

Veritas Storage Foundation™ In a VMware ESX Environment

Linux and Solaris x64 platforms

March 2009

TABLE OF CONTENTS

Introduction	3
Executive Summary.....	4
Overview	5
Virtual Machine File System	6
Virtual Disk	6
Raw Device Mapping.....	6
When to use Raw Device Mapping and Veritas Storage Foundation.....	7
Storage Foundation benefits in a VMware environment	8
Limitations.....	10
Volume Manager	10
Dynamic Multi-pathing.....	10
Disk format	10
Disk initialization.....	10
Volume Replicator	10
Storage Foundation.....	10
Portable Data Containers	10
File System.....	11
Cluster File System.....	11
I/O Fencing.....	11
Storage Foundation functionality and compatibility matrix	11
VMware.....	12
Snapshots	12
VMotion	12
NPIV (N-Port ID Virtualization).....	12
Support matrix	14
Licensing and pricing policy	15
Storage Foundation Standard and Enterprise.....	15
Single-node configuration examples:	15
Multi-node configuration example.....	15
Common configurations explored	16
FAQ 17	

Introduction

This paper describes how to configure and use Veritas Storage Foundation in a VMware ESX environment. Because most features in Storage Foundation are supported in a VMware environment, this paper focuses on the features that work differently in a VMware environment and why certain features are not supported.

The operating systems covered are Linux and Solaris x64, Windows is not covered in this paper. In addition, the paper is limited to Storage Foundation. Veritas Cluster Server information is not included in this white paper.

It is assumed that the reader is familiar with VMware and the VMware architecture in addition to being familiar with Storage Foundation. Only in places where it is absolutely necessary to share VMware specific information is there an explanation of the involved components.

The support outlined in this paper strictly covers VMware ESX version 3.0.1 and later versions. There is no support planned for Storage Foundation with VMware Server, Workstation or Virtual Desktop.

Storage consumed by a virtual machine can be allocated directly over the virtual machine network interface (NFS, CIFS, iSCSI to mention a few protocols), bypassing the virtualization layer, and is therefore not affected by virtualization. Support of storage configured in such a way is not covered in this white paper and any statement of support is limited to Storage Foundation running in a VMware virtual machine.

Use of the VMware certification mark means that Version 4.1 and 5.0 of Veritas Storage Foundation has been certified to run in a VMware Virtual Machine environment (see Table 2 for additional details). As of this date, Symantec makes no representation with respect to future product versions and future certification standards.

Executive Summary

VMware server virtualization is a technology that has significant benefits for a data center environment and it is changing the architecture of data centers. However, while it is great at virtualizing operating system instances, it does not solve or address storage virtualization. VMware has simplistic storage functionality that successfully addresses the current VMware market however as more storage intensive applications are being deployed in virtual environments; true storage management technology is required.

Deploying Storage Foundation in a VMware environment makes storage easier to manage and use by creating a functional storage management layer that the storage administrator or system administrator can access and use as it is not hidden inside the VMware ESX server. Benefits include:

- **Reducing complexity and management overhead** of day to day operation by standardizing on a single toolset for both physical and virtual environments

- **Ensuring predictable performance.** This is crucial for mission critical data center applications such as databases and mail servers. The architecture of the solution must be able to handle and support the same behavior as in a physical environment.

Intrusive operations such as **array migration** and **data migration** can cause significant downtime for a large amount of servers, especially in a virtual environment where one physical server is hosting as many as 30-40 virtual servers. Having the tools and infrastructure available to support such operations is essential to the long-term success of the data center.

Storage Foundation delivers on these items as well as numerous other functions that help the administrator manage the **data center** more efficiently.

Using Veritas Storage Foundation in a VMware environment means that Storage Foundation is running in the operating system, inside the Virtual Machine (see Figure 1 for high-level architecture diagram). Storage Foundation does not run inside the VMware ESX kernel or in the Hypervisor.

