

Veritas InfoScale™ Virtualization Guide - Linux on ESXi

Last updated: 2019-02-26

Legal Notice

Copyright © 2019 Veritas Technologies LLC. All rights reserved.

Veritas and the Veritas Logo are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas is required to provide attribution to the third-party ("Third-Party Programs"). Some of the Third-Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the third-party legal notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Veritas Technologies LLC and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. VERITAS TECHNOLOGIES LLC SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Veritas Technologies LLC
500 E Middlefield Road
Mountain View, CA 94043

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

<https://sort.veritas.com/documents>

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

infoscaledocs@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Section 1	Overview	8
Chapter 1	About Veritas InfoScale solutions in a VMware environment	9
	Overview of the Veritas InfoScale Products Virtualization Guide	9
	How Veritas InfoScale solutions work in a VMware environment	10
	How Veritas InfoScale product components enhance VMware capabilities	12
	When to use Raw Device Mapping and Storage Foundation	13
	Array migration	14
	Veritas InfoScale component limitations in an ESXi environment	15
	I/O fencing considerations in an ESXi environment	17
	About Veritas InfoScale solutions support for the VMware ESXi environment	18
	Virtualization use cases addressed by Veritas InfoScale products	18
Section 2	Deploying Veritas InfoScale products in a VMware environment	22
Chapter 2	Getting started	23
	Storage configurations and feature compatibility	23
	About setting up VMware with Veritas InfoScale products	24
	Veritas InfoScale products support for VMware environments	25
	Installing and configuring storage solutions in the VMware virtual environment	25
Chapter 3	Understanding Storage Configuration	26
	Configuring storage	26
	Enabling disk UUID on virtual machines	28
	Installing Array Support Library (ASL) for VMDK on cluster nodes	29

	Excluding the boot disk from the Volume Manager configuration	30
	Creating the VMDK files	30
	Mapping the VMDKs to each virtual machine (VM)	31
	Enabling the multi-write flag	33
	Getting consistent names across nodes	34
	Creating a file system	35
Section 3	Use cases for Veritas InfoScale product components in a VMware environment	37
Chapter 4	Application availability using Cluster Server	38
	About application availability with Cluster Server (VCS) in the guest	38
	About VCS support for Live Migration	38
Chapter 5	Multi-tier business service support	40
	About Virtual Business Services	40
	Sample virtual business service configuration	40
Chapter 6	Improving storage visibility, availability, and I/O performance using Dynamic Multi-Pathing	43
	Use cases for Dynamic Multi-Pathing (DMP) in the VMware environment	43
	How DMP works	44
	How DMP monitors I/O on paths	48
	Load balancing	49
	About DMP I/O policies	50
	Achieving storage visibility using Dynamic Multi-Pathing in the hypervisor	52
	Achieving storage availability using Dynamic Multi-Pathing in the hypervisor	54
	Improving I/O performance with Dynamic Multi-Pathing in the hypervisor	56
	Achieving simplified management using Dynamic Multi-Pathing in the hypervisor and guest	57

Chapter 7	Improving data protection, storage optimization, data migration, and database performance	59
	Use cases for Veritas InfoScale product components in a VMware guest	59
	Protecting data with Veritas InfoScale product components in the VMware guest	61
	About point-in-time copies	61
	Point-in-time snapshots for Veritas InfoScale products in the VMware environment	62
	Optimizing storage with Veritas InfoScale product components in the VMware guest	62
	About SmartTier in the VMware environment	62
	About compression with Veritas InfoScale product components in the VMware guest	63
	About thin reclamation with Veritas InfoScale product components in the VMware guest	64
	About SmartMove with Veritas InfoScale product components in the VMware guest	64
	About SmartTier for Oracle with Veritas InfoScale product components in the VMware guest	65
	Migrating data with Veritas InfoScale product components in the VMware guest	65
	Types of data migration	65
	Improving database performance with Veritas InfoScale product components in the VMware guest	66
	About Veritas InfoScale product components database accelerators	66
Chapter 8	Setting up virtual machines for fast failover using Storage Foundation Cluster File System High Availability on VMware disks	68
	About use cases for InfoScale Enterprise in the VMware guest	68
	Storage Foundation Cluster File System High Availability operation in VMware virtualized environments	69
	Storage Foundation functionality and compatibility matrix	70
	About setting up Storage Foundation Cluster File High System High Availability on VMware ESXi	71
	Planning a Storage Foundation Cluster File System High Availability (SFCFSHA) configuration	72
	Enable Password-less SSH	73

	Enabling TCP traffic to coordination point (CP) Server and management ports	74
	Configuring coordination point (CP) servers	75
	Deploying Storage Foundation Cluster File System High Availability (SFCFSHA) software	79
	Configuring Storage Foundation Cluster File System High Availability (SFCFSHA)	81
	Configuring non-SCSI3 fencing	82
Section 4	Reference	85
Appendix A	Known issues and limitations	86
	Prevention of Storage vMotion	86
Appendix B	Where to find more information	87
	Veritas InfoScale documentation	87
	Service and support	87
	About Veritas Services and Operations Readiness Tools (SORT)	88

Overview

- [Chapter 1. About Veritas InfoScale solutions in a VMware environment](#)

About Veritas InfoScale solutions in a VMware environment

This chapter includes the following topics:

- [Overview of the Veritas InfoScale Products Virtualization Guide](#)
- [How Veritas InfoScale solutions work in a VMware environment](#)
- [About Veritas InfoScale solutions support for the VMware ESXi environment](#)
- [Virtualization use cases addressed by Veritas InfoScale products](#)

Overview of the Veritas InfoScale Products Virtualization Guide

This document provides information about Veritas InfoScale products support for VMware virtualization technology. It contains:

- High-level conceptual information for Veritas InfoScale products and how they function in ESXi environments.
- High level implementation information for setting up Veritas InfoScale products in ESXi environments.
- Use case chapters with examples of how Veritas InfoScale products can improve performance outcomes for common ESXi use cases.

How Veritas InfoScale solutions work in a VMware environment

Using Veritas InfoScale solutions in a VMware environment means that the Veritas InfoScale product runs in the operating system, inside the Virtual Machine (VM).

The Veritas InfoScale component, such as Storage Foundation, does not run inside the VMware ESXi kernel or in the Hypervisor.

Figure 1-1 shows an example of the high-level architecture diagram with Storage Foundation running in the VM.

Figure 1-1 Architecture overview





Figure 1-2 shows the I/O path from the Virtual Machine to the storage.

Figure 1-2 I/O path from Virtual Machine to storage



VMware has several different methods to allocate block storage to a virtual machine:

- File-based virtual disks created in VMFS or from NFS - Virtual Disk
- Block storage mapped from local disk, Fibre Channel LUNs or iSCSI – Raw Device Mapping

VMware must be configured to use Raw Device Mapping for certain features of Veritas Volume Manager (VxVM) to operate as they do in a physical server environment.

See [“When to use Raw Device Mapping and Storage Foundation”](#) on page 13.

Dynamic Multi-Pathing (DMP) can be used in a Virtual Machine, either as a stand-alone product or as a component of another Veritas InfoScale product. In either case, DMP does not perform multi-pathing in the VM. The VMware architecture presents the Virtual Machine as a single data path and the Hypervisor layer takes care of the multi-pathing. Technically, you could possibly configure the same disk,

with raw device mapping, over two different host bus adapters to a virtual machine. The resulting configuration is not supported because it would run two multi-pathing solutions on top of each other.

Although DMP does not perform multi-pathing in the VM, DMP is an integral part of the data path of Veritas InfoScale products and cannot be disabled. DMP performs device management tasks such as device discovery and thin reclamation.

How Veritas InfoScale product components enhance VMware capabilities


In VMware vSphere, VMFS does not have the capability to mirror storage. This forces users to use mirrored LUN's to provide this functionality to the virtual machines. With Veritas Volume Manager running in the virtual machine, utilizing raw device mapping, data can be protected with mirroring in the virtual machine, including the ability to mirror between storage arrays.

With SFCFSA and Flexible Storage Sharing, you can also mirror storage on local disks directly attached to an ESXi host to local disks directly attached to a remote ESXi host, whether these disks are presented as raw device mappings or VMDK files.

Storage Foundation can make painful migrations of data from physical to virtual environments easier and safer to execute. With Storage Foundation there is no need to actually copy any data from source to destination; rather, the administrator re-assigns the same storage (or a copy of it for a test migration) to the virtual environment. Once the storage is assigned and configured into the virtual machine, Veritas Volume Manager will scan the device tree and discover the disk group and volume structure.

Figure 1-3 describes an example workflow.

Figure 1-3 Migration workflow



Veritas Volume Manager is agnostic to the actual physical device entry. That is, Veritas Volume Manager does not care if the device is `/dev/sdb` or `/dev/sdaz`. This

transparency makes it easy to move storage from one node to another, or between physical and virtual machines.

When to use Raw Device Mapping and Storage Foundation

Raw Device Mapping (RDM) enables a virtual machine to have direct access to the storage rather than going through VMFS. RDM is configured per physical storage device, i.e. a disk or LUN is assigned to one or more virtual machines. It is not possible to assign a part of a physical storage device to a virtual machine. Different types of storage (local SCSI disks, iSCSI disks, Fibre Channel disks) can be used with raw device mapping; Veritas Volume Manager supports all three types of disks.

Note: The Storage Foundation products work well with the iSCSI disks mapped directly to the Virtual Machines.

VMware provides two different modes for raw device mapping:

- Logical mode offers the same functionality and compatibility as a Virtual Disk with respect to VMware ESXi features.
- Physical mode is the most similar method to storage access in a non-virtual environment. Only one SCSI command, REPORT_LUNS, is virtualized as it is required to enable vMotion and a few other features in VMware.
With Storage Foundation, physical mode is recommended as it enables maximum functionality of Veritas Volume Manager in a VMware environment.

The different modes affect the functionality and behavior of Storage Foundation. It is important to use the correct mode for the desired functionality. The benefit of each storage access method is dependent on the workload in the virtual machine. It is easy to get started with one way of deploying storage without considering the long-term implications because of the ease of use of the virtual environment.

For applications with little to no storage need, using raw device mapping is overkill and not recommended. Also, if your environment depends on VMware snapshots, using Raw Device Mapping in physical mode is not possible as it is not supported by VMware.

Raw Device Mapping is a great fit for:

- Applications with large storage needs
- Applications that need predictable and measurable performance
- Multi-node clusters using disk quorums
- Applications with storage that is currently managed by Storage Foundation but is moving into a virtual environment

- Applications that require direct access to storage, such as storage management applications


Array migration

When a virtual machine (VM) is selected for movement of disk files to a new disk storage (DS) using Storage vMotion, the following takes place:

- The VM home directory (config, log, swap, snapshots) is copied to the destination DS.
A "shadow" VM is started on the destination DS using the copied files. The "shadow" VM idles waiting for the copying of the VM disk file(s) to complete.
- An initial copy of the VMs disk file(s) is made to the target DS. During the copy changes made to the source are tracked (change block tracking).
- Storage vMotion iteratively repeats this process of copying the changed blocks from the source DS to the destination DS.
- When the amount of outstanding changed blocks is small enough, vMotion invokes a Fast Suspend and Resume (FSR) of the VM (similar to vMotion) to transfer the running VM over to the idling shadow VM. As is the case with regular vMotion, this transfer normally happens so quickly that it will be completely transparent to the VM.
- After the FSR completes the old home directory the VM disk files are deleted from the source DS.

Figure 1-4 describes the high level process for VMware.


Figure 1-4 Virtual machine data migration with ESXi tools



In VMware, you can add disks to the ESXi server and the virtual machine without rebooting. This functionality makes it possible to offer a better process using Storage Foundation together with raw device mapped storage for online data migration.

Figure 1-5 describes the high level process for Storage Foundation.

Figure 1-5 Virtual machine data migration with Storage Foundation



Data migration for Storage Foundation can be executed either locally in the virtual machine with Veritas Volume Manager or in a central location, migrating all storage from an array utilized by Storage Foundation managed hosts. This powerful, centralized data migration functionality is available in InfoScale Operations Manager.

See the following Web site for information about InfoScale Operations Manager.

<https://www.veritas.com/product/storage-management/infoscale-operations-manager>

Veritas InfoScale component limitations in an ESXi environment

Some limitations apply for using Storage Foundation and Storage Foundation Cluster File System High Availability with VMware ESXi:

- **Dynamic Multi-Pathing**
 Dynamic Multi-Pathing (DMP) is supported in the VMware environment. Due to the architecture of VMware, DMP does not perform multi-pathing for the Virtual Machine. However, DMP is an integral part of the data path of Veritas InfoScale products and performs device management tasks such as device discovery and thin reclamation.
 See “How Veritas InfoScale solutions work in a VMware environment” on page 10.
- **Sharing VMDK files between virtual machines**
 When sharing VMDK files between virtual machines, SCSI BUS sharing mode for the corresponding SCSI controllers can be set to either “Physical” or “Virtual” modes. Setting this mode to “Physical” causes SCSI reservation conflict and I/O error on DMP. This issue occurs in LVM and raw disks also.

