



VMware® vStorage Virtual Machine File System

Technical Overview and Best Practices

A VMWARE TECHNICAL WHITE PAPER UPDATED FOR VMWARE VSPHERE™ 4
VERSION 2.0

Contents

Introduction3

Background3

VMFS Technical Overview4

Features of VMFS5

Benefits of VMFS.....5

 Enables Automated Cluster File System Capability5

 Optimizes Virtual Machine Access5

 Encapsulates the Entire Virtual Machine State in a Single Directory5

 Simplifies Provisioning and Administration of Virtual Machines5

 Provides Distributed Infrastructure Services for Multiple VMware ESX Servers5

 Facilitates Dynamic Growth6

 Intelligent Cluster Volume Management6

 Optimizes Storage Utilization6

 Enables High Availability with Lower Management Overhead6

 Simplifies Disaster Recovery.....6

 Comparing VMFS to Conventional File Systems.....7

Best Practices for Deployment and Use of VMFS7

 How Large a LUN?7

 Isolation or Consolidation?8

 Isolated Storage Resources8

 Consolidated Pools of Storage8

 Best Practice: Mix Consolidation with Some Isolation!9

 Use of RDMs or VMFS?9

 About RDMs9

 Why Use VMFS?10

 Why Use RDMs?10

 RDM scenario 1: Migrating an existing application to virtual server11

 RDM scenario 2: Using Microsoft Cluster Service in a virtual environment12

 When and How to Use Disk Spanning12

 Gaining Additional Throughput As Well As Storage Capacity12

 Suggestions for Rescanning12

Conclusion.....13

 About the Author13

Glossary14

Introduction

The VMware® vStorage Virtual Machine File System (VMFS) is a high-performance cluster file system (CFS) that enables virtualization to scale beyond the boundaries of a single system. Designed, constructed and optimized for the virtual server environment, VMFS increases resource utilization by providing multiple virtual machines (VMs) with shared access to a consolidated pool of clustered storage. And it provides the foundation for virtualization spanning multiple servers, enabling services such as virtual machine snapshot, VMware vStorage Thin Provisioning (Thin Provisioning), VMware VMotion™ (VMotion), VMware Distributed Resource Scheduler (VMware DRS), VMware High Availability (VMware HA) and VMware Storage VMotion (Storage VMotion).

VMFS reduces management overhead by providing a highly effective virtualization management layer, one that is especially suitable for large-scale enterprise datacenters. Administrators employing VMFS find it easy and straightforward to use, and they benefit from the greater efficiency and increased storage utilization offered by the use of shared resources.

This paper provides a technical overview of VMFS, including a discussion of features and their benefits. It highlights how VMFS capabilities enable greater scalability and decreased management overhead, and shares best practices and architectural considerations for deployment of VMFS.

Background

In today's IT environment, systems administrators must balance competing goals: finding ways to scale and consolidate their environment while decreasing the management overhead required to provision and monitor resources. Virtualization provides the answer to this challenge. VMware vSphere™ 4 ("vSphere") enables administrators to run more workloads on a single server, and it facilitates virtual machine mobility without downtime.

A key feature of vSphere is the ability for all machines to dynamically access shared resources such as a pool of storage. VMware vCenter™ provides a management interface that can easily provision, monitor and leverage the shared disk resources. Without such an intelligent interface, the operational costs of scaling virtual machine workloads and their storage resources might affect the benefits of virtualization.

VMware has addressed these needs by developing VMFS to increase the benefits gained from sharing storage resources in a virtual environment. VMFS plays a key role in making the virtual environment easy to provision and manage. It provides the foundation for storage access to virtual machines by making available an automated CFS along with cluster volume management capabilities for the virtual environment.

VMFS Technical Overview

VMware vStorage Virtual Machine File System is a high-performance cluster file system that provides storage virtualization that is optimized for virtual machines. Each virtual machine is encapsulated in a small set of files; VMFS is the default storage management interface for these files on physical SCSI disks and partitions.