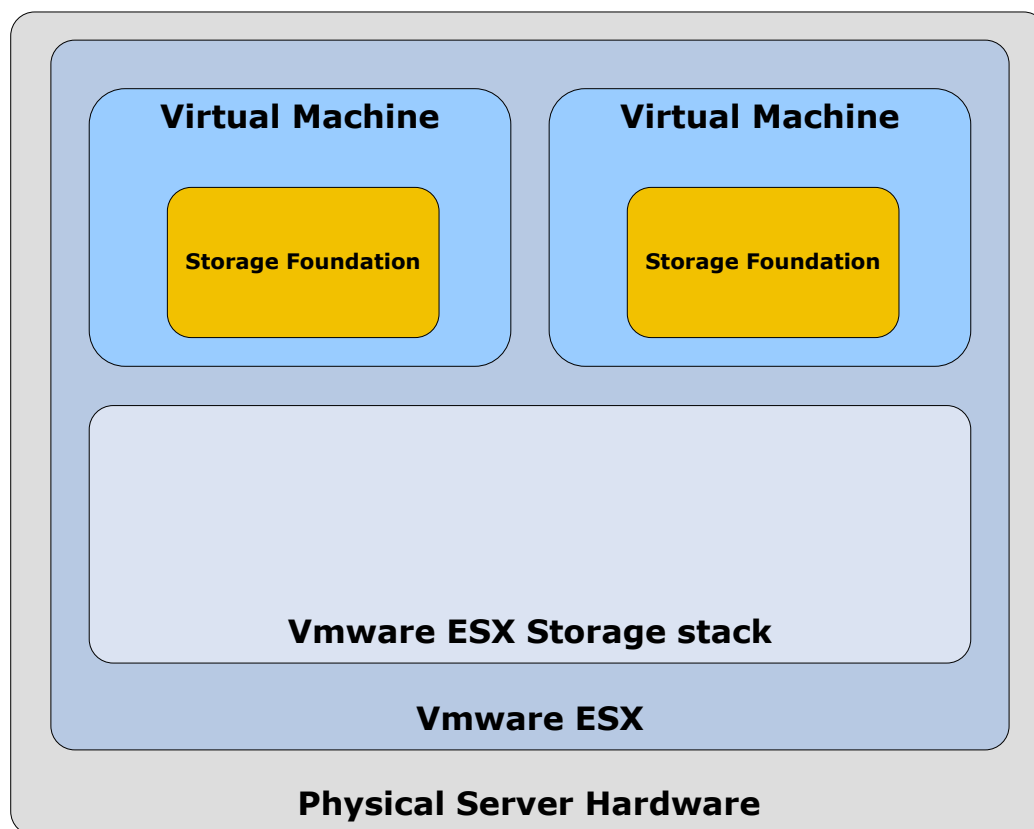


Figure 1 Architecture overview

VMware has several different methods to allocate block storage to a virtual machine¹

- File-based virtual disks created in VMFS or from NFS - **Virtual Disk**
- Block storage mapped from local disk, Fiber Channel LUN's or iSCSI – **Raw Device Mapping**

Specific configuration of VMware is required for certain features of Veritas™ Volume Manager to operate as it does in a physical server environment. For example, the basic operation of initializing a disk in Veritas Volume Manager requires a change of the default disk format (except in the case of raw device mapping, physical mode; see page 9 for more detail) and Dynamic Multi-pathing does not operate in VMware virtual machines. The details for Storage Foundation as well as how the specific VMware configuration affects VMware functionality are listed in the next section.

¹ If you are not familiar with these terms, please visit www.vmware.com and read through the VMware documentation, VMware product details are only covered briefly in this document.

Virtual Machine File System

Virtual Machine File System (VMFS) is VMware's clustered file system that is integrated into the VMKernel, providing a type of storage infrastructure for virtual machines. Its primary purpose is to store virtual machine disk images (Virtual Disks) however virtual machine configuration files and REDO-logs are also stored in the file system. Virtual machines are never exposed to VMFS.

Virtual Disk

A Virtual Disk is stored as a file (also referred to as a VMDK) in VMFS or in an NFS mount point on the VMware ESX server. For the virtual machine, a virtual disk appears as a SCSI disk, independent of what storage is used for VMFS.