Solution:

Disable simultaneous write protection for the shared disk using the multi-writer flag. The procedure is described in the following VMware Knowledge Base article:

<http://kb.vmware.com/kb/1034165>

- Volume Replicator (VVR) option

VVR is supported inside a Virtual Machine. Keep in mind that VVR can use a significant amount of network bandwidth (depends on the amount of data written to disk) and this can reduce the available network bandwidth for other virtual machines. Since the network card is a shared resource, ensure that enough bandwidth is available to sustain the writes happening to disk.
- VMware snapshots

The following limitations apply for VMware snapshots:

 - VMware snapshots are not supported when raw device mapping is used in physical mode, regardless of whether Veritas InfoScale components are installed or not. The REDO-log functionality that is required for VMware Snapshots is not available with raw device mapping used in the physical mode.

Veritas InfoScale components support VMDK file utilization as a backend storage to overcome this limitation.

Veritas InfoScale components also support raw device mapping in the logical mode and VMware snapshots, since RDM-logical mode uses the same level of SCSI virtualization as VMDK files.

See “[When to use Raw Device Mapping and Storage Foundation](#)” on page 13.
 - VMware snapshots are not supported when using VMDK files as backend storage and the multi-writer flag has been set. For example, when using InfoScale Enterprise with VMDK files. This is a VMware limitation.
 - VMware snapshots are not supported with any VM using disks managed by the VMwareDisks agent.
- vMotion (Live Migration)
 - For InfoScale Foundation: vMotion is supported.
 - For InfoScale Availability and Cluster Server (VCS):

VMware vMotion has a limitation that affects all clustering software. vMotion is not supported when a virtual SCSI controller is set to have sharing enabled. Virtual SCSI controller sharing is a virtual machine attribute and is required to be set for the virtual machines that share storage between each other (on the same physical ESXi server or between physical ESXi servers). Essentially all clustering products that rely on SAN storage require this attribute to be set.

VCS provides the VMwareDisks agent to override this limitation and enable shared storage for InfoScale Availability to operate. InfoScale Availability supports VMDK files and therefore vMotion.

See “[About setting up Storage Foundation Cluster File High System High Availability on VMware ESXi](#)” on page 71.

This limitation does not affect the virtual machines that do not have the sharing attribute turned on for their virtual SCSI controllers.

- For InfoScale Enterprise: vMotion is supported.
- N-Port ID Virtualization (NPIV)
NPIV used with InfoScale Foundation is fully supported. No additional setup tasks are required for InfoScale Foundation when the storage is NPIV-enabled. VMware currently does not support I/O fencing with NPIV for any other 3rd party clustering software other than MSCS. In VMware environments, InfoScale products supports I/O fencing using the Coordination Point server as an arbitration mechanism.

I/O fencing considerations in an ESXi environment

VMware does not support SCSI-3 Persistent Reservations (and hence I/O Fencing) with any other 3rd party clustering software with RDM logical mode or VMDK-based virtual disks. In VMware environments, SFHA and SFCFSHA support the following methods of fencing:

- disk-based fencing with RDM-P mode.
Available starting with SFHA and SFCFSHA version 5.1 Service Pack 1 Rolling Patch 1.
See the following tech note for details.
https://www.veritas.com/support/en_US/article.TECH169366
- non-SCSI-3 PR-based fencing using the Coordination Point (CP) server.
The CP server provides arbitration amongst the multiple nodes.

I/O fencing utilizes HBA World Wide Numbers (WWNs) to create registrations on the storage; this has implications in a virtual environment where the HBA is shared between virtual servers on the same physical ESXi host as the WWN used for I/O fencing ends up being the same for each virtual machine. Therefore, SFCFSHA virtual machines (in the same SFCFSHA cluster) cannot share physical servers as the I/O fencing behavior will result in all nodes from that physical ESXi host being fenced out if an event triggers the fencing functionality. In short, if I/O fencing is configured, the SFCFSHA nodes (in the same SFCFSHA cluster) have to be running on separate physical ESXi hosts.

About Veritas InfoScale solutions support for the VMware ESXi environment

Veritas InfoScale components support the VMware ESXi environment as follows:

Table 1-1 Veritas InfoScale solutions support for ESXi virtualization components

Veritas InfoScale components	ESXi host	ESXi guest
Dynamic Multi-Pathing (DMP) for Linux/UNIX	N	Y *
Storage Foundation (SF)	N	Y
Cluster Server (VCS)	N	Y
Storage Foundation and High Availability (SFHA)	N	Y
Storage Foundation Cluster File System High Availability (SFCFSHA)	N	Y
Storage Foundation for Oracle RAC (SF Oracle RAC)	N	Y
Replicator Option	N	Y

Note: * DMP has limited functionality in the guest: multi-pathing works, however it shows only a single path. DMP can encapsulate the disk if mapping the physical disk to the guest.

Virtualization use cases addressed by Veritas InfoScale products

Veritas InfoScale product components support the following VMware environment use cases:

Table 1-2 Virtualization use cases addressed by Veritas InfoScale products in an ESXi environment

Virtualization use case	Veritas solution	Implementation details
Application management and availability	InfoScale Availability or InfoScale Enterprise in guest	How to manage application monitoring and failover on virtual machines using VCS. See “About application availability with Cluster Server (VCS) in the guest” on page 38.
High availability for live migration	InfoScale Availability or InfoScale Enterprise in guest	How to use VCS to provide high availability for live migration with vMotion. See “About VCS support for Live Migration” on page 38.
Storage visibility	InfoScale Foundation or InfoScale Storage or InfoScale Enterprise in ESXi host	How to use DMP features to improve storage visibility. See “Achieving storage visibility using Dynamic Multi-Pathing in the hypervisor” on page 52. See DMP documentation.
Storage availability	InfoScale Foundation or InfoScale Storage or InfoScale Enterprise in ESXi host	How to use DMP features to improve storage availability. See “Achieving storage availability using Dynamic Multi-Pathing in the hypervisor” on page 54. See DMP documentation.
Improved I/O performance	InfoScale Foundation or InfoScale Storage or InfoScale Enterprise in ESXi host	How to use DMP features to improve I/O performance.
Simplified path management with DMP	InfoScale Foundation or InfoScale Storage or InfoScale Enterprise in guest	How to use DMP features for end-to-end storage path visibility, simplified management, and improved performance. See DMP documentation.

Table 1-2 Virtualization use cases addressed by Veritas InfoScale products in an ESXi environment (*continued*)

Virtualization use case	Veritas solution	Implementation details
Data protection	InfoScale Storage or InfoScale Enterprise in guest	<p>How to use Storage Foundation backup and recovery features for data protection.</p> <p>See “Protecting data with Veritas InfoScale product components in the VMware guest” on page 61.</p>
Storage optimization	InfoScale Storage or InfoScale Enterprise in guest	<p>How to use Storage Foundation thin provisioning, FileSnap, SmartTier, and SmartIO features to optimize storage in a VMware environment.</p> <p>See “Optimizing storage with Veritas InfoScale product components in the VMware guest” on page 62.</p>
Data migration	InfoScale Storage or InfoScale Enterprise in guest	<p>How to use Storage Foundation Portable Data Containers to migrate data safely and easily in a VMware guest environment.</p> <p>See “Migrating data with Veritas InfoScale product components in the VMware guest” on page 65.</p>
Improved database performance	InfoScale Storage or InfoScale Enterprise in guest	<p>How to use Storage Foundation database accelerators to improve database performance in a VMware guest environment.</p> <p>See “Improving database performance with Veritas InfoScale product components in the VMware guest” on page 66.</p>
Simplified storage management	InfoScale Foundation or InfoScale Storage or InfoScale Enterprise in guest	<p>How to use DMP for VMware and Storage Foundation features for end-to-end storage visibility, performance, optimization, and enhanced ease of management.</p>

Table 1-2 Virtualization use cases addressed by Veritas InfoScale products in an ESXi environment *(continued)*

Virtualization use case	Veritas solution	Implementation details
Application high availability and fast failover	InfoScale Enterprise in guest	How to manage application high availability and fast failover See “About use cases for InfoScale Enterprise in the VMware guest” on page 68.

Deploying Veritas InfoScale products in a VMware environment

- [Chapter 2. Getting started](#)
- [Chapter 3. Understanding Storage Configuration](#)

Getting started

This chapter includes the following topics:

- [Storage configurations and feature compatibility](#)
- [About setting up VMware with Veritas InfoScale products](#)
- [Veritas InfoScale products support for VMware environments](#)
- [Installing and configuring storage solutions in the VMware virtual environment](#)

Storage configurations and feature compatibility

All block storage topologies that are supported with ESXi are supported when Storage Foundation is running inside a Virtual machine. The storage specific details are hidden for Storage Foundation by VMware hence FC, iSCSI and locally attached disks are supported.

Table 2-1 vSphere features compatible with shared storage configurations

VM storage configurations	Compatible vSphere features			
	vMotion	DRS	VMware HA	VMware Snapshots
VMDK on VMFS	Y	N	Y	N
Physical and Virtual RDM with FC SAN	N	N	N	N
Physical and Virtual RDM with iSCSI SAN	N	N	N	N
VMDK on NFS	N	N	N	N
iSCSI inside guest	Y	Y	Y	Y *

Table 2-1 vSphere features compatible with shared storage configurations
(continued)

VM storage configurations	Compatible vSphere features			
	vMotion	DRS	VMware HA	VMware Snapshots
NFS inside guest (using Mount agent)	Y	Y	Y	Y*

* Taking snapshots is possible if VCS and all applications are fully stopped. Before reverting to the snapshot, shutdown VCS and all applications and then revert to the snapshot. Ensure that no software updates or configuration changes related to VCS have been applied post the snapshot.

Refer to appendix A for more on known limitations before moving to a configuration with VCS in the guest.

About setting up VMware with Veritas InfoScale products

Before setting up your virtual environment, verify that your planned configuration meets the requirements to install Veritas InfoScale products.

Refer to the following information for the system requirements, licensing, and other considerations for installation of Veritas InfoScale product components.

- **Licensing:** Customers running Storage Foundation or Storage Foundation Cluster File System High Availability (SFCFSHA) in a VMware environment are entitled to use an unlimited number of guests on each licensed server or CPU.
- **Veritas InfoScale product requirements:** Each Veritas InfoScale product has system requirements and supported software. Refer to the *Release Notes* for each product.
 In addition, Veritas InfoScale may require specific configuration in the VMware environment.
 See “[Storage configurations and feature compatibility](#)” on page 23.
- **Release Notes:** For each Veritas InfoScale product, the release notes contains last-minute news and important details, including updates to system requirements and supported software. Review the *Release Notes* for the latest information before you start installing the product.
 The product documentation is available on the web at the following location:
<https://sort.veritas.com/documents>

Veritas InfoScale products support for VMware environments

Veritas InfoScale products are a set of components that provide storage administration and management in a heterogeneous storage environment. Veritas supports the Veritas InfoScale product components running within the virtual machine on VMware ESXi.

Veritas tests and supports the Veritas InfoScale product components on the following versions of VMware ESXi:

- VMware ESXi 5.1U1
- VMware ESXi 5.5

The guest operating system can be any of the Linux operating systems that are supported for this release of Veritas InfoScale products. All of the Linux operating systems that are supported by the Veritas InfoScale products in a physical machine are supported in a virtual machine.

For information about the supported operating systems, see the Linux Release Notes for the Veritas InfoScale product version you are using.

See [“About Veritas InfoScale solutions support for the VMware ESXi environment”](#) on page 18.

Installing and configuring storage solutions in the VMware virtual environment

To set up a guest in VMware environment with Veritas InfoScale products, install and configure the Veritas InfoScale product component on the required virtual machines.

Refer to the installation guide for your Veritas InfoScale product for installation and configuration information.

See [“Veritas InfoScale documentation”](#) on page 87.

Understanding Storage Configuration

This chapter includes the following topics:

- [Configuring storage](#)
- [Enabling disk UUID on virtual machines](#)
- [Installing Array Support Library \(ASL\) for VMDK on cluster nodes](#)
- [Excluding the boot disk from the Volume Manager configuration](#)
- [Creating the VMDK files](#)
- [Mapping the VMDKs to each virtual machine \(VM\)](#)
- [Enabling the multi-write flag](#)
- [Getting consistent names across nodes](#)
- [Creating a file system](#)

Configuring storage

There are two options to provide storage to the Virtual Machines (VMs) that will host the Cluster File System:

- The first option, Raw Device Mapping Protocol (RDMP), uses direct access to external storage and supports parallel access to the LUN, but does not allow vMotion or DRS. For RDMP configuration, you must map the raw device to each VM and make sure you select the Physical (RDM-P) configuration, so SCSI-3 PGR commands are passed along to the disk.

- The second option, VMFS virtual disk (VMDK), provides a file that can only be accessed in parallel when the VMFS multi-writer option is enabled. This option supports server vMotion and DRS, but does not currently support SCSI-3 PR IO fencing. The main advantage of this architecture is the ability to move VMs around different ESXi servers without service interruption, using vMotion.