VMFS enables IT organizations to greatly simplify virtual machine provisioning by efficiently storing the entire machine state in a central location. It allows multiple instances of VMware ESX Server to access shared virtual machine storage concurrently. It also enables virtualization-based distributed infrastructure services such as VMware DRS, VMware HA, VMotion and Storage VMotion to operate across a cluster of ESX Servers. In short, VMFS provides the foundation that enables the scaling of virtualization beyond the boundaries of a single system.

The following figure shows how multiple ESX Servers with several virtual machines running on them can use VMFS to share a common clustered pool of storage.




Figure 1. VMFS Enables Multiple ESX Servers to Share Storage

Each of the three ESX Servers has two virtual machines running on it. The lines connecting them to the disk icons for the virtual machine disks (VMDKs) are logical representations of the association between and allocation of the larger VMFS volume, which is made up of one large logical unit number (LUN). A VM detects the VMDK as a local SCSI target. The virtual disks are really just files on the VMFS volume, shown in the illustration as a dashed oval.

Each ESX Server stores its virtual machine files in a specific subdirectory on the VMFS file system. When a VM is operating, VMFS has a lock on those files so that other ESX Servers cannot update them. VMFS ensures that the VM cannot be opened by more than one ESX Server in the cluster.

Each of the three ESX Servers detects the entire LUN. The LUN is a clustered volume, and VMFS provides the distributed lock management that arbitrates access, allowing ESX Servers to share the clustered pool of storage. The point of control moves from the storage area network (SAN) to the VMkernel with no loss of security.

Features of VMFS

The technical features of VMFS that make it suitable for use in a virtual environment include:

- Automated file system with hierarchical directory structure
- Optimization for virtual machines in a clustered environment
- Lock management and distributed logical volume management
- Dynamic datastore expansion by spanning multiple storage extents
- CFS with journal logging for fast recovery
- Thin-provisioned virtual disk format for space optimization
- Virtual machine-level point-in-time snapshot copy management
- Encapsulation of the entire virtual machine state in a single directory

Benefits of VMFS

As an intelligent and automated storage interface for virtual machine environments, VMFS provides both an automated CFS capability and intelligent cluster volume management functions. VMFS has a number of benefits that make it particularly well suited as a CFS for the virtual environment. It is included with vSphere at no additional cost and is tailored to VM performance patterns.

Enables Automated Cluster File System Capability

VMFS is automated and optimized for virtual machines. It allows multiple instances of ESX Server to access the same virtual machine storage. Virtual machines can be dynamically and automatically migrated between instances of ESX Server.

Optimizes Virtual Machine Access

VMFS provides the SCSI access layer for virtual machines to efficiently read and write data on the underlying disk. It uses adaptive block sizing for large I/Os, and subblock allocation for small files and directories. VMFS is rigorously tested and certified for a wide range of Fibre Channel and iSCSI storage systems, and it is optimized to support large files while also performing many small concurrent writes.

Encapsulates the Entire Virtual Machine State in a Single Directory

VMFS stores all of the files that make up the virtual machine in a single directory, and automatically creates a new subdirectory for each new virtual machine. This location is often referred to as the VMhome.

Simplifies Provisioning and Administration of Virtual Machines

VMFS reduces the number of steps required to provision storage for a VM. It also reduces the number of interactions required between virtualization administration (vSphere admin) and the storage administration team to allocate storage to a new VM. vSphere administrators appreciate the automated file naming and directory creation as well as the user-friendly hierarchical file system structure that eases navigation through the files that form the virtual machine environment.