Virtual Disks are supported with Volume Manager but with some limitations, for more details go to the Limitations chapter.

Raw Device Mapping

Raw Device Mapping (RDM) enables a virtual machine to have direct access to the storage rather than going via VMFS. RDM is configured on a 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, Fiber Channel disks) can be used with raw device mapping; Veritas Volume Manager supports all three types of disks. It is important to understand that there are two different modes for raw device mapping; the different modes affect the functionality and behavior of Storage Foundation and it is important to use the correct mode for the desired functionality.

Logical mode offers the same functionality and compatibility as a Virtual Disk with respect to VMware ESX features. Every SCSI command is virtualized in the virtual SCSI layer (see

Figure 5), giving the same behavior as a virtual disk but without having to bring the disk into a VMFS file system.

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. Physical mode is recommended as it enables maximum functionality of Volume Manager in a VMware environment.

When to use Raw Device Mapping and Veritas Storage Foundation

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 utilizing 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

Storage Foundation benefits in a VMware environment

Standardization of tools

Independent of how an operating system is hosted, consistent storage management tools will save the administrator time and reduce the complexity of the environment. Storage Foundation in the virtual machine can provide the same command set, storage namespace and environment as in a non-virtual environment.

Array migration

Array migration is a task that all organizations do more and more often as storage needs continue to grow. With VMware it is impossible to do online array migration, causing significant downtime for users of virtual machines. The high level process for VMware is described in Figure 2.

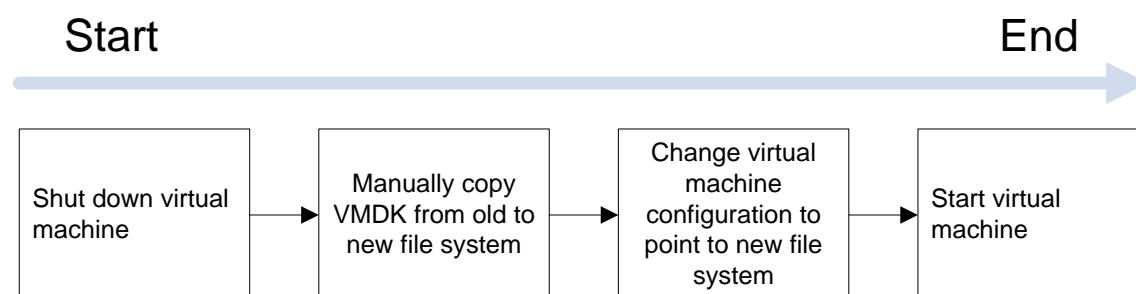


Figure 2 Virtual machine data migration with ESX tools

In VMware ESX 3.x it is possible to add disks to the ESX server and the virtual machine without reboot. This makes it possible to offer a better process utilizing Storage Foundation together with raw device mapped storage for online data migration (see Figure 3 for details).

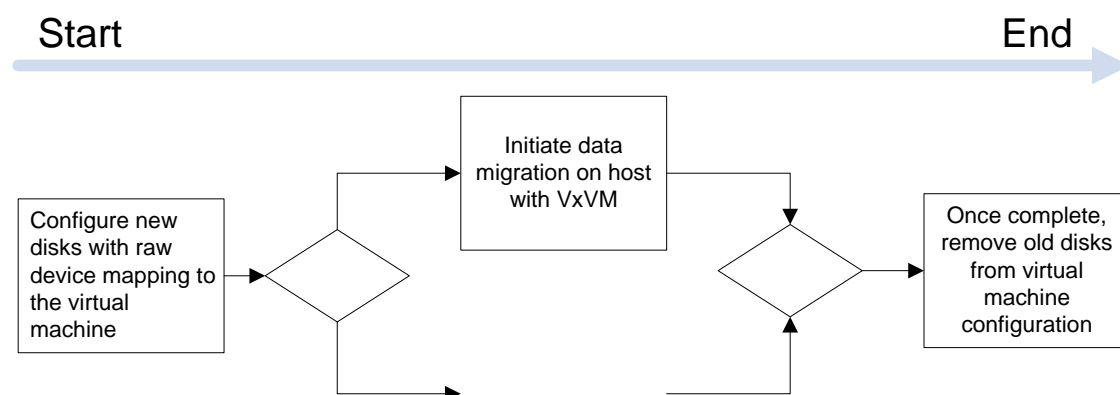


Figure 3 Virtual machine data migration with Storage Foundation

Data migration for Storage Foundation can be executed either locally in the virtual machine with 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 Veritas Storage Foundation Manager 1.1 (and later versions) (<http://www.symantec.com/sfm>).