This deployment example uses VMDK files with the multi-writer option enabled. In this section we will show how to configure the ESXi server and virtual machines to share a VMDK file and how to configure SFCFSHA to consume that storage and create a file system. Support for VMDK files is based on the multi-writer option described in this VMware article: <http://kb.vmware.com/kb/1034165> By default, one VMDK file can only be mounted by one VM at a time. By following the steps in the VMware article, simultaneous write protection provided by VMFS is disabled using the multi-writer flag. When choosing this configuration, users should be aware of the following limitations and advantages.

Limitations:

- Virtual disks must be eager zeroed thick
- VMDK Sharing is limited to eight ESXi servers
- Linked clones and snapshots are not supported. Be aware that other vSphere activities utilize cloning and that backup solutions leverage snapshots via the vAPIs, so backups may be adversely impacted.
- SCSI-3 PR IO fencing is not supported by VMDK files. Special care needs to be taken when assigning VMDKs to VMs. Inadvertently assigning a VMDK file already in use to the wrong VM will likely result in data corruption.
- Storage vMotion is not supported

The advantage is that server vMotion is supported.

The lack of SCSI-3 PR IO fencing support requires the usage of at least three Coordination Point servers, to provide non-SCSI-3 fencing protection. In case of a split brain situation, CP servers will be used to determine what part of the sub-cluster will continue providing service. Once the multi-writer flag is enabled on a VMDK file, any VM will be able to mount it and write, so special care in the provisioning phase needs to be taken.

Note that if the number of SFCFSHA nodes is greater than eight, several nodes will have to run in the same ESXi server, based on the limitation that a maximum of eight ESXi servers can share the same VMDK file. For example, if you are running at the SFCFSHA maximum of 64 nodes, those 64 VMs would share the same VMDK file, but you could only use eight ESXi servers to host the cluster.

These are the steps that need to be taken when configuring VMDKs as shared backed storage and that will be presented in the next sections:

Table 3-1 Steps to configure VMDK

Storage deployment task	Deployment steps
Enabling Disk UUID on virtual machines (VMs)	See “Enabling disk UUID on virtual machines” on page 28.
Installing Array Support Library (ASL) for VMDK on cluster nodes	See “Installing Array Support Library (ASL) for VMDK on cluster nodes” on page 29.
Excluding the boot disk from the Volume Manager configuration	See “Excluding the boot disk from the Volume Manager configuration” on page 30.
Creating the VMDK files	See “Excluding the boot disk from the Volume Manager configuration” on page 30.
Mapping VMDKs to each virtual machine	See “Mapping the VMDKs to each virtual machine (VM)” on page 31.
Enabling the multi-write flag	See “Enabling the multi-write flag” on page 33.
Getting consistent names across nodes	See “Getting consistent names across nodes” on page 34.
Creating a Cluster File System	See “Creating a file system” on page 35.

Enabling disk UUID on virtual machines

You must set the `disk.EnableUUID` parameter for each VM to “TRUE”, if you plan to setup a high availability configuration. This step is necessary so that the VMDK always presents a consistent UUID to the VM, thus allowing the disk to be mounted properly. For each of the virtual machine nodes (VMs) that will be participating in the cluster, follow the steps below from the vSphere client:

To enable disk UUID on a virtual machine

- 1 Power off the guest.
- 2 Select the guest and select **Edit Settings**.
- 3 Select the **Options** tab on top.
- 4 Select **General** under the **Advanced** section.
- 5 Select the **Configuration Parameters...** on right hand side.

- 6 Check to see if the parameter **disk.EnableUUID** is set, if it is there then make sure it is set to **TRUE**.
 If the parameter is not there, select **Add Row** and add it.
- 7 Power on the guest.

Installing Array Support Library (ASL) for VMDK on cluster nodes

In order for the cluster file system to work properly with the VMDK files, an ASL must be installed in the virtual server. The ASL package (VRTSaslapm) version that contains the VMDK ASL is 6.0.100.100.

Note: Any future updates to the VMDK ASL will be published in <http://sort.veritas.com> and will have a higher revision than 6.0.100.100.

To download the ASL package

- 1 Go to <http://sort.veritas.com>.
- 2 Under **Downloads**, click **ASL/APM/DDI/DDL**.
- 3 Select the ASL package by appropriately filtering the array information.
- 4 Click **Download** in the **Array Support Library (ASL)/Array Policy Module (APM) Details** page.

On each clustered file system node, perform the following steps. The steps are illustrated with the example installation and your details may vary.

To install the ASL package

- 1 To install the package, follow the instructions outlined in the Readme file (VRTSaslap_readme.txt) which is displayed towards the end of the **Array Support Library (ASL)/Array Policy Module (APM) Details** page.
 You can also save this Readme file by selecting the **Save As...** option.
- 2 Follow the steps outlined in the *Installation Procedure* section of the Readme.

In the case of the example installation, the version for RHEL 6.

After installing the ASL, you will notice that the disk has been renamed from disk_0 to vmdk0_0. Before ASL:

```
# vxdisk list
DEVICE          TYPE          DISK          GROUP          STATUS
disk_0          auto:none    -             - online       invalid
```

After ASL has been deployed:

```
# vxdisk list
DEVICE          TYPE          DISK          GROUP          STATUS
vmdk0_0         auto:none     -             - online       invalid
```

vmdk0_0 is the boot disk that is to be excluded from Volume Manager configuration.

Excluding the boot disk from the Volume Manager configuration

It is a best practice to exclude the boot disk from Volume Manager. This allows the shared VMDK files to be configured to use the same name. In order to exclude the disk, run the command `vxdmpadm` with the name of the boot disk. In the example installation:

```
[root@cfs01 RHEL6]# vxdmpadm exclude dmpnodename=vmdk0_0
```

Verify that the boot disk is no longer reported under the VxVM configuration:

```
[root@cfs01 RHEL6]# vxdisk list
DEVICE          TYPE          DISK          GROUP          STATUS
[root@cfs01 RHEL6]#
```

Creating the VMDK files

The VMDKs that will be used by Storage Foundation Cluster File System High Availability (SFCFSHA) can be created either by the vSphere GUI or using the command line. Using the GUI, there is no control for the name of the file used, and they will be stored under the folder belonging to the VM that is creating the files. We would prefer in this case to control those file names, so we will use the command line to create the following configuration:

Table 3-2 Virtual disk configuration

Data Store	Virtual Disk on ESXi	VMDK NAME	Virtual device	SCSI Driver	Virtual size (GB)
DS1	Hard disk 2	cfs0/shared1.vmdk	SCSI 1:0	Paravirtual	90
DS2	Hard disk 3	cfs0/shared2.vmdk	SCSI 1:1	Paravirtual	90
DS3	Hard disk 4	cfs0/shared3.vmdk	SCSI 1:2	Paravirtual	90
DS4	Hard disk 5	cfs0/shared4.vmdk	SCSI 1:3	Paravirtual	90

Table 3-2 Virtual disk configuration (*continued*)

Data Store	Virtual Disk on ESXi	VMDK NAME	Virtual device	SCSI Driver	Virtual size (GB)
DS5	Hard disk 6	cfs0/shared5.vmdk	SCSI 1:4	Paravirtual	90

To create the infrastructure

- 1** Connect to one of the ESXi virtual machines.
- 2** Create a folder called cfs0 (the name of the cluster) in each of the datastores:

```
mkdir /vmfs/volumes/DS1/cfs0
mkdir /vmfs/volumes/DS2/cfs0
mkdir /vmfs/volumes/DS3/cfs0
mkdir /vmfs/volumes/DS4/cfs0
mkdir /vmfs/volumes/DS5/cfs0
```

- 3** Create each of the VMDKs that will be used:

```
vmkfstools -c 90G -d eagerzeroedthick
/vmfs/volumes/DS1/cfs0/shared1.vmdk
vmkfstools -c 90G -d eagerzeroedthick
/vmfs/volumes/DS2/cfs0/shared2.vmdk
vmkfstools -c 90G -d eagerzeroedthick
/vmfs/volumes/DS3/cfs0/shared3.vmdk
vmkfstools -c 90G -d eagerzeroedthick
/vmfs/volumes/DS4/cfs0/shared4.vmdk
vmkfstools -c 90G -d eagerzeroedthick
/vmfs/volumes/DS5/cfs0/shared5.vmdk
```

Mapping the VMDKs to each virtual machine (VM)

Map each of the created VMDK files to each VM. The example procedure illustrates mapping the VMDKs to the cfs01 node: all steps should be followed for each of the other nodes.

To map the VMDKs to each VM

- 1** Shut down the VM.
- 2** Select the VM and select **Edit Settings...**
- 3** Select **Add** , select **Hard disk** and click **Next**.
- 4** Select **Use an existing virtual disk** and click **Next**.
- 5** Select **Browse** and choose DS1 data store.
- 6** Select the folder **cfs0** and select **shared1.vmdk file** and click **Next**.
- 7** On **Virtual Device Node** select **SCSI (1:0)** and click **Next**.
- 8** Review the details to verify they are correct and click **Finish**.
- 9** Since this is the first disk added under SCSI controller 1, a new SCSI controller is added.

Modify the type to **Paravirtual**, if that is not the default, and check that **SCSI Bus Sharing** is set to **None**, as this is key to allow vMotion for the VMs.

- 10** Follow steps 3 to 8 for the rest of disks that will be added to each of the VMs.

For the example configuration, the parameters for steps 5-7 are given in the table below:

Data Store	VMDK Name	Virtual Device
DS1	cfs0/shared1.vmdk	SCSI 1:0
DS2	cfs0/shared2.vmdk	SCSI 1:1
DS3	cfs0/shared3.vmdk	SCSI 1:2
DS4	cfs0/shared4.vmdk	SCSI 1:3
DS5	cfs0/shared5.vmdk	SCSI 1:4

The final configuration for the first node of the example cluster (cfs01):

Now follow the same steps for each node of the cluster and map each VMDK file to the VM following the instructions above. Once all the steps are completed, all the VMs should have access to the same VMDK files. Note that at this point, all the VMs are still powered off and that multi-writer flag has not been enabled yet (it will be done in the next step). Any attempt to power on the VMs in this state will prevent a second VM start because it will violate the restrictions to access a VMDK by only a host at a time.

Enabling the multi-write flag

Enable the multi-write flag if you plan to setup an SFCFSHA configuration. For detailed instructions on how to enable the multi-writer flag, see the steps in the following VMware article:

<http://kb.vmware.com/kb/1034165>

The steps given below illustrate the example where five VMDK files are configured and are shared by four virtual machines (VMs). These VMs constitute the four nodes of the cluster and they are powered off. Now it is time to enable the multi-writer flag for each of the VMs.

To enable the multi-write flag for a virtual machine

- 1 On the vSphere Client, right-click on the cfs01 virtual machine. Go to **Edit Settings > Options > Advanced > General > Configuration Parameters....**
- 2 Select **Add Row**.
- 3 Enter **scsi1:0.sharing** on the **Name** column.
- 4 Enter **multi-writer** on the **Value** column.
- 5 Repeat steps 2 through 4 and enter the multi-writer value for the rest of the SCSI controllers and targets. For the example configuration:

```
scsil:1.sharing multi-writer
scsil:2.sharing multi-writer
scsil:3.sharing multi-writer
scsil:4.sharing multi-writer
```

Once those steps are done, the VM configuration will resemble:

disk.EnableUUID	true
scsi1:0.sharing	multi-writer
scsi1:1.sharing	multi-writer
scsi1:2.sharing	multi-writer
scsi1:3.sharing	multi-writer
scsi1:4.sharing	multi-writer

- 6 Press **OK** to confirm.

- 7 Repeat steps 1 to 6 for the other virtual machines (cfs02, cfs03 and cfs04 in the example configuration).
- 8 Once all the virtual machines have been configured correctly, power them on and verify that there are no issues. Note that the disks have been added to each of the hosts.

Example configuration for cfs01:

```
# vxdisk list
DEVICE          TYPE          DISK          GROUP          STATUS
vmdk0_0         auto:none     -             - online       invalid
vmdk0_1         auto:none     -             - online       invalid
vmdk0_2         auto:none     -             - online       invalid
vmdk0_3         auto:none     -             - online       invalid
vmdk0_4         auto:none     -             - online       invalid
vmdk0_5         auto:none     -             - online       invalid
[root@cfs01 ~]#
```

Getting consistent names across nodes

It is likely that the VMDK files are presented in a different order on each system and that the names given by Volume Manager may vary. The recommended best practice for a consistent deployment is to rename the disk so the configuration is clear.