Provides Distributed Infrastructure Services for Multiple VMware ESX Servers

VMFS provides on-disk locking that enables concurrent sharing of virtual machine storage resources across many ESX nodes. In fact, VMFS enables storage resource pools to be shared by as many as 32 ESX Servers. Furthermore, VMFS manages storage access for multiple ESX instances and enables them to read and write to the same storage pool at the same time. It also provides the means

by which VMware DRS and VMotion can dynamically move an entire VM from one ESX Server to another in the same cluster without having to restart the VM. And Storage VMotion provides the means to move a VMhome from one datastore to another without VM downtime. This capability provides a means of migrating VMs off an overcrowded datastore or to a datastore with a different performance capability.

Facilitates Dynamic Growth

As a logical volume manager, VMFS enables an interface to storage resources so that several types of storage (SAN, iSCSI and NAS) can be presented as datastores on which VMs can reside. Enabling dynamic growth of those datastores through aggregation of storage resources provides the ability to increase a shared storage resource pool without incurring downtime. With the addition of VMFS Volume Grow introduced in vSphere, a datastore on block-based storage can now be expanded on an underlying LUN that was expanded within the storage array. And VMware also enables dynamic growth of the virtual disk for guest operating systems that support this capability.

Intelligent Cluster Volume Management

VMFS simplifies administration with an intelligent interface that makes it easy to manage allocation and access of virtual disk resources, providing the ability to recognize and mount snapshot copies at the datastore or LUN level. VMFS has a volume signature that can be resignatured to manage additional but convergent copies of a given datastore on block-based storage.

Optimizes Storage Utilization

VMFS enables virtual disk thin provisioning as a means to dramatically increase the storage utilization. With dynamic allocation and intelligent provisioning of available storage capacity in a datastore, vStorage Thin Provisioning cuts down on the amount of space that is allocated but not used in a datastore.

Enables High Availability with Lower Management Overhead

VMFS enables portability of VMs across ESX Servers to provide high availability while lowering management overhead. As a CFS and cluster volume manager (CVM), VMFS enables unique virtualization services that leverage live migration of running virtual machines from one physical server to another. VMFS also facilitates automatic restart of a failed virtual machine on a separate physical server, and supports clustering virtual machines across different physical servers. File-level lock management provides the foundation needed for the multiserver virtualization that enables VMware HA, VMware DRS, VMotion and VMware Fault Tolerance (VMware FT), causing less downtime and faster recovery.

Simplifies Disaster Recovery

Because VMFS stores a VM's files in a single subdirectory, disaster recovery, testing and cloning are greatly simplified. The entire state of the VM can be remotely mirrored and easily recovered in the event of a disaster.

And with automated handling of virtual machine files, VMFS provides encapsulation of the entire VM so that it can easily become part of a disaster recovery solution. The VMFS features that are especially useful in disaster recovery include:

- Hardware independence between primary and secondary sites
- Encapsulation—all files for a VM in one place
- Robust journal file system capability for CFS metadata
- Integration of raw disk maps (RDMs) in the VMFS structure
- Resignature option to handle storage array-based snapshots

Comparing VMFS to Conventional File Systems

Conventional file systems allow only one server to have read/write access to a specific file at a given time. In contrast, VMFS is a CFS that leverages shared storage to allow multiple instances of ESX Server to have concurrent read and write access to the same storage resources. VMFS also has distributed journaling of changes to the VMFS metadata to enable fast and resilient recovery across these multiple ESX Server clusters (sometimes referred to as a VMware cluster).

On-disk locking in VMFS ensures that a virtual machine is not powered on by multiple installations of ESX Server at the same time. With VMware HA enabled, if a server fails, the on-disk lock for each virtual machine is released, allowing the virtual machine to be restarted on other physical servers. Moreover, VMFS provides virtual machine-level snapshot capability, enabling disaster recovery. VMware Consolidated Backup (VCB) leverages this feature to provide proxy backup of virtual environments.

VMFS does not have every feature found today in other CFS and CVM systems. However, there is no other CFS or CVM that provides the capabilities of VMFS. Its distributed locking methods forge the link between the VM and the underlying storage resources in a manner that no other CFS or CVM can equal. The unique capabilities of VMFS enable VMs to join a VMware cluster seamlessly, with no management overhead.