Mirroring

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 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.

Moving storage between physical and virtual environments

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, Volume Manager will scan the device tree and discover the diskgroup and volume structure. An example workflow is described in Figure 4.

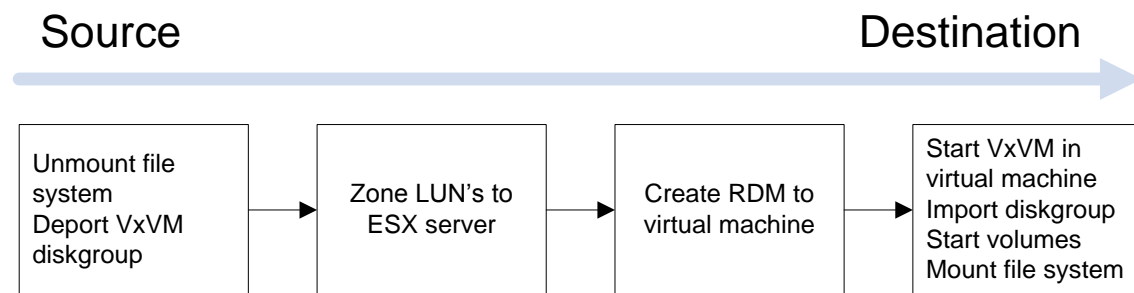


Figure 4 Migration workflow

Volume Manager is ignorant of the actual physical device entry, i.e. 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.

Volume Manager

Dynamic Multi-pathing

Dynamic Multi-pathing (DMP) is an integral part of the data path of Storage Foundation and can't be disabled. However, in a VMware environment, due to the architecture of VMware, DMP is not supported.

Figure 5 illustrates the data path for a Virtual Machine. As can be seen in

Figure 5, the Virtual Machine is only provided with a single path and the Hypervisor layer takes care of the multi-pathing, hence DMP is not used in the Virtual Machine.

It is technically possible to configure the same disk, with raw device mapping, over two different host bus adapters to a virtual machine. This is not supported because running two multi-pathing solutions on top of each other is a bad practice and will yield un-predictable failover results.

Disk format

Volume Manager supports three disk formats, `simple`, `sliced` and `cdsdisk`. `Cdsdisk` was introduced in Veritas Storage Foundation 4.0 and is the default disk format in Veritas Storage Foundation 4.1 and 5.0. `Cdsdisk` is required for Portable Data Containers functionality and is only supported with Raw Device Mapping in Physical mode (see Table 1 for details).

Disk initialization

When initializing a disk in VMware, the Volume Manager disk format limitation comes into play. Selecting the wrong disk format will give an error message when trying to initialize the disk. Table 1 outlines what disk formats are supported on the different types of VMware storage access methods.

Volume Replicator

Veritas Volume Replicator 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.

There are two configurations with VVR that are not supported at this time:

1. Bunker replication with the bunker node being a virtual machine
2. CVM/VVR replication. Replication of a SF CFS cluster that is running in virtual machines

Storage Foundation

Portable Data Containers

Portable Data Containers combines functionality from Veritas File System and Volume Manager, enabling movement of data (volume/file system) between supported operating systems. While the Portable Data Containers functionality is executed by a file system command, it also depends on Volume Manager disk format, `cdsdisk`. Portable Data Containers can therefore only be used with raw device mapping in physical mode.

File System

All features in Veritas File System are supported in a VMware virtual machine environment.

Cluster File System

Storage Foundation Cluster File System is supported when running inside a virtual machine however if I/O fencing is used then it requires special attention (see I/O fencing section).