As an example of the initial discrepancies between cfs01 and cfs03, cfs01 the disk name associated to device ending on serial number 226 is vmdk0_5:

```
[root@cfs01 ~]# /etc/vx/bin/vxgetdmpnames
enclosure vendor=VMware product=disk serial=vmdk name=vmdk0
  dmpnode serial=6000c2993a8d6030ddf71042d4620cec name=vmdk0_1
  dmpnode serial=6000c29ac083abd0a86fa46b509d69f5 name=vmdk0_2
  dmpnode serial=6000c29e13f6aff58ac3d543b022dfe2 name=vmdk0_3
  dmpnode serial=6000c29f2a8d6030ddf71042d4620cec name=vmdk0_4
  dmpnode serial=6000c2993a8d6030ddf71042d4620cec name=vmdk0_5
```

Observe how cfs03 named the same device vmdk_0_0:

```
[root@cfs01 ~]# /etc/vx/bin/vxgetdmpnames
enclosure vendor=VMware product=disk serial=vmdk name=vmdk0
  dmpnode serial=6000c2993a8d6030ddf71042d4620cec name=vmdk0_1
  dmpnode serial=6000c29ac083abd0a86fa46b509d69f5 name=vmdk0_2
  dmpnode serial=6000c29e13f6aff58ac3d543b022dfe2 name=vmdk0_3
```

```
dmpnode serial=6000c29f2a8d6030ddf71042d4620cec name=vmdk0_4  
dmpnode serial=6000c2993a8d6030ddf71042d4620cec name=vmdk0_5
```

In order to get the same names across all the cluster nodes the command `vxddladm` is used. For each node of the cluster, run the command:

```
# vxddladm assign names
```

Observe now how `cfs03` got the right name for device ending at 226 serial number:

```
[root@cfs01 ~]# /etc/vx/bin/vxgetdmpnames  
enclosure vendor=VMware product=disk serial=vmdk name=vmdk0  
dmpnode serial=6000c2993a8d6030ddf71042d4620cec name=vmdk0_1  
dmpnode serial=6000c29ac083abd0a86fa46b509d69f5 name=vmdk0_2  
dmpnode serial=6000c29e13f6aff58ac3d543b022dfe2 name=vmdk0_3  
dmpnode serial=6000c29f2a8d6030ddf71042d4620cec name=vmdk0_4  
dmpnode serial=6000c2993a8d6030ddf71042d4620cec name=vmdk0_5
```

Creating a file system

The next step will be to configure a common mount point across all the nodes, mounted on the same storage. In order to simplify the examples given here, a single disk group containing all the disks and a single volume will be created. Depending on the application requirements the number of disk groups and volumes may vary.

The boot disk has been excluded from Volume Manger configuration, so the 5 available disks (`vmdk0_1`, `vmdk0_2`, `vmdk0_3`, `vmdk0_4` and `vmdk0_5`) will be the ones added to the disk group. These are the steps:

To create a clustered file system

- 1 Initialize the disks:
- 2 Create a new disk group and add the disks.
- 3 Verify the configuration. Note the DISK and GROUP information.
- 4 Create a striped volume with the 5 disks available.
- 5 Create a File System.
- 6 If you plan to configure a clustered file system environment, add the newly created file system to the cluster configuration. Given that this will be mounted by all the nodes at the same time, we will add it as a cluster resource, and commands `cfsmntadm` and `cfsmount` will be used.

- 7** In case of a clustered file system, verify that the new directory is available in all the nodes by running the `cfscluster status` command or by verifying with `df` in each of the nodes.

Use cases for Veritas InfoScale product components in a VMware environment

- [Chapter 4. Application availability using Cluster Server](#)
- [Chapter 5. Multi-tier business service support](#)
- [Chapter 6. Improving storage visibility, availability, and I/O performance using Dynamic Multi-Pathing](#)
- [Chapter 7. Improving data protection, storage optimization, data migration, and database performance](#)
- [Chapter 8. Setting up virtual machines for fast failover using Storage Foundation Cluster File System High Availability on VMware disks](#)

Application availability using Cluster Server

This chapter includes the following topics:

- [About application availability with Cluster Server \(VCS\) in the guest](#)
- [About VCS support for Live Migration](#)

About application availability with Cluster Server (VCS) in the guest

Using VCS virtual-to-virtual or in-guest clustering in a VMware environments provides high availability of applications inside the guest by providing protection from host failures, hardware failures, OS crashes and also application failures at software layer. For example, in cases of application hang, file-level corruption at the OS level cannot be resolved with a reboot.

Since there is a cost involved in maintaining standby virtual machines (VMs), you may choose to protect only specific applications using VCS in-guest and protect the remaining applications using VMware HA. By using VMware HA in conjunction with VCS in the guest, when a host fails, standby VCS nodes running on that host are automatically restarted by VMware HA on a new host without the need for user-intervention, potentially eliminating the need to maintain multiple standbys.

About VCS support for Live Migration

VCS in-guest clustering continues to provide high availability of applications on virtual machines, in live migration scenarios initiated by the virtualization technology. You can use Live migration to perform a stateful migration of a virtual machine in a VCS environment. During this period, you may see notifications if the migrating

node is unable to heartbeat with its peers within LLT's default peer inactive timeout. To avoid false failovers, determine how long the migrating node is unresponsive in your environment. If that time is less than the default LLT peer inactive timeout of 16 seconds, VCS operates normally. If not, increase the peer inactive timeout to an appropriate value on all the nodes in the cluster before beginning the migration. Reset the value back to the default after the migration is complete.

Multi-tier business service support

This chapter includes the following topics:

- [About Virtual Business Services](#)
- [Sample virtual business service configuration](#)

About Virtual Business Services

The Virtual Business Services feature provides visualization, orchestration, and reduced frequency and duration of service disruptions for multi-tier business applications running on heterogeneous operating systems and virtualization technologies. A virtual business service represents the multi-tier application as a consolidated entity that helps you manage operations for a business service. It builds on the high availability and disaster recovery provided for the individual tiers by Veritas InfoScale products such as Cluster Server.

Application components that are managed by Cluster Server or Microsoft Failover Clustering can be actively managed through a virtual business service.


You can use the Veritas InfoScale Operations Manager Management Server console to create, configure, and manage virtual business services.

Sample virtual business service configuration

This section provides a sample virtual business service configuration comprising a multi-tier application. [Figure 5-1](#) shows a Finance application that is dependent on components that run on three different operating systems and on three different clusters.

- Databases such as Oracle running on Solaris operating systems form the database tier.
- Middleware applications such as WebSphere running on AIX operating systems form the middle tier.
- Web applications such as Apache and IIS running on Windows and Linux virtual machines form the Web tier.
 Each tier can have its own high availability mechanism. For example, you can use Cluster Server for the databases and middleware applications for the Web servers.

Figure 5-1 Sample virtual business service configuration



Each time you start the Finance business application, typically you need to bring the components online in the following order – Oracle database, WebSphere, Apache and IIS. In addition, you must bring the virtual machines online before you start the Web tier. To stop the Finance application, you must take the components offline in the reverse order. From the business perspective, the Finance service is unavailable if any of the tiers becomes unavailable.

When you configure the Finance application as a virtual business service, you can specify that the Oracle database must start first, followed by WebSphere and the Web servers. The reverse order automatically applies when you stop the virtual business service. When you start or stop the virtual business service, the components of the service are started or stopped in the defined order.

For more information about Virtual Business Services, refer to the *Virtual Business Service–Availability User’s Guide*.

Improving storage visibility, availability, and I/O performance using Dynamic Multi-Pathing

This chapter includes the following topics:

- [Use cases for Dynamic Multi-Pathing \(DMP\) in the VMware environment](#)
- [How DMP works](#)
- [Achieving storage visibility using Dynamic Multi-Pathing in the hypervisor](#)
- [Achieving storage availability using Dynamic Multi-Pathing in the hypervisor](#)
- [Improving I/O performance with Dynamic Multi-Pathing in the hypervisor](#)
- [Achieving simplified management using Dynamic Multi-Pathing in the hypervisor and guest](#)

Use cases for Dynamic Multi-Pathing (DMP) in the VMware environment

Using DMP in the VMware environment enables the following use cases:

Table 6-1 DMP support for VMware environment use cases

Virtualization use case	Veritas solution	Implementation details
Storage visibility	Dynamic Multi-Pathing (DMP) for VMware in the ESXi hypervisor	How to use DMP features to improve storage visibility. See “Achieving storage visibility using Dynamic Multi-Pathing in the hypervisor” on page 52.
Storage availability	DMP in the ESXi hypervisor	How to use DMP features to improve storage availability. See “Achieving storage availability using Dynamic Multi-Pathing in the hypervisor” on page 54.
Improved I/O performance	DMP in the ESXi hypervisor	How to use DMP features to improve I/O performance.
Simplified management	Dynamic Multi-Pathing (DMP) in the ESXi hypervisor and in the VMware guest	How to use DMP features for end-to-end storage visibility, simplified management, and improved performance.

How DMP works

Dynamic Multi-Pathing (DMP) provides greater availability, reliability, and performance by using the path failover feature and the load balancing feature. These features are available for multiported disk arrays from various vendors.

Disk arrays can be connected to host systems through multiple paths. To detect the various paths to a disk, DMP uses a mechanism that is specific to each supported array. DMP can also differentiate between different enclosures of a supported array that are connected to the same host system.

The multi-pathing policy that DMP uses depends on the characteristics of the disk array.

DMP supports the following standard array types:

Table 6-2

Array type	Description
Active/Active (A/A)	Allows several paths to be used concurrently for I/O. Such arrays allow DMP to provide greater I/O throughput by balancing the I/O load uniformly across the multiple paths to the LUNs. In the event that one path fails, DMP automatically routes I/O over the other available paths.
Asymmetric Active/Active (A/A-A)	A/A-A or Asymmetric Active/Active arrays can be accessed through secondary storage paths with little performance degradation. The behavior is similar to ALUA, except that it does not support the SCSI commands that an ALUA array supports.
Asymmetric Logical Unit Access (ALUA)	DMP supports all variants of ALUA.
Active/Passive (A/P)	<p>Allows access to its LUNs (logical units; real disks or virtual disks created using hardware) via the primary (active) path on a single controller (also known as an access port or a storage processor) during normal operation.</p> <p>In implicit failover mode (or autotrespass mode), an A/P array automatically fails over by scheduling I/O to the secondary (passive) path on a separate controller if the primary path fails. This passive port is not used for I/O until the active port fails. In A/P arrays, path failover can occur for a single LUN if I/O fails on the primary path.</p> <p>This array mode supports concurrent I/O and load balancing by having multiple primary paths into a controller. This functionality is provided by a controller with multiple ports, or by the insertion of a SAN switch between an array and a controller. Failover to the secondary (passive) path occurs only if all the active primary paths fail.</p>

Table 6-2 (continued)


Array type	Description
Active/Passive in explicit failover mode or non-autotrespass mode (A/PF)	<p>The appropriate command must be issued to the array to make the LUNs fail over to the secondary path.</p> <p>This array mode supports concurrent I/O and load balancing by having multiple primary paths into a controller. This functionality is provided by a controller with multiple ports, or by the insertion of a SAN switch between an array and a controller. Failover to the secondary (passive) path occurs only if all the active primary paths fail.</p>
Active/Passive with LUN group failover (A/PG)	<p>For Active/Passive arrays with LUN group failover (A/PG arrays), a group of LUNs that are connected through a controller is treated as a single failover entity. Unlike A/P arrays, failover occurs at the controller level, and not for individual LUNs. The primary controller and the secondary controller are each connected to a separate group of LUNs. If a single LUN in the primary controller's LUN group fails, all LUNs in that group fail over to the secondary controller.</p> <p>This array mode supports concurrent I/O and load balancing by having multiple primary paths into a controller. This functionality is provided by a controller with multiple ports, or by the insertion of a SAN switch between an array and a controller. Failover to the secondary (passive) path occurs only if all the active primary paths fail.</p>

An array policy module (APM) may define array types to DMP in addition to the standard types for the arrays that it supports.

Veritas InfoScale uses DMP metanodes (DMP nodes) to access disk devices connected to the system. For each disk in a supported array, DMP maps one node to the set of paths that are connected to the disk. Additionally, DMP associates the appropriate multi-pathing policy for the disk array with the node.

Figure 6-1 shows how DMP sets up a node for a disk in a supported disk array.


Figure 6-1 How DMP represents multiple physical paths to a disk as one node



DMP implements a disk device naming scheme that allows you to recognize to which array a disk belongs.

Figure 6-2 shows an example where two paths, `vmhba1:C0:T0:L0` and `vmhba2:C0:T0:L0`, exist to a single disk in the enclosure, but VxVM uses the single DMP node, `enc0_0`, to access it.

Figure 6-2 Example of multi-pathing for a disk enclosure in a SAN environment



How DMP monitors I/O on paths

DMP maintains a pool of kernel threads that are used to perform such tasks as error processing, path restoration, statistics collection, and SCSI request callbacks.

One kernel thread responds to I/O failures on a path by initiating a probe of the host bus adapter (HBA) that corresponds to the path. Another thread then takes the appropriate action according to the response from the HBA. The action taken can be to retry the I/O request on the path, or to fail the path and reschedule the I/O on an alternate path.

The restore kernel task is woken periodically (by default, every 5 minutes) to check the health of the paths, and to resume I/O on paths that have been restored. As some paths may suffer from intermittent failure, I/O is only resumed on a path if the path has remained healthy for a given period of time (by default, 5 minutes). DMP can be configured with different policies for checking the paths.