Best Practices for Deployment and Use of VMFS

This section offers some best practices, insight and experience in addressing several questions that often arise when deploying and using VMFS volumes. It is not intended to provide the definitive answer for every question, because often there is no one single right answer. The intent here is to discuss what the trade-offs and considerations are, as well as to offer some insights in choosing the answer that best fits a specific configuration.

The topics addressed are:

- How big should LUNs be made for a given VMFS volume?
- Should we isolate storage for VMs or share a consolidated pool of storage?
- Should we use RDMs or VMFS volumes?
- Should we use disk spanning? If so, are there any concerns or suggestions?
- How much of a CVM is VMFS in terms of discovery?

How Large a LUN?

The best way to configure a LUN for a given VMFS volume is to size for throughput first and capacity second. That is, you should aggregate the total I/O throughput for all applications or VMs that might run on a given shared pool of storage; then make sure you have provisioned enough back-end disk spindles (disk array cache) and appropriate storage service to meet the requirements.

This is actually no different from what most system administrators do in a physical environment. It just requires an extra step, to consider when to consolidate a number of workloads onto a single ESX Server, or onto a collection of ESX Servers that are addressing a shared pool of storage.

Each storage vendor likely has its own recommendation for the size of a provisioned LUN, so it's best to check with the vendor. However, if the vendor's stated optimal LUN capacity is backed with a single disk that has little or no storage array write cache, the configuration might result in low performance in a virtual environment. In this case, a better solution might be a smaller LUN striped within the storage array across many physical disks, with some write cache in the array. The RAID protection level also factors into the I/O throughput performance.

Because there is no single correct answer to the question of how large your LUNs should be for a VMFS volume, the more important question to ask is, "How does one determine if a certain datastore is over-provisioned or under-provisioned?" There are many performance screens and metrics that can be investigated within VMware vCenter to monitor datastore I/O rates and latency. Monitoring these metrics is the best way to determine whether a LUN is properly sized and loaded. As workload can vary over time, periodic tracking is an important consideration.

Isolation or Consolidation?

The decision whether to “isolate” or “consolidate” storage resources for virtual environments is a topic of some debate. The basic answer comes down to the nature of the I/O access patterns of that VM. If you have a very heavy I/O-generating application, then it might be worth the potentially inefficient use of resources to allocate a single LUN to a single VM. This can be accomplished using either an RDM or a VMFS volume that is dedicated to a single VM. These two types of volumes perform similarly (within 5 percent of each other), with varying read and write sizes and I/O access patterns. Figure 2 illustrates the differences between isolation and consolidation.