It is possible to configure SF CFS with virtual disks or raw device mapping in logical mode but it is not supported for production environments, only for test and dev.

I/O Fencing

SCSI-3 Persistent reservations are only supported with raw device mapping (RDM) in physical mode. I/O fencing is not supported with RDM logical mode or with VMDK-based virtual disks.

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 ESX host as the WWN used for I/O fencing ends up being the same for each virtual machine.

CFS virtual machines (in the same CFS cluster) cannot share physical servers as the I/O fencing behavior will result in all nodes from that physical ESX host being fenced out if an event triggers the fencing functionality. In short, if I/O fencing is configured, the CFS nodes (in the same CFS cluster) has to be running on separate physical ESX hosts.

NPIV is currently under evaluation for I/O fencing testing when using multiple Virtual Machines on the same physical ESX host.

Storage Foundation functionality and compatibility matrix

Storage Foundation	VMware ESX disk modes		
	Virtual Disk (VMDK)	Raw Device Mapping Logical mode	Raw Device Mapping Physical mode
VxVM Disk format: simple, sliced	Yes	Yes	Yes
VxVM Disk format: cdsdisk	No	No	Yes
I/O fencing	No	No	Yes (see details above)
Portable Data Containers	No	No	Yes
Dynamic Multi Pathing	No	No	No
Volume Replicator	Yes	Yes	Yes
CVM/VVR	No	No	No
Bunker node	No	No	No

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

VMware

Snapshots

VMware snapshots are point in time copies of a virtual machine. Snapshots allow the administrator to go back to a specific point in time, including the memory state. VMware snapshots should not be confused with array based snapshot technologies or Volume Manager snapshot functionality.

VMware snapshots are not supported when raw device mapping is used in physical mode. This is unrelated to whether Storage Foundation is installed or not. The REDO-log functionality that is required for VMware Snapshots is not available with Raw Device Mapping - physical mode.

Raw device mapping in logical mode and VMware snapshots are supported as RDM-logical mode use the same level of SCSI virtualization as VMDK files (see

Figure 5).

VMotion (Live Migration)

VMware VMotion allows virtual machines to be migrated between physical servers without noticeable downtime. VMotion works at the Hypervisor level hence the operating system is not even aware that it has been moved to a different physical server and applications running within the operating system are typically not affected by VMotion.

Storage Foundation (standalone)

All features of standalone Storage Foundation works with VMotion. However, make sure you read the next paragraph to understand the limitations of VMotion and virtual SCSI controller sharing.

Storage Foundation Cluster File System & Cluster Server

VMware VMotion has a limitation that affects all clustering software; it 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 virtual machines that share storage between each other (on the same physical ESX server or between physical ESX servers), essentially all clustering products that rely on SAN storage require this to be set.

Storage Foundation Cluster File System requires shared storage to operate hence VMotion functionality is not available for Virtual Machines running Storage Foundation Cluster File System.

Virtual Machines that do not have the sharing attribute turned on their virtual SCSI controllers are not affected by this limitation.

NPIV (N-Port ID Virtualization)

NPIV used with Storage Foundation is fully supported. No additional setup tasks are required for Storage Foundation when using NPIV-enabled storage.

I/O fencing with NPIV is currently under evaluation.

Storage Foundation 5.0 MP3 and ESX 3.5 update 2 are minimum versions required for NPIV support.

NPIV does not enable any previously unavailable functionality such as Veritas Dynamic Multi Pathing.