The statistics-gathering task records the start and end time of each I/O request, and the number of I/O failures and retries on each path. DMP can be configured to use this information to prevent the SCSI driver being flooded by I/O requests. This feature is known as I/O throttling.

See [“Path failover mechanism”](#) on page 48.

See [“I/O throttling”](#) on page 48.

Path failover mechanism

DMP enhances system availability when used with disk arrays having multiple paths. In the event of the loss of a path to a disk array, DMP automatically selects the next available path for I/O requests without intervention from the administrator.

DMP is also informed when a connection is repaired or restored, and when you add or remove devices after the system has been fully booted (provided that the operating system recognizes the devices correctly).

If required, the response of DMP to I/O failure on a path can be tuned for the paths to individual arrays. DMP can be configured to time out an I/O request either after a given period of time has elapsed without the request succeeding, or after a given number of retries on a path have failed.

I/O throttling

If I/O throttling is enabled, and the number of outstanding I/O requests builds up on a path that has become less responsive, DMP can be configured to prevent new I/O requests being sent on the path either when the number of outstanding I/O requests has reached a given value, or a given time has elapsed since the last successful I/O request on the path. While throttling is applied to a path, the new I/O

requests on that path are scheduled on other available paths. The throttling is removed from the path if the HBA reports no error on the path, or if an outstanding I/O request on the path succeeds.

Subpaths Failover Group (SFG)

A subpaths failover group (SFG) represents a group of paths which could fail and restore together. When an I/O error is encountered on a path in an SFG, DMP does proactive path probing on the other paths of that SFG as well. This behavior adds greatly to the performance of path failover thus improving I/O performance. Currently the criteria followed by DMP to form the subpaths failover groups is to bundle the paths with the same endpoints from the host to the array into one logical storage failover group.

Low Impact Path Probing (LIPP)

The restore daemon in DMP keeps probing the LUN paths periodically. This behavior helps DMP to keep the path states up-to-date even when no I/O occurs on a path. Low Impact Path Probing adds logic to the restore daemon to optimize the number of the probes performed while the path status is being updated by the restore daemon. This optimization is achieved with the help of the logical subpaths failover groups. With LIPP logic in place, DMP probes only a limited number of paths within a subpaths failover group (SFG), instead of probing all the paths in an SFG. Based on these probe results, DMP determines the states of all the paths in that SFG.

Load balancing

By default, DMP uses the Minimum Queue I/O policy for load balancing across paths for all array types. Load balancing maximizes I/O throughput by using the total bandwidth of all available paths. I/O is sent down the path that has the minimum outstanding I/Os.

For Active/Passive (A/P) disk arrays, I/O is sent down the primary paths. If all of the primary paths fail, I/O is switched over to the available secondary paths. As the continuous transfer of ownership of LUNs from one controller to another results in severe I/O slowdown, load balancing across primary and secondary paths is not performed for A/P disk arrays unless they support concurrent I/O.

For other arrays, load balancing is performed across all the currently active paths.

You can change the I/O policy for the paths to an enclosure or disk array. This operation is an online operation that does not impact the server or require any downtime.

About DMP I/O policies

The DMP I/O policy indicates how DMP distributes I/O loads across multiple paths to a disk array or enclosure. You can set the I/O policy for an enclosure (for example, `HDS01`), for all enclosures of a particular type (such as `HDS`), or for all enclosures of a particular array type (such as `A/A` for Active/Active, or `A/P` for Active/Passive).

The following policies may be set:

`adaptive`

This policy attempts to maximize overall I/O throughput to or from the disks by dynamically scheduling I/O on the paths. It is suggested for use where I/O loads can vary over time. For example, I/O to or from a database may exhibit both long transfers (table scans) and short transfers (random look ups). The policy is also useful for a SAN environment where different paths may have different number of hops. No further configuration is possible as this policy is automatically managed by DMP.

`adaptiveminq`

Similar to the `adaptive` policy, except that I/O is scheduled according to the length of the I/O queue on each path. The path with the shortest queue is assigned the highest priority.

`balanced`
`[partitionsize=size]`

This policy is designed to optimize the use of caching in disk drives and RAID controllers. The size of the cache typically ranges from 120KB to 500KB or more, depending on the characteristics of the particular hardware. During normal operation, the disks (or LUNs) are logically divided into a number of regions (or partitions), and I/O from/to a given region is sent on only one of the active paths. Should that path fail, the workload is automatically redistributed across the remaining paths.

You can use the `size` argument to the `partitionsize` attribute to specify the partition size. The partition size in blocks is adjustable in powers of 2 from 2 up to 231. A value that is not a power of 2 is silently rounded down to the nearest acceptable value.

Specifying a partition size of 0 is equivalent to specifying the default partition size.

The default value for the partition size is 512 blocks (256k). Specifying a partition size of 0 is equivalent to the default partition size of 512 blocks (256k).

The default value can be changed by adjusting the value of the `dmp_pathswitch_blks_shift` tunable parameter.

Note: The benefit of this policy is lost if the value is set larger than the cache size.

For example, the suggested partition size for an Hitachi HDS 9960 A/A array is from 32,768 to 131,072 blocks (16MB to 64MB) for an I/O activity pattern that consists mostly of sequential reads or writes.

`minimumq`

This policy sends I/O on paths that have the minimum number of outstanding I/O requests in the queue for a LUN. No further configuration is possible as DMP automatically determines the path with the shortest queue.

This is the default I/O policy for all arrays.

`priority`

This policy is useful when the paths in a SAN have unequal performance, and you want to enforce load balancing manually. You can assign priorities to each path based on your knowledge of the configuration and performance characteristics of the available paths, and of other aspects of your system.

<code>round-robin</code>	This policy shares I/O equally between the paths in a round-robin sequence. For example, if there are three paths, the first I/O request would use one path, the second would use a different path, the third would be sent down the remaining path, the fourth would go down the first path, and so on. No further configuration is possible as this policy is automatically managed by DMP.
<code>singleactive</code>	This policy routes I/O down the single active path. This policy can be configured for A/P arrays with one active path per controller, where the other paths are used in case of failover. If configured for A/A arrays, there is no load balancing across the paths, and the alternate paths are only used to provide high availability (HA). If the current active path fails, I/O is switched to an alternate active path. No further configuration is possible as the single active path is selected by DMP.

Achieving storage visibility using Dynamic Multi-Pathing in the hypervisor

When DMP is installed in the hypervisor, DMP provides instant visibility to the storage attributes of a LUN. You can use the storage attributes to determine the characteristics of the LUN.

The following scenario describes a typical use case for storage visibility.

As the ESX administrator, you notice that your existing pool of storage is getting near capacity, so you put in a requisition to your storage team to add additional storage. You request two LUNs to meet the storage needs of different tiers of applications:

- 500GB RAID 5 LUN
- 500GB RAID 0+1 LUN


After the storage team processes the request, you get an email from the storage admin stating "Your storage is ready". Without any more information, you need to guess which LUN is which, and has the appropriate storage configuration (RAID 5 vs. RAID 0+1).

You need to determine which LUN is which, so that you can set up the applications on the appropriate storage. If you allocate mission critical, high performance applications to a RAID 5 LUN, SLA's could be impacted, and the applications may run slower than needed. Conversely, you want to avoid allocating low-priority applications on the high-end storage tier.

DMP provides instant visibility to the storage attributes of a LUN so you can determine exactly which device is which.


Navigate to the storage view, and you can see the LUNs. The attributes for each LUN show which LUN is RAID 5 and which is RAID 0+1.

Figure 6-3 Storage view showing the LUN RAID configuration



The storage attributes also let you identify the LUN using the AVID name. The AVID name provides a clear way to refer to the device, and can reduce confusion between the ESXi system administrator and the storage administrator.

Figure 6-4 Storage attributes showing AVID name



If an issue occurs for the LUN, you can refer to the storage device by the AVID name. For example, if one of the LUNs has a performance issue, you can communicate that to the platform or storage administration team. The AVID name provides a common, concise language between the platform and storage administration teams, and minimizes any confusion. It also speeds time to remediation of issues, and reduces ambiguity in communications.

Achieving storage availability using Dynamic Multi-Pathing in the hypervisor

When DMP is installed in the hypervisor, DMP allows you to easily provision new storage to ESXi servers. DMP ensures that all available paths to the storage are used, without any additional server-side configuration. The template-based configuration can ensure that the same settings and tuning is applied to each ESX server in an environment.


The following scenario describes a typical use case for storage availability.

You have a storage array with two array side controllers, and two HBA controllers on the host. In a typical SAN configuration, this configuration results in 4 paths to the back-end storage array, and ensures that the loss of an HBA, SAN switch, or array controller does not result in the loss of access to the storage array.

Without DMP, VMware defaults to use the Fixed I/O policy, which usually uses a single path for I/O regardless of the number of paths to the back-end storage array. To change this default behavior, the VMware administrator must change the Path Selection policy on all LUNs on a given host, which is a time consuming process.


This process must be performed any time new storage LUNs are provisioned to the host.

Figure 6-5 Example of VMware native multi-pathing



By contrast, DMP determines all the paths that may be used for I/O, and makes use of those paths without administrator intervention. As new storage LUNs are added, the administrator does not need to remember to make any changes to the multi-pathing configuration, because the new LUNs under DMP control automatically inherit the same configuration characteristics as the other LUNs in use from the same storage enclosure. This feature reduces the opportunity for mis-configuration, and helps streamline the storage provisioning process.

Figure 6-6 DMP optimized path configuration



Improving I/O performance with Dynamic Multi-Pathing in the hypervisor

When you install DMP in the hypervisor, DMP provides visibility into the I/O paths and performance. You can use the datacenter view or the individual statistics to determine if certain LUNs are not performing as well as other LUNs. If you determine a bottleneck, you can move some of the I/O workload (and thus Virtual Machines) to another datastore to alleviate the I/O performance contention.

The following scenario describes a typical use case for improving I/O performance.

Use the datacenter view of DMP to determine which storage arrays are attached to which ESXi servers. In this scenario, one of the storage arrays is not performing well.

You can select the storage array that is not performing well to view additional I/O performance details at the VMDK level. You can use this information to understand I/O loads within individual guests and the actual subcomponents of the disks assigned to the guests.

Figure 6-7 shows the statistics for several LUNs. In the "Avg Time Write" column for the clariion0_5 device, its performance is significantly slower than the clariion0_1, despite similar workloads reflected by the "virtual machine" column.

Figure 6-7 Example I/O performance statistics

LUN	Virtual Machine	Virtual Disk	Operations		Blocks		Avg Time (ms)	
			Reads	Writes	Reads	Writes	Reads	Writes
enc_dataion0_1			4469597	1349346	211980271	77086062	7.26	1.95
	lab34-vmcat	-	15602	1790	615370	57078	7.25	3.93
	lab49-vmDMP vC	-	3683331	1008408	135988552	76074117	7.59	1.92
enc_dataion0_5			112921	122425	9298305	41230258	6.61	66.49
	lab33-vmg	-	44170	106342	3222759	41047323	13.94	75.68
	lab24-vmvmb	-	9020	8001	367435	152562	3.67	4.54
enc_dataion0_6			7896	0	2	0	0.09	0


For this example, DMP is not running in the guests. If you also have DMP running in the individual guests, you can correlate the I/O statistics to the particular guests.

If you determine a high I/O workload for a particular LUN, you may need to move some of the I/O workload (and thus Virtual Machines) to another datastore to alleviate the I/O performance contention.

Achieving simplified management using Dynamic Multi-Pathing in the hypervisor and guest

It is very common in VMware environments to have clusters of VMware systems that share common storage and configurations. It is not always easy to keep those configuration in sync across a number of servers, and this can place a burden on the administrators to ensure that the same settings are in place on all systems in the same cluster. DMP for VMware provides a mechanism to create and apply a template across systems to ensure that the same settings are in place. After an administrator configures a single host, the DataCenter View provides the ability to save those settings as a template, and then apply the common template to other hosts in the VMware DataCenter. This mechanism provides an easy and simplified process to make sure that all hosts are identically configured, with minimal administrator intervention.

Figure 6-8 DMP tuning template options



With native multi-pathing, administrators must often manually configure each storage LUN to select a multi-pathing policy. Even then, the best policy available is round robin, which does little to work around I/O congestion during unusual SAN workload or contention for a shared storage path.

With DMP, as storage is discovered, it is automatically claimed by the appropriate DMP Array Support Library (ASL). The ASL is optimized to understand the storage characteristics and provide seamless visibility to storage attributes. You do not need to perform any additional management tasks to ensure that all paths to a storage device are used to their optimum level. By default, DMP uses the Minimum Queue algorithm, which ensures that I/Os are optimally routed to a LUN path that has available capacity. Unlike round robin policy, the minimum queue algorithm does not force I/O onto overloaded paths (which round robin would do, as it has no visibility to contention).

As you configure or add additional storage, no additional administration is required beyond configuring the storage into a VMware datastore. If you made any optimizations on the storage enclosure, a new LUN automatically inherits the optimizations for that enclosure. You do not need to configure new storage LUNs to use any specific multi-pathing policy. DMP automatically leverages all paths available, with no need for administrator intervention.