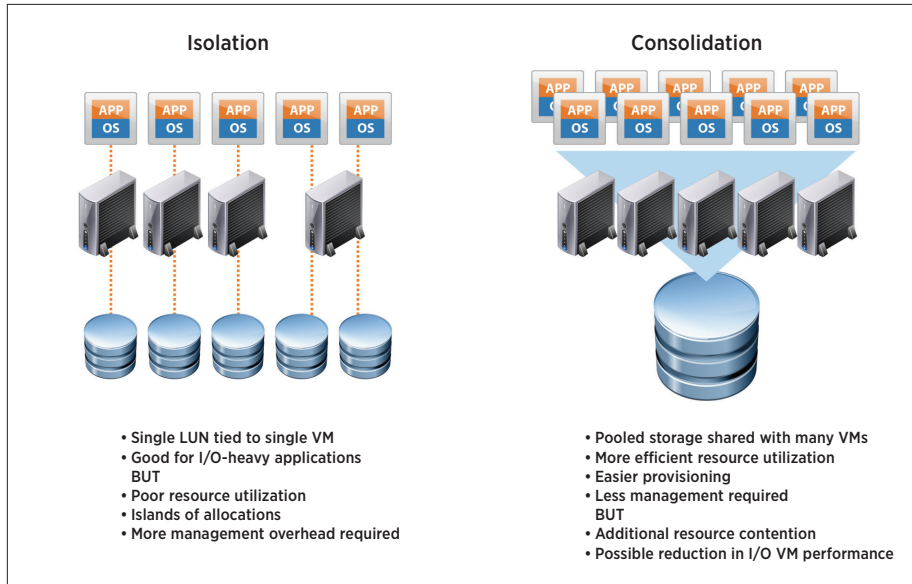


Figure 2. Consolidation Has More Advantages than Isolation

The arguments regarding isolation and consolidation are detailed as follows:

Isolated Storage Resources

One school of thought suggests limiting the access of a single LUN to a single VM. In the physical world, this is quite common. When using RDMs, such isolation is implicit, as each RDM volume is mapped to a single VM.

The downside to this approach is that as you scale the virtual environment, you soon reach the upper limit of 256 LUNs. You also need to provision an additional disk/LUN each time you want to increase storage capacity for the VM. This can lead to significant management overhead; in some environments, provisioning a disk/LUN might take days to have the request serviced by the storage administration team.

Another consideration is that every time you must grow the capacity for a VM, your minimum commit size is that allocation of a LUN. Although many arrays allow a LUN to be any size, the storage administration team might balk at carving up lots of small LUNs because this configuration makes it harder for them to manage the array. Most storage teams prefer to allocate LUNs that are fairly large; they like the system administration or applications teams to divide those LUNs into smaller chunks higher up in the stack. VMFS suits this allocation scheme perfectly and is one of the reasons VMFS is so effective in the virtualization storage management layer.

Consolidated Pools of Storage

The “consolidation” school wants to gain additional management productivity and resource utilization by pooling the storage resource and sharing it, with many VMs running on several ESX Servers. Dividing this shared resource between many VMs enables better flexibility, as well as easier provisioning and ongoing management of the storage resources for the virtual environment.

Compared to strict isolation, consolidation normally offers better utilization of storage resources. The cost is additional resource contention, which under some circumstances can lead to reduction in VM I/O performance.

At this time, there are no clear rules of thumb regarding the limits of scalability, and best practices are still being developed. For most environments, the ease of storage resource management and incremental provisioning offers gains that outweigh any performance impact. As you will see later in this paper, however, there are limits to the extent of consolidation.

Best Practice: Mix Consolidation with Some Isolation!

In general, it's wise to separate heavy I/O workloads from the shared pool of storage, to optimize the performance of those high-transactional throughput applications—an approach best characterized as “consolidation with some level of isolation.”

When looking at the number of VMs per LUN, there is no exact rule to determine the limits of performance and scalability, as workloads can vary significantly. These limits also depend on the number of ESX Servers sharing concurrent access to a given VMFS volume. The key is to remember the upper limit of 256 LUNs per ESX Server and consider that this number can limit the consolidation ratio if you take the concept of “1 LUN per VM” too far.

Many different applications can easily and effectively share a clustered pool of storage. And what little might be lost to increased contention can clearly be compensated for by the increase in disk utilization and improvements in management efficiency.

Use of RDMs or VMFS?

Another question is when to use VMFS and when to use RDMs. This section explains the trade-offs.

About RDMs

First, a little more detail about RDMs. As illustrated in Figure 3, an RDM file is a special file in a VMFS volume that manages metadata for its mapped device. The mapping file is presented to the management software as an ordinary disk file, available for the usual file system operations. To the virtual machine, the storage virtualization layer presents the mapped device as a virtual SCSI device. Key contents of the metadata in the mapping file include the location of the mapped device (name resolution) and the locking state of the mapped device.