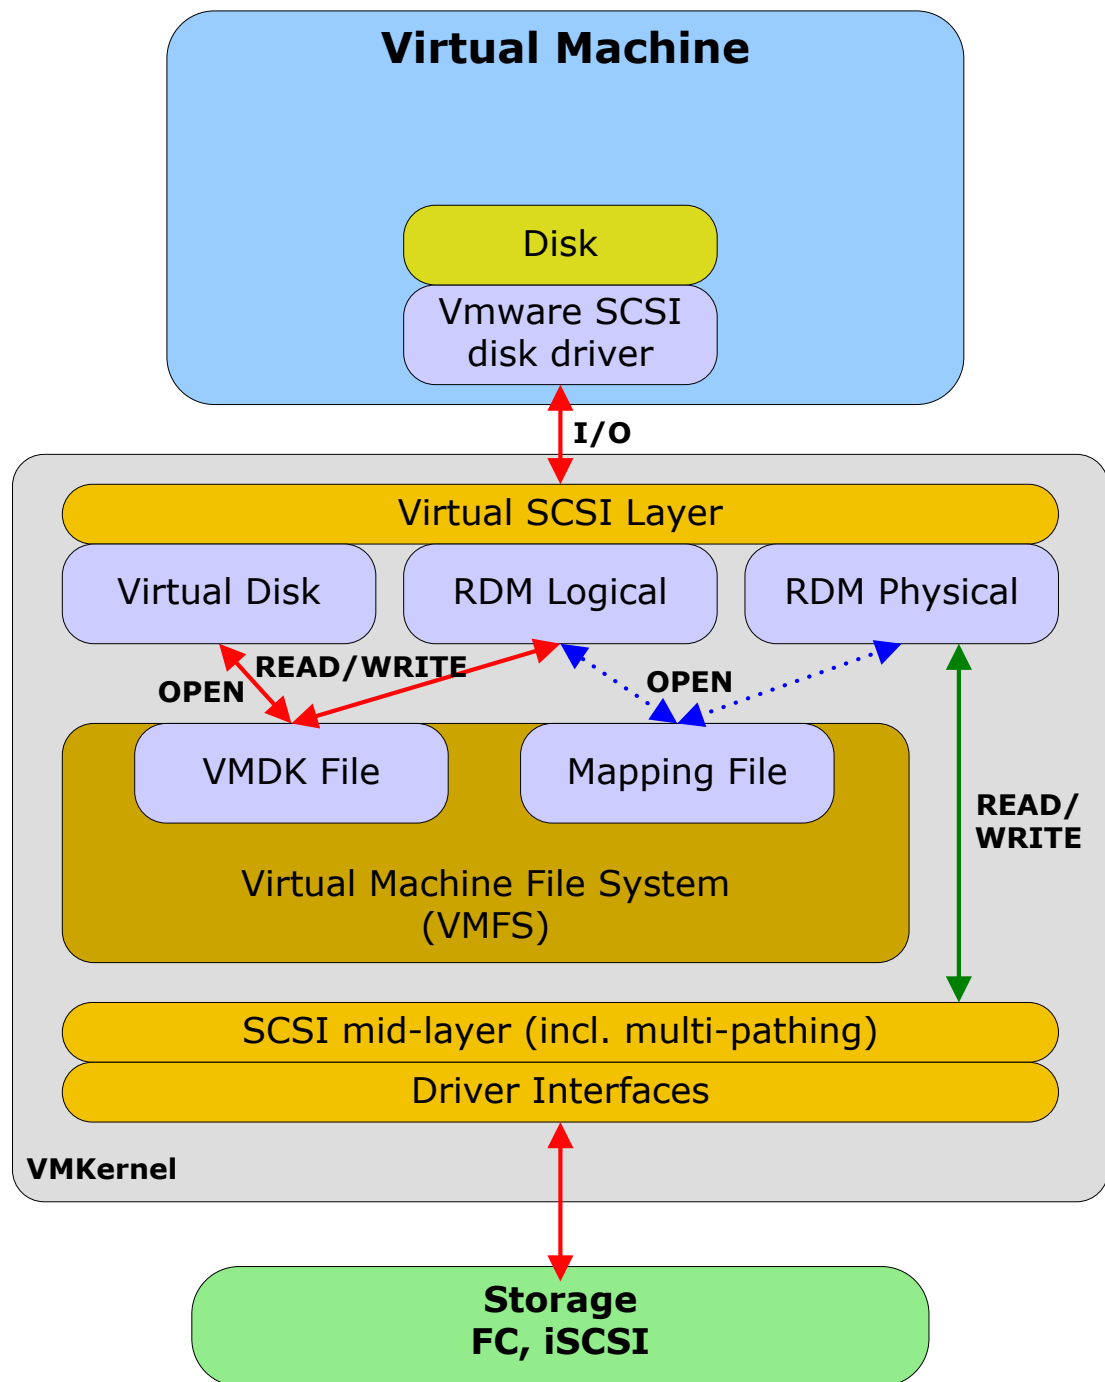


Figure 5 I/O path from Virtual Machine to Storage

Support matrix

The platforms shown in Table 2 have been tested and are supported with Storage Foundation running within the virtual machine on VMware ESX 3.0.x. Note that this paper is focused on supported Storage Foundation Linux and Solaris platforms and does not include Windows information.

VMware ESX versions earlier than 3.0.1 have not been tested and are not officially supported.

VMware ESX 3.5 update 2 was used for verification of the Storage Foundation product line; earlier versions of ESX 3.5 have not been tested and are not supported.

Future major releases of VMware ESX are not automatically supported.

Storage Foundation for Databases (Oracle, DB2 and Sybase) are supported with Storage Foundation 5.0 MP3, earlier versions of Storage Foundation have not been tested.

Later versions of maintenance packs are automatically supported unless otherwise noted. For example, Storage Foundation 4.1 MP3 was used in testing the Linux support, Storage Foundation 4.1 MP4 is automatically supported when it is released.

Operating System	SF 4.1		SF CFS 4.1	SF 5.0
	X86 (32-bit)	X86_64 (64-bit)	X86, X86_64 (32/64-bit)	X86_64 (64-bit)
Red Hat Enterprise Linux 4	Yes ¹	Yes ¹	Yes ¹	Yes ²
Red Hat Enterprise Linux 5	Yes ⁴	Yes ⁴	Yes ⁴	Yes ⁵
SUSE Linux Enterprise Server 9	Yes ¹	Yes ¹	Yes ¹	Yes ²
SUSE Linux Enterprise Server 10	Yes ¹	Yes ¹	Yes ¹	Yes ⁵
Sun Solaris 10 x64	N/A	Yes ³	Yes ³	Yes

Table 2 Support Matrix for Storage Foundation in a VMware ESX environment

1. Storage Foundation 4.1 MP3 (minimum)
2. Storage Foundation 5.0 GA (minimum)
3. Storage Foundation 4.1 MP1 (minimum)
4. Storage Foundation 4.1 MP4 (minimum)
5. Storage Foundation 5.0 MP3 (minimum)

Licensing and pricing policy