Improving data protection, storage optimization, data migration, and database performance

This chapter includes the following topics:

- [Use cases for Veritas InfoScale product components in a VMware guest](#)
- [Protecting data with Veritas InfoScale product components in the VMware guest](#)
- [Optimizing storage with Veritas InfoScale product components in the VMware guest](#)
- [Migrating data with Veritas InfoScale product components in the VMware guest](#)
- [Improving database performance with Veritas InfoScale product components in the VMware guest](#)

Use cases for Veritas InfoScale product components in a VMware guest

Veritas InfoScale product components provide many features for enhancing storage in a VMware environment. You can use the InfoScale products for data protection, storage optimization, ease of data migration, and optimized performance.

Using the InfoScale product in VMware guest enables the following use cases:

Table 7-1 Storage Foundation support for VMware environment use cases

Virtualization use case	Veritas InfoScale product component	Implementation details
Data protection	Storage Foundation or SFHA in the guest	<p>How to use Storage Foundation point-in-time copy features for data protection.</p> <p>See “Protecting data with Veritas InfoScale product components in the VMware guest” on page 61.</p>
Storage optimization	Storage Foundation or SFHA in the guest	<p>How to use Storage Foundation thin provisioning, SmartTier, and SmartIO features to optimize storage in a VMware environment.</p> <p>See “Optimizing storage with Veritas InfoScale product components in the VMware guest” on page 62.</p>
Data migration	Storage Foundation or SFHA in the guest	<p>How to use Storage Foundation Portable Data Containers to migrate data safely and easily in a VMware guest environment.</p> <p>See “Migrating data with Veritas InfoScale product components in the VMware guest” on page 65.</p>
Improved database performance	Storage Foundation or SFHA in the guest	<p>How to use Storage Foundation database accelerators to improve database performance in a VMware guest environment.</p> <p>See “Improving database performance with Veritas InfoScale product components in the VMware guest” on page 66.</p>
Simplified storage management with Storage Foundation or SFHA	Dynamic Multi-Pathing (DMP) in the hypervisor and Storage Foundation or SFHA in the guest	<p>How to use DMP for VMware and Storage Foundation features for end-to-end storage visibility, performance, optimization, and enhanced ease of management.</p>

Protecting data with Veritas InfoScale product components in the VMware guest

When you install InfoScale in the VMware guest, you can use the following Veritas InfoScale product point-in-time copy technologies to protect your data:

Table 7-2 Data protection options in Veritas InfoScale product components

Data protection feature	Description
FlashSnap	Volume-based point-in-time snapshot and recovery method
Database FlashSnap	Database-optimized point-in-time volume snapshots and recovery method
Storage Checkpoints	File system-based point-in-time copy and recovery method
Database Storage Checkpoints	Database-optimized point-in-time file system copy and recovery method
FileSnap	File-level point-in-time snapshot and recovery method
Volume replication	Ongoing method of volume replication to a remote location
File replication	Ongoing method of file system replication to a remote location

About point-in-time copies

Storage Foundation offers a flexible and efficient means of managing business-critical data. Storage Foundation lets you capture an online image of an actively changing database at a given instant, called a point-in-time copy.

More and more, the expectation is that the data must be continuously available (24x7) for transaction processing, decision making, intellectual property creation, and so forth. Protecting the data from loss or destruction is also increasingly important. Formerly, data was taken out of service so that the data did not change while data backups occurred; however, this option does not meet the need for minimal down time.

A point-in-time copy enables you to maximize the online availability of the data. You can perform system backup, upgrade, or perform other maintenance tasks on the point-in-time copies. The point-in-time copies can be processed on the same host as the active data, or a different host. If required, you can offload processing of the point-in-time copies onto another host to avoid contention for system resources on your production server. This method is called off-host processing. If implemented correctly, off-host processing solutions have almost no impact on the performance of the primary production system.

Point-in-time snapshots for Veritas InfoScale products in the VMware environment

The point-in-time snapshots are fully supported for Veritas InfoScale products in the VMware environment. In addition to using point-in-time snapshots to create back-ups or perform off-host processing, you can use them to provision additional virtual machines.

For more information about point-in-time snapshots, see *Veritas InfoScale Solutions Guide*.

Optimizing storage with Veritas InfoScale product components in the VMware guest

When you install Storage Foundation or Storage Foundation High Availability in the VMware guest, you can use the following Veritas InfoScale storage optimization technologies to maximize your storage utilization:

Table 7-3 Storage optimization options in Veritas InfoScale product components

Storage optimization feature	Description
Compression	Maximize storage utilization by reducing file sizes.
Thin Reclamation	Optimize your thin array usage by setting up thin storage.
SmartMove	Optimize your thin array usage with thin reclamation solutions to maintain thin storage.
SmartTier	Maximize storage efficiency by moving data to storage tiers based on age, priority, and access rate criteria.
SmartTier for Oracle	Oracle-optimized method to maximize storage efficiency by moving data to storage tiers based on age, priority, and access rate criteria.

About SmartTier in the VMware environment

SmartTier is a VxFS feature that enables you to allocate file storage space from different storage tiers according to rules you create. SmartTier provides a more flexible alternative compared to current approaches for tiered storage. Static storage tiering involves a manual one-time assignment of application files to a storage class, which is inflexible over a long term. Hierarchical Storage Management solutions typically require files to be migrated back into a file system name space before an

application access request can be fulfilled, leading to latency and run-time overhead. In contrast, SmartTier allows organizations to:

- Optimize storage assets by dynamically moving a file to its optimal storage tier as the value of the file changes over time
- Automate the movement of data between storage tiers without changing the way users or applications access the files
- Migrate data automatically based on policies set up by administrators, eliminating operational requirements for tiered storage and downtime commonly associated with data movement

Note: SmartTier is the expanded and renamed feature previously known as Dynamic Storage Tiering (DST).

You can use the SmartTier functionality of InfoScale in a VMware environment similar to how you use it in the physical environment. For example, you may have two different types of storage arrays, some SSD and some HDD. You can allocate one VMFS on each type of disk. You can use the SSD devices as one storage tier and the HDD devices as the second storage tier.

For more information about SmartTier, see *Storage Foundation Cluster File System High Availability Administrator's Guide* or *Storage Foundation Administrator's Guide*.

About compression with Veritas InfoScale product components in the VMware guest

Compressing files reduces the space used by files, while retaining the accessibility of the files and being transparent to applications. Compressed files look and behave almost exactly like uncompressed files: the compressed files have the same name, and can be read and written as with uncompressed files. Reads cause data to be uncompressed in memory, only; the on-disk copy of the file remains compressed. In contrast, after a write, the new data is uncompressed on disk.

The compression functionality is the same in the VMware environment as in the physical devices.

For more information about file compression, see *Storage Foundation Cluster File System High Availability Administrator's Guide* or *Storage Foundation Administrator's Guide*.

About thin reclamation with Veritas InfoScale product components in the VMware guest

Veritas InfoScale product components support reclamation of the unused storage on thin-reclamation capable arrays and LUNs. InfoScale Solutions automatically discover LUNs that support thin reclamation.

If you configure Veritas InfoScale product components in Raw Device Mapping mode, Veritas InfoScale product components can control the storage in the VMware guests directly. With this configuration, you can use the thin reclamation functionality of Veritas InfoScale product components on the VMware guest. This configuration may limit some VMware functionality such as vMotion. You cannot use the thin reclamation functionality of Veritas InfoScale product components if the storage is configured in VMDK mode. Veritas InfoScale product components do not have the visibility to the physical storage under the VMware guest to reclaim the freed storage.

For more information about thin reclamation, see *Storage Foundation Cluster File System High Availability Administrator's Guide* or *Storage Foundation Administrator's Guide*.

About SmartMove with Veritas InfoScale product components in the VMware guest

The SmartMove utility optimizes move and copy operations. The SmartMove utility leverages the knowledge that Veritas File System (VxFS) has of the Veritas Volume Manager (VxVM) storage. VxFS lets VxVM know which blocks have data. When VxVM performs an operation that copies or moves data, SmartMove enables the operation to only copy or move the blocks used by the file system.

This capability improves performance for synchronization, mirroring, and copying operations because it reduces the number of blocks that are copied. SmartMove only works with VxFS file systems that are mounted on VxVM volumes. If a file system is not mounted, the utility has no visibility into the usage on the file system.

The SmartMove functionality is the same in the VMware environment as in the physical devices.

For more information about SmartMove, see *Storage Foundation Cluster File System High Availability Administrator's Guide* or *Storage Foundation Administrator's Guide*.

About SmartTier for Oracle with Veritas InfoScale product components in the VMware guest

SmartTier is a VxFS feature that enables you to allocate file storage space from different storage tiers according to rules you create. SmartTier provides a more flexible alternative compared to current approaches for tiered storage.

In an Oracle database environment, the access age rule can be applied to archive log files and Flashback files. Oracle updates the header of each datafile at every database checkpoint and hence access age rules cannot be used for datafiles. For a partitioned table, we can use the name base rule to relocate files belonging to a given partition, for instance last year, to the secondary storage tier. However if a database does not have partitioned tables, current methods for relocation do not fit to the Oracle database environment.

To adopt a successful SmartTier policy and save storage costs, a method is needed for relocating Oracle objects which are accessed infrequently to the secondary tier. Relocating the entire file is not useful in all cases.

SmartTier for Oracle in the VMware environment works the same as in the physical environment.

For more information about SmartTier for Oracle, see *Veritas InfoScale Storage and Availability Management for Oracle Databases*.

Migrating data with Veritas InfoScale product components in the VMware guest

When you install the InfoScale products Storage Foundation or Storage Foundation High Availability in the VMware guest, you can use the following Veritas InfoScale data migration technologies to migrate your data safely and easily:

- Portable Data Containers
- Volume mirroring

Types of data migration

This section describes the following types of data migration:

- Migrating data between platforms using Cross-platform Data Sharing (CDS)
Storage Foundation lets you create disks and volumes so that the data can be read by systems running different operating systems. CDS disks and volumes cannot be mounted and accessed from different operating systems at the same time. The CDS functionality provides an easy way to migrate data between one system and another system running a different operating system.

- Migrating data between arrays
Storage Foundation supports arrays from various vendors. If your storage needs change, you can move your data between arrays.

Note: The procedures are different if you plan to migrate to a thin array from a thick array.

Improving database performance with Veritas InfoScale product components in the VMware guest

When you install the InfoScale products in the VMware guest, you can use the following Veritas InfoScale database accelerator technologies to enhance your database performance:

- Veritas Oracle Disk Manager
- Veritas Cached Oracle Disk Manager
- Veritas Concurrent I/O

Database accelerators enable your database to achieve the speed of raw disk while retaining the management features and convenience of a file system.

About Veritas InfoScale product components database accelerators

The major concern in any environment is maintaining respectable performance or meeting performance service level agreements (SLAs). Veritas InfoScale product components improve the overall performance of database environments in a variety of ways.

Table 7-4 Veritas InfoScale product components database accelerators

Veritas InfoScale database accelerator	Supported databases	Use cases and considerations
Oracle Disk Manager (ODM)	Oracle	<ul style="list-style-type: none"> ■ To improve Oracle performance and manage system bandwidth through an improved Application Programming Interface (API) that contains advanced kernel support for file I/O. ■ To use Oracle Resilvering and turn off Veritas Volume Manager Dirty Region Logging (DRL) to increase performance, use ODM. ■ To reduce the time required to restore consistency, freeing more I/O bandwidth for business-critical applications, use SmartSync recovery accelerator.
Cached Oracle Disk Manager (Cached ODM)	Oracle	To enable selected I/O to use caching to improve ODM I/O performance, use Cached ODM.
Concurrent I/O	DB2 Sybase	<p>Concurrent I/O (CIO) is optimized for DB2 and Sybase environments</p> <p>To achieve improved performance for databases run on VxFS file systems without restrictions on increasing file size, use Veritas InfoScale Concurrent I/O.</p>

These database accelerator technologies enable database performance equal to raw disk partitions, but with the manageability benefits of a file system. With the Dynamic Multi-pathing (DMP) feature of Storage Foundation, performance is maximized by load-balancing I/O activity across all available paths from server to array. DMP supports all major hardware RAID vendors, hence there is no need for third-party multi-pathing software, reducing the total cost of ownership.

Veritas InfoScale database accelerators enable you to manage performance for your database with more precision.

Setting up virtual machines for fast failover using Storage Foundation Cluster File System High Availability on VMware disks

This chapter includes the following topics:

- [About use cases for InfoScale Enterprise in the VMware guest](#)
- [Storage Foundation Cluster File System High Availability operation in VMware virtualized environments](#)
- [Storage Foundation functionality and compatibility matrix](#)
- [About setting up Storage Foundation Cluster File High System High Availability on VMware ESXi](#)

About use cases for InfoScale Enterprise in the VMware guest

In addition to the application availability and storage capabilities, InfoScale Enterprise adds the fast failover capability of a parallel cluster.