RDM files contain metadata used to manage and redirect disk accesses to the physical device. Employing RDMs provides the advantages of direct access to a physical device while keeping some advantages of a virtual disk in the VMFS file system. In effect, the RDM merges VMFS manageability with raw device access.

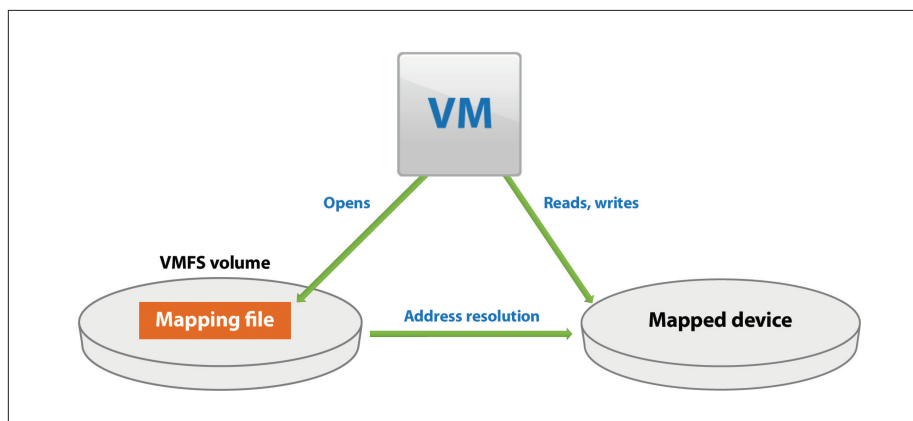


Figure 3. The RDM File in a VMFS Volume Provides a Symbolic Link to the Raw Mapped Device

Think of an RDM as a symbolic link from a VMFS volume to a raw volume. The mapping makes volumes appear as files in a VMFS volume. The mapping file—not the raw volume—is referenced in the virtual machine configuration. The mapping file, in turn, contains a reference to the raw volume.

Using RDMs, you can:

- Use VMotion to migrate virtual machines using raw volumes.
- Add raw volumes to virtual machines using the VI client.
- Use file system features such as distributed file locking, permissions and naming.

RDMs have two compatibility modes:

- Virtual compatibility mode allows a mapping to act exactly like a virtual disk file, including the use of storage array snapshots.
- Physical compatibility mode allows direct access of the SCSI device, for those applications needing lower level control.

VMware VMotion, VMware DRS and VMware HA are all supported for RDMs that are in both physical and virtual compatibility modes.

Why Use VMFS?

For most applications, VMFS is the clear choice. It provides the automated file system capabilities that make it easy to provision and manage storage for virtual machines running on a cluster of ESX Servers. VMFS has an automated hierarchical file system structure with user-friendly file-naming access. It automates the subdirectory naming process to make administration more efficient in managing RDMs. It enables a higher disk utilization rate by facilitating the process of provisioning the virtual disks from a shared pool of clustered storage.

As you scale the number of ESX Servers and the total capacity of shared storage, VMFS greatly simplifies the process. It also enables a larger pool of storage than might be addressed via RDMs. Because the number of LUNs that a given cluster of ESX Servers can discover is currently capped at 256, you can reach this number rather quickly if mapping a set of LUNs to every VM running on the ESX Server cluster.

When using RDMs, each LUN must be sized for the needs of each specific VM to which it is mapped. For this reason, using RDMs usually requires more frequent and varied dependence on the storage administration team.

With VMFS, however, you can carve out many smaller virtual machine disks (VMDKs) for virtual machines from a single VMFS volume. This enables the partitioning of a larger VMFS volume (or a single LUN) into several smaller virtual disks, which facilitates a centralized management utility (VMware vCenter) to be used as a control point. The control point resides at the ESX Server level, between the storage array and the VM.

