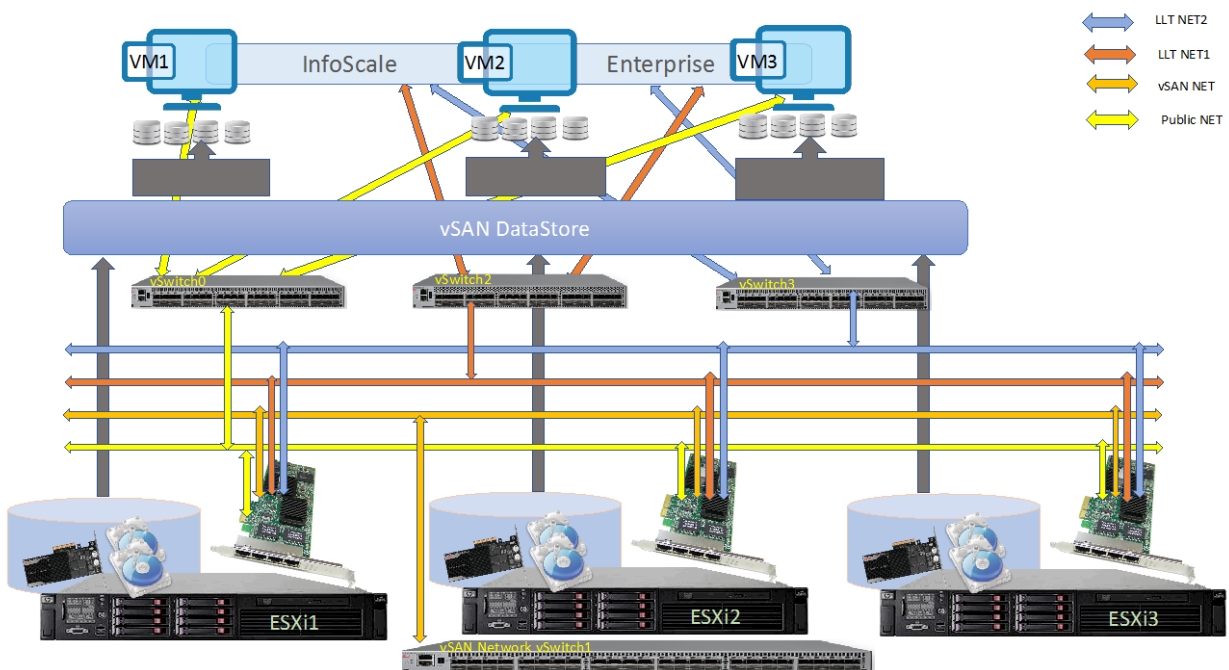
### Hardware Details:

1. 3 x Dell servers
2. HDDs: 4 x SAS HDDs, Capacity - 278 GB
3. SSDs: 1 x INTEL(R) PCIe Device, Capacity – 400 GB

### Software Details:

1. ESXi 6.0 Update 1 GA bits (VMware ESXi 6.0.0 build-3380124)
2. vCenter 6.0 Update 1 GA bits (version 6.0.0 Build 3018524)
3. InfoScale Version: 7.0 CFS (3-node CFS Cluster)
4. Guest OS: Red Hat 7.1

### Setup Diagram



# vSAN datastore support for VMDK shared using multi-writer and InfoScale Enterprise

## **Creating VMDKs on vSAN datastore:**

Creation of eager-zeroed disks on the vSAN datastore was not available through the vCenter web-client version 6.0.0 Build 2997665 but is available in the subsequent versions. Contact VMware support for more information. To create multi-writer enabled VMDKs on the vSAN datastore, follow the procedure as per the [VMware Knowledge Base](#).

In the release used for this testing, the vSphere web client cannot be used to create eager-zeroed disks on the vSAN datastore. While the vSAN datastore supports eager-zeroed disks, initially this functionality was not exposed in the vSphere web client. To accommodate the limitation in this release, we must use either PowerCLI or the ESXi command line to create eager-zeroed disks.