Storage Foundation pricing policy for Linux and Solaris x64 platforms change when used in a VMware virtual machine environment. The current pricing guidelines will not be repeated in this document.

Storage Foundation Standard and Enterprise

This policy covers both Storage Foundation Standard/Enterprise version 4.1 and 5.0.

For each license of the licensed Software, you may run one (1) instance of the licensed Software on one (1) physical server and an unlimited number of virtual Servers located on such physical Server

Single-node configuration examples:

1 physical node (2 CPU's):

10 virtual machines running Linux with Storage Foundation Enterprise 4.1.

This requires the purchase of 2 CPU licenses of Storage Foundation Enterprise 4.1. This purchase gives the right to run Storage Foundation Enterprise 4.1 in all virtual machines hosted on that physical server.

1 physical node (2 CPU's):

10 virtual machines all running mixed versions of Linux (RHEL 3 and 4 for example). This requires the customer to run two different versions of Storage Foundation because of the operating system mix. In this case, Storage Foundation 4.0 and 5.0 have been chosen.

This requires the purchase of 2 CPU licenses of Storage Foundation 4.0 as well as 2 CPU licenses of Storage Foundation 5.0. This purchase gives the right to run Storage Foundation 4.0 as well as 5.0 in all virtual machines hosted on that physical server.

1 physical node (2 CPU's):

10 virtual machines running Linux and Solaris x64 with Storage Foundation running in both operating systems.

This requires the purchase of 2 CPU licenses of Storage Foundation for Linux as well as 2 CPU licenses of Storage Foundation covering Solaris x64. This purchase gives the right to run Storage Foundation for both operating systems in all virtual machines hosted on that physical server.