With RDMs, there is no way to break up the LUN and address it as anything more than a single disk for a given VM. One example of this limitation is a case where a user provisioned several 50GB LUNs and wanted to test relative performance on a few VMs. The plan called for testing with 10GB virtual disks. With an RDM, the only choice was to address the entire 50GB RDM to the VM and use only the first 10GB. This wasted the other 40GB of that LUN. Using VMFS with a 50GB volume, on the other hand, enabled the creation of five directly addressable virtual disks of 10GB each on the shared VMFS volume.

Even if performance is the main consideration, you can always employ a single VMFS volume for each VM in much the same way as an RDM volume is isolated to a single VM. (Used this way, the VMFS and the RDM volumes will provide similar performance.) The bigger question is whether to isolate or consolidate, and that is not limited to use of RDMs for isolation and VMFS for consolidation.

Why Use RDMs?

Even with all the advantages of VMFS, there are still some cases where it makes more sense to use RDM storage access. Two scenarios that call for raw disk mapping are:

- Migrating an existing application from a physical environment to virtualization
- Using Microsoft Cluster Service (MSCS) for clustering in a virtual environment

RDM scenario 1: Migrating an existing application to virtual server:

Figure 4 shows a typical migration from a physical server to a virtual one. Before migration, the application running on the physical server has two disks (LUNs) associated with it. One disk is for the operating system and application files, and a second disk is for the application data.

To begin, you use the VMware utility P2V to build the VM and to load the OS and application data into the new VM.

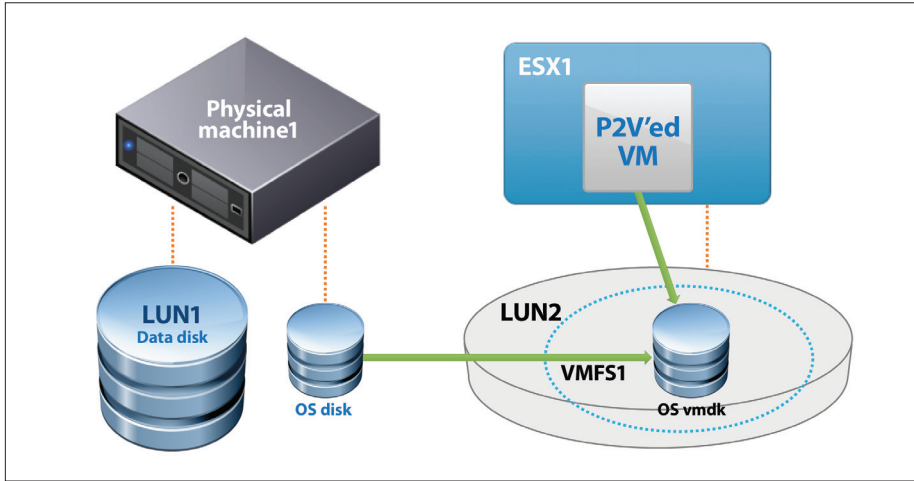


Figure 4: Use P2V to Build the New Virtual Machine and to Load the OS

Next you remove access to the data disk from the physical machine and make sure the disk is properly zoned and accessible from the ESX Server. Then create an RDM for the new VM pointing to the data disk. This allows the contents of the existing data disk to be accessed just as they are, without the need to copy them to a new location.