Storage Foundation Cluster File System High Availability operation in VMware virtualized environments

Storage Foundation Cluster File System High Availability (SFCFSHA) has two supported modes of operation when used inside a guest operating system operating system (OS) that is running on a VMware hypervisor:

- Mode 1: SFCFSHA is connected to external storage via RDM-P
Use SFCFSHA running in a guest OS and connected to external storage via RDM-P when you require highly reliable SCSI-3 PGR keys for split-brain protection and data fencing
- Mode 2: SFCFSHA connected to VMFS with the VMFS multi-writer flag enabled
Use SFCFSHA running in a guest OS and connected to VMFS with the VMFS multi-writer flag enabled when you require guest OS mobility via VMware vMotion.

Requirements for Mode 1:

- SFCFSHA must be connected to a physical LUN has been mapped to a virtual machine disk (VMDK) using the VMware raw disk mapping physical (RDMP) compatibility mode.
- This option provides full SFCFSHA high availability support for fast failover, split brain protection, and data fencing, but does not allow use of VMware snapshots, vMotion, or other VMware HA features.
- SFCFSHA snapshots, clones, deduplication, and other storage optimization features are fully supported in this configuration.

Requirements for Mode 2:

- The guest OS must use a VMFS virtual machine disk (VMDK).
- VMFS normally prohibits multiple guest OSes from connecting to the same virtual disk, thus precluding use of that VMDK with a parallel file system such as SFCFSHA. However, by enabling a new VMFS option, multi-writer (VMware k/b article 1034165), an administrator can create a VMDK that can be read/written to by multiple guest OSes simultaneously.
- Since VMFS does not allow SFCFSHA to see the SCSI-3 persistent group reservation (PGR) bit on the LUN, SCSI-based I/O fencing will not work. To use SFCFSHA with the VMFS multi-writer option, Coordination Point Servers (CPS) must be used for split-brain protection.

For information on configuring CPS fencing, see the *Storage Foundation Cluster File System High Availability Administrator's Guide*.

Storage Foundation functionality and compatibility matrix

- CPS does not support data protection via fencing, so care must be taken to prevent inadvertent data corruption caused by a non-SFCFSHA guest OS writing to a VMDK. SFCFSHA snapshots, clones, deduplication, and other storage optimization features are fully supported in this configuration.

Veritas support for SFCFSHA in a guest OS:

- Veritas will attempt to resolve any issues as if they were occurring in a physical OS environment. Once a potential problem has been identified, Veritas support personnel will recommend an appropriate solution that would be applicable on the native OS on a physical server.
- If that solution does not work in the VMware virtualized environment, Veritas reserves the right to ask the customer to replicate the problem in a physical environment. In some cases the customer may be referred to VMware for support.

Storage Foundation functionality and compatibility matrix

Table 8-1 shows the Storage Foundation functionality and compatibility with VMware ESXi disk modes.

Table 8-1 Storage Foundation functionality and compatibility matrix with VMware disk modes

Storage Foundation	VMware ESXi disk mode: Virtual Disk (VMDK)	VMware ESXi disk mode: Raw Device Mapping Logical mode	VMware ESXi disk mode: Raw Device Mapping Physical mode
VxVM Disk format: simple, sliced	Yes	Yes	Yes
VxVM Disk format: cdsdisk	Yes	Yes	Yes
I/O fencing	Yes (with non-SCSI3-PR based fencing) ¹	Yes (with non-SCSI3-PR based fencing) ¹	Yes (with disks configured in RDM-P mode) ¹
Portable Data Containers	No	No	Yes
Dynamic Multi-Pathing in the guest	No ²	No ²	No ²

Table 8-1 Storage Foundation functionality and compatibility matrix with VMware disk modes *(continued)*

Storage Foundation	VMware ESXi disk mode: Virtual Disk (VMDK)	VMware ESXi disk mode: Raw Device Mapping Logical mode	VMware ESXi disk mode: Raw Device Mapping Physical mode
Volume Replicator	Yes	Yes	Yes
CVM/VVR	Yes	Yes	Yes
Bunker node (non-CVM environment)	Yes	Yes	Yes
DDL extended attributes	No	No	Yes
Thin reclamation	No	No	Yes

See [“I/O fencing considerations in an ESXi environment”](#) on page 17.

About setting up Storage Foundation Cluster File High System High Availability on VMware ESXi

This sample deployment illustrates how to install and configure Storage Foundation Cluster File System High Availability (SFCFSHA) in a VMware virtual server using VMware filesystem (VMFS) virtual disks (VMDKs) as the storage subsystem.


The information provided here is not a replacement or substitute for SFCFSHA documentation nor for the VMware documentation. This is a deployment illustration which complements the information found in other documents.

The following product versions and architecture are used in this example deployment:

- RedHat Enterprise Linux (RHEL) Server 6.2
- Storage Foundation Cluster File System High Availability
- ESXi 5.1

A four node virtual machine cluster will be configured on two VMware ESXi servers. Shared storage between the two ESXi servers using Fibre Channel has been setup. The Cluster File System will exist across four virtual machines: cfs01, cfs02, cfs03, and cfs04. Three Coordination Point (CP) servers will be used: cps1, cps2, and cps3 (this one placed in a different ESXi server). For storage, five data stores will be used and one shared VMDK file will be placed in each data store.

Figure 8-1 SFCFSHA virtual cluster on VMware ESXi



Two private networks will be used for cluster heartbeat. They are called PRIV1 and PRIV2. Virtual switch vSwitch2 also has the VMkernel Port for vMotion enabled. vSwitch0 is used for management traffic and the public IP network.

Some blades have a two network limit. If this is the case, configure one network for heartbeats and the other one as a heartbeat backup (low priority setting).

Planning a Storage Foundation Cluster File System High Availability (SFCFSHA) configuration

The deployment example illustrates the recommended steps for installing and configuring SFCFSHA using CP servers. To achieve a successful SFCFSHA implementation on VMware, follow the recommended steps. Different versions of RHEL or ESXi may require different steps.

See the SFCFSHA documentation for additional information.

Table 8-2 Tasks for deploying SFCFSHA on VMware

SFCFSHA deployment task	Deployment steps
Deploy virtual machines (VMs)	See RHEL and ESXi documentation.
Enable password-less SSH	See “Enable Password-less SSH” on page 73.
Enable traffic to CP server and management ports	See “Enabling TCP traffic to coordination point (CP) Server and management ports” on page 74.

About setting up Storage Foundation Cluster File High System High Availability on VMware ESXi

Table 8-2 Tasks for deploying SFCFSHA on VMware (*continued*)

SFCFSHA deployment task	Deployment steps
Configure CP servers	See “Enabling TCP traffic to coordination point (CP) Server and management ports” on page 74.
Deploy SFCFSHA software	See “Deploying Storage Foundation Cluster File System High Availability (SFCFSHA) software” on page 79.
Configure SFCFSHA	See “Configuring Storage Foundation Cluster File System High Availability (SFCFSHA)” on page 81.
Configure non-SCSI3 fencing	

Enable Password-less SSH

The installer will be able to password-less configure SSH/RSH among the cluster nodes, but it will not be able to enable this required functionality between the cluster nodes and the CP servers. For example, later setting changes, like modifying the IO fencing configuration, may need password-less SSH/RSH. In this configuration we are going to configure password-less SSH between the node that we will be using to configure the cluster and the rest of the nodes.

The following are the instances where Password-less SSH will be enabled:

Table 8-3 Instances to enable password-less SSH

Source	Target	Reason
cfs01	cfs02	Cluster configuration
cfs01	cfs03	Cluster configuration
cfs01	cfs04	Cluster configuration
cfs01	cps01	Non-SCSI-3 fencing configuration
cfs01	cps2	Non-SCSI-3 fencing configuration
cfs01	cps3	Non-SCSI-3 fencing configuration
cps01	cps01	Secure cluster configuration
cps2	cps01	Secure cluster configuration

Table 8-3 Instances to enable password-less SSH (*continued*)

Source	Target	Reason
cps3	cps01	Secure cluster configuration

If you do not enable Password-less SSH, then follow the manual configuration instructions given at the end of this document.

Enabling TCP traffic to coordination point (CP) Server and management ports

For successful intra-cluster communication, make sure that cluster nodes and CP servers can be reached on port 14250 (or any other if you changed the default). If RedHat Firewall has been enabled, make sure there is a rule to allow the connection to ports 14250 and 14149.

To enable TCP traffic to CP server and management ports

1 Stop iptables service:

```
[root@cps3 sysconfig]# service iptables stop
iptables: Flushing firewall rules: [ OK ]
iptables: Setting chains to policy ACCEPT: filter [ OK ]
iptables: Unloading modules: [ OK ]
```

2 Enter the following lines at the /etc/sysconfig/iptables file:

```
-A INPUT -p tcp -m tcp --dport 14250 -j ACCEPT
-A INPUT -p tcp -m tcp --dport 14149 -j ACCEPT
```

3 Start the service:

```
[root@cps2 ~]# service iptables restart
iptables: Flushing firewall rules: [ OK ]
iptables: Setting chains to policy ACCEPT: filter [ OK ]
iptables: Unloading modules: [ OK ]
iptables: Applying firewall rules: [ OK ]
[root@cps2 ~]#
```

4 Verify the new rule is in place:

```
[root@cps3 sysconfig]# iptables --list

Chain INPUT (policy ACCEPT)
target          prot opt source          destination
ACCEPT         tcp  -- anywhere        anywhere         tcp dpt:cps
ACCEPT         tcp  -- anywhere        anywhere         tcp dpt:vrts-tdd
ACCEPT         tcp  -- anywhere        anywhere         tcp dpt:xprtld
```

For the SFCFSHA deployment example, these rules must be enabled on cfs01, cfs02, cfs03, cfs04, cps1, cps2 and cps3

Configuring coordination point (CP) servers

In order to protect the cluster from data corruption caused by a cluster failure that results in “split-brain,” coordination point (CP) servers must be deployed. If you do not have CP servers configured already, you must configure the CP servers that are used to provide non-SCSI-3 fencing protection at this point in the deployment process.

CP servers can be used for multiple clusters, so this step will only be necessary the first time a cluster is configured and no CP servers are available.

Tasks for configuring a new CP server:

Table 8-4

Step	Description
Configure a new Coordination Point server for Storage Foundation Cluster File System High Availability (SFCFSHA)	See “Configuring a Coordination Point server for Storage Foundation Cluster File System High Availability (SFCFSHA)” on page 76.
Configure a Cluster Server (VCS) single node cluster	See “Configuring a Cluster Server (VCS) single node cluster” on page 78.

Table 8-4 (continued)

Step	Description
Configure a Coordination Point server service group	See “Configuring a Coordination Point server service group” on page 77.

Configuring a Coordination Point server for Storage Foundation Cluster File System High Availability (SFCFSHA)

In the example SFCFSHA deployment, three CP servers are used. Each CP server will be housed in a virtual machine, with each one forming a single node CP server cluster. Each of the two physical ESX Servers will contain one CP server, and a third CP server will be located in another location. To illustrate the process, steps are given to deploy the cps1 node.

The installer script will be used to perform the deployment of Cluster Server (option 2). When running the Pre-Installation Check, the following packages will be detected as not installed:

```
CPI Error V-(-30-2225 The following required OS rpms (or higher
version) were not found on cfs01: nss-softokn-freebl-3.12.9-3.el6.i686
glibc-2.12-1.25.el6.i686 pam-1.1.1-8.el6.i686 libstdc++-4.4.5-6.el6.i686
libgcc-4.4.5-6.el6.i686 ksh-20100621-6.el6.x86_64
```

The requirement for those missing packages SF 6.0 and 6.0.1 on RHEL 6.2/6.3 is documented in article [TECH196954](#).

As explained, you must install the following packages before deploying Storage Foundation:

- glibc-2.12-1.25.el6.i686
- libgcc-4.4.5-6.el6.i686
- libstdc++-4.4.5-6.el6.i686
- nss-softokn-freebl-3.12.9-3.el6.i686

We are deploying on top of RedHat 6.2, and these are the RPMs installed:

- rpm -ivh --nodeps glibc-2.12-1.47.el6.i686.rpm
- rpm -ivh --nodeps libgcc-4.4.6-3.el6.i686.rpm
- rpm -ivh --nodeps libstdc++-4.4.6-3.el6.i686.rpm
- rpm -ivh --nodeps nss-softokn-freebl-3.12.9-11.el6.i686.rpm
- rpm -ivh --nodeps pam-1.1.1-10.el6.i686.rpm

About setting up Storage Foundation Cluster File High System High Availability on VMware ESXi

- `rpm -ivh --nodeps ksh-20100621-12.el6.x86_64.rpm`

Verify that CP servers can listen on port 14250. Disable the firewall rules or enter a new rule as explained at the beginning of this document to allow communication to this port

Verify that the CP servers have Password-less SSH connection to the cluster node where fencing configuration will be run.

Configuring a Coordination Point server service group

Even in a single node cluster, a virtual IP address (VIP) is used. This enables the creation of a VCS resource to control the VIP availability. For the example configuration, a VIP for each CP server is assigned to illustrate the process.