Multi-node configuration example

2 physical nodes (Each node: 2 CPU's):

10 virtual machines all running Linux together with Storage Foundation 4.1

This requires the purchase of 4 CPU licenses of Storage Foundation 4.1 (2 CPU's per node). This purchase gives the right to run Storage Foundation Enterprise 4.1 in all virtual machines hosted on the two physical servers.

Common configurations explored

This chapter explores common configurations and determines if they are supported.

Note that even if the configuration is supported, some constraints may be present hence assuming the product will work as in a physical environment may be incorrect. The previous pages in this document has explored all the constraints for running Storage Foundation in a ESX environment and should be read in detail.

Common Configurations

I would like to run Storage Foundation with ESX 3.5 and use VMotion?

Fully supported configuration.

I would like to run Storage Foundation with ESX 3.5 and use VMotion and NPIV?

Fully supported configuration.

I would like to run Storage Foundation Cluster File System with ESX 3.5 and use NPIV?

Fully supported configuration.

I would like to run Storage Foundation Cluster File System with ESX 3.5 and use NPIV and I/O Fencing on the same physical ESX server?

Fully supported configuration.

I would like to run Storage Foundation Cluster File System with ESX 3.5 and use NPIV and I/O Fencing on two different physical ESX servers?

Fully supported configuration.

I would like to run Storage Foundation Cluster File System with ESX 3.5 and use NPIV and VMotion?

Not supported due to VMware limitations with VMotion.

I would like to run Storage Foundation Cluster File System with ESX 3.5 and use NPIV, I/O fencing and VMotion on the same physical ESX server?

Not supported due to VMware limitations with VMotion.

I would like to run Storage Foundation Cluster File System with ESX 3.5 and use NPIV, I/O fencing and VMotion on two different physical ESX servers?

Not supported due to VMware limitations with VMotion.

I would like to run Storage Foundation HA (not VCS for ESX) with ESX 3.5 and use VMotion?

Not supported due to VMware limitations with VMotion.

Does Storage Foundation have any components that run at the Hypervisor level, i.e. inside VMware ESX?

No, Storage Foundation has no components and does not require any components to run at the Hypervisor level or in VMware ESX Service Console.

Is iSCSI storage supported?

All block storage topologies that are supported with ESX 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. If I/O fencing is used then the storage must be able to handle SCSI-3 persistent reservations. The Storage Foundation Hardware Compatibility List (HCL) includes that information and the HCL can be found here:

<http://www.symantec.com/business/support/documentation.jsp?language=english&view=comp&pid=15208>

What Storage Foundation products are supported in a VMware ESX environment?

Storage Foundation and Storage Foundation Cluster File System have both been tested and are supported in VMware ESX Virtual Machines. For specifics, please read this white paper.

Are there any plans for support of older versions of VMware ESX?

There are no plans to support older versions (earlier than 3.0.1) of VMware ESX.

Is there any way to get DMP to work in a Virtual machine?

The VMware architecture does not allow for VMware Multi Pathing to be turned off or bypassed, hence there is no way to reliably enable DMP in the virtual machine.

Is Veritas Volume Replicator supported in VMware?

VVR is supported in a virtual machine with two exceptions:

1. Bunker replication with the bunker node being a virtual machine
2. CVM/VVR replication. Replication of a SF CFS cluster that is running in virtual machines

Is Storage Foundation Cluster File System supported in VMware?

Yes, Storage Foundation Cluster File System is supported running inside a virtual machine utilizing storage via raw device mapping in physical mode. Please read the details in this white paper with regards to I/O fencing limitations.

What is the pricing for Storage Foundation in a VMware environment?

See the pricing and licensing section of this document.

When will VMware ESX 3.5 be supported?

ESX 3.5 testing has been completed and it is now supported. See this document for details regarding new features such as NPIV.

Is NPIV supported?

NPIV has been tested with SF 5.0 MP3 and ESX 3.5 update 2 and is fully supported. It is not supported with earlier versions of Storage Foundation or ESX.