## **Creating eager-zeroed disks (PowerCLI)**

1. Launch PowerCLI and connect to vCenter with the `Connect-VIServer`. For more guidance on PowerCLI conventions, login, and so on, refer to the [VMware vSphere PowerCLI Documentation](#).
2. Create the new hard disk with the `New-HardDisk` cmdlet in the form of:

```
New-HardDisk -VM VM_name -CapacityGB size -Datastore  
vSAN_datastore_name -StorageFormat EagerZeroedThick -Controller  
controller_name -Persistence IndependentPersistent
```

For example, to create a 10GB disk on the RAC\_0 VM, attached to SCSI Controller 1 on the datastore `vsanDatastore`, run this command:

```
New-HardDisk -VM "RAC_0" -CapacityGB 10 -Datastore "vsanDatastore" -  
StorageFormat EagerZeroedThick -Controller "SCSI Controller 1" -  
Persistence IndependentPersistent
```

The name of the storage controller will match the name of the device we added to the VMs in the previous section.

3. Repeat this process for any remaining shared disks to be created.

## **Creating eager-zeroed disks (ESXi command line)**

Note: You must enable the local ESXi shell or SSH access to the host and log in as privileged (root) user to complete the below procedure.

1. Navigate to the directory of the first VM in the Oracle RAC or CFS cluster:

```
cd /vmfs/volumes/vsan_datastore/VM_Name
```

For example:

```
cd /vmfs/volumes/vsanDatastore/RAC_0
```

## vSAN datastore support for VMDK shared using multi-writer and InfoScale Enterprise

2. Create an eager zeroed thick virtual disk to be shared using the vmkfstools:

```
vmkfstools -c size -W vsan -d eagerzeroedthick `pwd`/vmdk_file_name
```

For example:

```
vmkfstools -c 12G -W vsan -d eagerzeroedthick `pwd`/RAC_0_1.vmdk
```

3. Repeat step 2 for as many shared disks, that have to be created.

Once the VMDKs are created, add to each VM using the traditional procedure as mentioned in the above VMware knowledgebase article.

# vSAN datastore support for VMDK shared using multi-writer and InfoScale Enterprise

## Test cases

The following tests are performed and completed successfully to validate the functionality of vSAN VMDKs enabled with the multi-writer flag with InfoScale Enterprise product.

Scenario	Remarks
Installation/Configuration of InfoScale Enterprise (SFCFSHA) 7.0 on a 3-node guest cluster	Install and configure InfoScale Enterprise using common product installer (CPI).
Exporting VMDKs with MWF to all the three guests in sharing mode	Add the VMDKs with MWF using VMware KB mentioned above.
Configure CPS based Non-SCSI3 fencing for Node-Arbitration	As VMDK doesn't support SCSI3-PR, configure CPS based Non-SCSI3 fencing using common product installer (CPI).
Installation/Configuration Oracle 12CR1 database on CFS mounts.	To test application based high-availability, use Oracle database for qualification testing.
Configure VCS agent for Oracle in fast-failover mode	To reduce failover and failback time duration, configure Oracle VCS agent in suggested mode.
Test the failover/failback of VCS Oracle agent	To test the failover and failback of Oracle service while database workload running.
Network split-brain	Simulate the VCS cluster failure by disabling the LLT (Cluster interconnects) and check of database service fails over to available node.
vMotion	Test the live migration of virtual machine which is part InfoScale cluster without affecting any cluster failure.
ESXi node failure	Configure VMware-HA and test the ESXi node failure hosting virtual machine which is part of the InfoScale cluster. Make sure that the VMware-HA starts the failed virtual machine on another ESXi host and virtual machine joins the VCS cluster.



# **vSAN datastore support for VMDK shared using multi-writer and InfoScale Enterprise**

## ***Conclusion***

This qualification testing is to demonstrate that InfoScale supports multi-writer enabled VMDKs created on the vSAN shared datastore. The test scenarios are successfully performed to do a sanity check for the software environment.