To configure CP server service group

- 1 Verify that you have a VIP available for each of your CP servers.
- 2 Run the command:

```
# /opt/VRTS/install/installvcs<version> -configcps
```

Where *<version>* is the specific release version

- 3 When the installer prompts if you want to configure a CP server, select **Configure Coordination Point Server on single node VCS system.**
- 4 The name of the CP server is the same as the host name plus “v” at the end. For the example configuration CP server it is cps1v.
- 5 Enter the Virtual IP address that is associated to the single node cluster. In the example of node cps1, it is 10.182.99.124. Accept the default port suggested.
- 6 As discussed before, security is enabled for the example and is recommended as a best practice.

About setting up Storage Foundation Cluster File High System High Availability on VMware ESXi

- 7 When prompted, enter the location of the database. In the example, the database will be installed locally, so you can accept the default location.
- 8 After reviewing the configuration parameters, continue with the configuration of the CP server Service Group. The NIC used at cps1 is eth4. The example does not use NetworkHosts. Enter the netmask, and the configuration is complete.

The CPSSG Service Group is now online.

```
[root@cps1 rhel6_x86_64]# hastatus -sum
```

-- SYSTEM STATE				
-- System	State	Frozen		
A cps1	Running	0		
-- GROUP STATE				
-- Group	System	Probed	AutoDisabled	State
B CPSSG	cps1	Y	N	ONLINE

```
[root@cps1 rhel6_x86_64]#
```

Configuring a Cluster Server (VCS) single node cluster

Configure a VCS single node cluster for the Coordination Point server. The following procedure illustrates the process for the example configuration.

To configure a single-node cluster

- 1 Run the installer script and perform a Pre-Installation Check to verify that the server is ready to have VCS installed.
- 2 Select the product installation option, and when prompted to make a selection choose VCS.
- 3 Select the option to install all RPMs. The VRTScps package is included.
- 4 Enter the name of the host where you are running the configuration, cps1 in this example. After reviewing all the packages that will be installed, the package installation begins. In the example environment, the keyless licensing is required because a Veritas Operation Manager host is deployed. Global Cluster Option will not be enabled.
- 5 When prompted to configure VCS, enter **y**.

About setting up Storage Foundation Cluster File High System High Availability on VMware ESXi

- 6 When prompted to configure fencing, enter **n**. Fencing is not needed for a single-node cluster.
- 7 Although configuring a single node, you must enable LLT and GAB for Coordination Point Server (CP server) clients to connect to the CP server. When prompted to start GAB and LLT, choose **y**.
- 8 Enter the cluster name (same as node). For the configuration example, it is cps1.
- 9 Select **Configure heartbeat links using LLT over Ethernet**.
- 10 In the configuration example, an additional NIC must be configured for the private network. This step is necessary even with a single node cluster.
Enter a unique cluster ID that is not already in use. The installer verifies that the cluster ID is not already in use.
- 11 You do not need to enter a Virtual IP, as the same one used for the host will be fine to use.
- 12 You must configure secure mode or not according to your own configuration requirements. For the example configuration, SFCFSHA is configured to use secure mode, so the Coordination Point Servers must also use secure mode.
- 13 The FIPS option has not been qualified yet with CP servers. When prompted to choose, select **Configure the cluster in secure node without fips**.
This also matches the configuration done for the cluster nodes.
- 14 For the example configuration, SMTP and SNMP are not used. For your own configuration you must choose based on your requirements.
- 15 The VCS configuration starts after the selections are completed. Once finished, you can verify if VCS is running.

```
[root@cps1 rhel6_x86_64]# hastatus -sum

-- SYSTEM STATE
-- System                State                Frozen

A  cps1                  Running             0
[root@cps1 rhel6_x86_64]#
```

Deploying Storage Foundation Cluster File System High Availability (SFCFSHA) software

Deploy SFCFSHA in each of the four virtual machines that make up the cluster. To do that, you could have either selected **yes** in the previous CPI step, so the install

About setting up Storage Foundation Cluster File High System High Availability on VMware ESXi

could continue, or you could have run the installer script again. If you ran the installer script again, use the following procedure to deploy SFCFSHA.

To deploy SFCFSHA after running the installer script again

- 1** Make the following selections as you run the installer script:
 - 1) Install a Product
 - 5) Veritas InfoScale Enterprise
 - Agree the terms of the EULA.
- 2** When you are prompted to select the RPMs to install on all systems, select option 3, Install all RPMs.
- 3** Enter the name of the nodes of the cluster (cfs01 cfs02 cfs03 cfs04) on which to install the RPMs.

The installer verifies the pre-requisites again, and a list of all the RPMs that will be installed will be printed. The RPMs are installed.
- 4** Once the RPMs have been installed, the installer prompts you for the license. Select option 2 to enable keyless licensing so this cluster can be managed by VOM (Veritas Operation Manager):
- 5** You do not need licenses to enable the following options for this SFCFSHA deployment:
 - Replication
 - Global Cluster Option

Once the licenses have been registered, the installer finishes with the deployment.

You can now start the configuration. If you want a Virtual Machine template with the SFCFSHA software already deployed, stop here and take a snapshot or other copy of this image. The next step will be to run “installer –configure” to continue with the configuration.

Configuring Storage Foundation Cluster File System High Availability (SFCFSHA)

To configure SFCFSHA cluster settings

- 1 Run **installer –configure** or just continue from where we left in the previous step entering **y**.
- 2 Fencing would normally be the next step in configuring SFCFSHA. However, the I/O fencing configuration depends on other factors which are not yet determined:
 - VMDKs or RDMP storage devices are used
 - How I/O and network paths are configured
 - Configuration of coordination point (CP) server (or, in some cases, Coordination Disks)

For now you can enter **n** when prompted to configure IO fencing in enabled mode and come back to it later in the configuration process.

- 3 Configure the cluster name when prompted.

The cluster name for the example deployment is `cfs0`

- 4 Configure the NICs used for heartbeat when prompted.

LLT (Low Latency Protocol) can be configured over Ethernet or UDP. UDP is needed only when routing between the nodes is necessary. If UDP is not needed, then Ethernet is the clear recommendation.

In the example deployment, `eth4` and `eth5` are the private links. `Eth3` is the public link, and it will be only used as low priority heartbeat path (so it only will be used if the other two paths fail).

All media speed checking should succeed. If not, please review your node interconnections.

- 5 Configure the cluster ID when prompted. A unique cluster ID is needed: it is vital to choose a number that is not used in any other cluster. This is especially true when using the same network interconnection (both private and public). The CPI generates a random number, and checks the network to make sure that packets with that ID do not exist. However the CPI cannot guarantee that the ID is not being used in a cluster that is currently powered off. The best practice is to maintain a register of the cluster IDs used across the data center to avoid use of duplicate IDs. In the example configuration, no other clusters with that ID have been found.

About setting up Storage Foundation Cluster File High System High Availability on VMware ESXi

- 6 At this point a summary of the configuration to be deployed is presented. Examine the summary and enter **y** if everything is correct. If not enter **n** and go through the steps again.
- 7 The installer prompts for a Virtual IP to manage the cluster. This is not mandatory, and the cluster can be configured without that IP. Depending on your implementation, it might be a best practice.
- 8 Decide whether or not to use secure mode.
In the past, the difficulty in configuring Cluster Server secure mode deterred many users from using it. For SFCFSHA:
 - Secure mode configuration is much easier
 - The installer takes care of the entire configuration
 - A validated user and password from the OS is used instead of the traditional admin/password loginFor demonstration purposes, secure mode is used in the example deployment, but feel free to choose the option that best suits your needs.
FIPS is not used for the example configuration as it is not certified for deployment with CP servers. Option 1, **secure mode without FIPS** is used.
- 9 SMTP is not needed for the example.
- 10 SNMP notifications are not needed for the example.
At this point the cluster configuration will be initiated.

Configuring non-SCSI3 fencing

VMDK files do not currently support SCSI-3 Persistent Reservation and therefore non-SCSI-3 PR fencing must be used. Coordination point (CP) servers provide the required level of server based fencing. At this point in the configuration process, the three CP servers that are to be used with this cluster should be available and the CP service should be up and running.

To configure non-SCSI-3 fencing

- 1** If you started at the beginning of the installer process and selected the enable fencing option, you are prompted to configure fencing.

If you chose not to enable fencing at that point, then the cluster configuration is finished. You should now run `installsfcsfsha61 -fencing` to enable fencing in the cluster.
- 2** Regardless of how you navigated to the fencing configuration of the installer, select option 1 for Coordination Point client-based fencing.
- 3** When prompted if your storage environment supports SCSI-3 PR, select **n**, since VMDK files do not support SCSI-3 PR.
- 4** When prompted if you want to configure Non-SCSI-3 fencing, select **y**.
- 5** For production environments, three CP servers are recommended. Enter **3** when prompted for the number of coordination points.
- 6** Specify how many interfaces the CP servers will be listening on and the IP address of each interface. If a CP server is reachable via several networks, the best practice is to configure every interface. This allows the SFCFSHA nodes maximum communication flexibility, should a race condition occur.

Enter the host names and VIPs for the other CP servers and review the fencing configuration.
- 7** When prompted, select secure mode. All the trusted relationships between cluster nodes and CP servers are automatically set up.
- 8** Verify that the cluster information is correct. Each node is registered with each of the CP servers. Once this is done, the installer will restart VCS to apply the fencing configuration. At this point we don't have any file system configured yet.
- 9** When prompted, it is a recommended best practice to configure the Coordination Point Agent on the client, so CP servers are proactively monitored from the cluster. This step completes the fencing configuration.

Once fencing configuration is complete, you can verify if it is correct.

To verify the fencing configuration

- 1 Query each of the CP servers to verify each node has been registered.

```
# CCPS_USERNAME=CPSADM@VCS_SERVICES
# CPS_DOMAINTYPE=vx
[root@cfs01 install1]# cpsadm -s cps1v -a list_nodes
ClusterName UUID Hostname(Node ID) Registered
=====
cfs0 {38910d38-1dd2-11b2-a898-f1c7b967fd89} cfs01(0) 1
cfs0 {38910d38-1dd2-11b2-a898-f1c7b967fd89} cfs02(1) 1
cfs0 {38910d38-1dd2-11b2-a898-f1c7b967fd89} cfs03(2) 1
cfs0 {38910d38-1dd2-11b2-a898-f1c7b967fd89} cfs04(3) 1
[root@cfs01 install1]# cpsadm -s cps1v -a list_membership -c cfs0
List of registered nodes: 0 1 2 3
```

- 2 Run the same command against the each CP server.
- 3 Using the VCS Cluster Explorer screen, we can see that the vxfen service group has been created to monitor CP servers and that it is healthy.

Reference

- [Appendix A. Known issues and limitations](#)
- [Appendix B. Where to find more information](#)

Known issues and limitations

This appendix includes the following topics:

- [Prevention of Storage vMotion](#)

Prevention of Storage vMotion

In a configuration where VMDK files are used with the multi-writer flag, any attempt of migrating the VMDK file to another data store will be prevented with an error.

The operation is unable to succeed.

In order to migrate VMDK to different storage, SFCFSHA functionalities can be used to transparently migrate data between different disks.

Where to find more information

This appendix includes the following topics:

- [Veritas InfoScale documentation](#)
- [Service and support](#)
- [About Veritas Services and Operations Readiness Tools \(SORT\)](#)

Veritas InfoScale documentation

The latest documentation is available on the Veritas Services and Operations Readiness Tools (SORT) website in the Adobe Portable Document Format (PDF).

See the release notes for information on documentation changes in this release.

Make sure that you are using the current version of documentation. The document version appears on page 2 of each guide. The publication date appears on the title page of each document. The documents are updated periodically for errors or corrections.

<https://sort.veritas.com/documents>

You need to specify the product and the platform and apply other filters for finding the appropriate document.

Service and support

To access the self-service knowledge base, go to the following URL:

https://www.veritas.com/support/en_US.html

About Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a Web site that automates and simplifies some of the most time-consuming administrative tasks. SORT helps you manage your datacenter more efficiently and get the most out of your Veritas products.

SORT can help you do the following:

- | | |
|---|---|
| Prepare for your next installation or upgrade | <ul style="list-style-type: none">■ List product installation and upgrade requirements, including operating system versions, memory, disk space, and architecture.■ Analyze systems to determine if they are ready to install or upgrade Veritas products.■ Download the latest patches, documentation, and high availability agents from a central repository.■ Access up-to-date compatibility lists for hardware, software, databases, and operating systems. |
| Manage risks | <ul style="list-style-type: none">■ Get automatic email notifications about changes to patches, array-specific modules (ASLs/APMs/DDIs/DDLs), and high availability agents from a central repository.■ Identify and mitigate system and environmental risks.■ Display descriptions and solutions for hundreds of Veritas error codes. |
| Improve efficiency | <ul style="list-style-type: none">■ Find and download patches based on product version and platform.■ List installed Veritas products and license keys.■ Tune and optimize your environment. |

Note: Certain features of SORT are not available for all products. Access to SORT is available at no extra cost.

To access SORT, go to:

<https://sort.veritas.com>