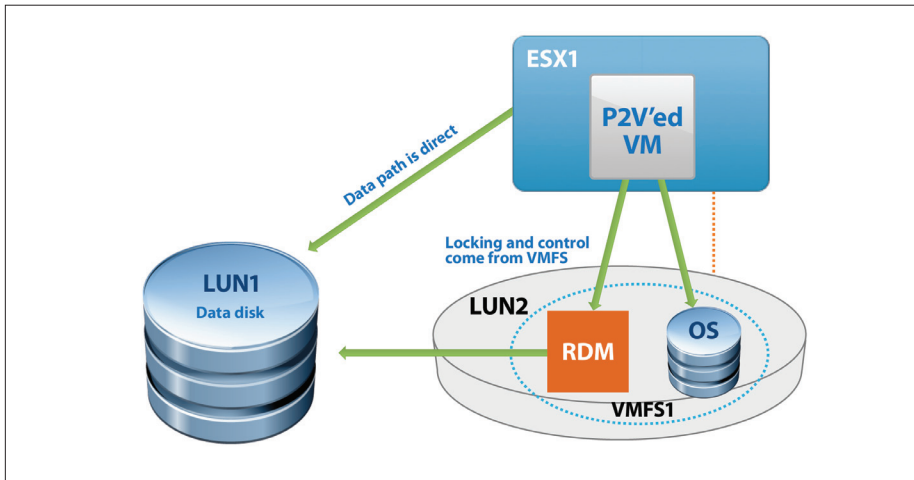


Figure 5: Then Create an RDM Pointing to the Data Disk from the Virtual Machine

The path to the data disk located on the RDM is stored in the VMFS. Although VMFS provides security access control and locking, the data path to the RDM is direct access. As with virtual disks, VMFS controls access to make sure there is no simultaneous access to the data disk from other VMs. Because RDMs enable VMotion, the VMFS can transfer control of the RDM to the destination ESX Server when a VM migrates.

RDM scenario 2: Using Microsoft Cluster Service in a virtual environment:

Another common use of RDMs is for MSCS configurations. These can use clusters in a single ESX Server (cluster in a box); clusters between virtual machines on separate ESX Servers (cluster across boxes); or clusters across both physical and virtual machines. Each of these scenarios has different requirements for shared storage, summarized in the following table.

	CLUSTER IN A BOX	CLUSTER ACROSS BOXES	N+1 CLUSTERING
Virtual disks	Yes	No	No
Passthrough RDM (physical compatibility mode)	No	Yes	Yes
Non-passthrough RDM (virtual compatibility mode)	Yes	Yes	No

Table 1. Summary of MSCS Shared Storage

Using MSCS in a virtual environment requires a disk for each node in which to store files specific to that node. Other files require shared access for the quorum disk. Those disks must support native file system access, which requires the use of RDMs in physical compatibility mode. This is another example where RDMs provide a more flexible solution for the use of virtualization technology.

For more information on MSCS configurations supported with VMware ESX Servers, refer to the VMware *Setup for Microsoft Cluster Service* documentation at: http://www.vmware.com/pdf/vsphere4/r40/vsp_40_mscs.pdf or http://www.vmware.com/pdf/vi3_301_201_mscs.pdf. More details can be found in Knowledge Base article # 1004617.

When and How to Use Disk Spanning

It’s generally best to begin with a single LUN in a VMFS volume. To increase the size of that resource pool, you can provide additional capacity by either 1) adding a new VMFS extent to the VMFS volume or 2) increasing the size of the VMFS volume on an underlying LUN that has been expanded in the array (via a dynamic expansion within the storage array). Adding a new extent to the existing VMFS volume will result in the existing VMFS volume’s spanning across more than one LUN. However, until the initial capacity is filled, that additional allocation of capacity is not yet put to use. The VMFS does not stripe I/O across LUNs when more than one LUN is allocated to a given VMFS volume.

Growing the VMFS volume on an existing, larger LUN will also increase the size of the VMFS volume but should not be confused with spanning. However, the VMFS Volume Grow can be used to expand a VMFS volume that spans a few VMFS extents and can grow a VMFS volume that spans multiple LUNs, provided there is space on those LUNs to grow the VMFS extents onto.

Gaining Additional Throughput As Well As Storage Capacity

Additional capacity with disk spanning does not necessarily increase I/O throughput capacity for that VMFS volume. It does, however, result in increased storage capacity. If properly provisioned on the underlying storage array, the additional capacity can be allocated on LUNs other than the first LUN and will result in additional throughput capacity as well. It is very important to be certain you are adding LUNs of similar performance capability (RAID level and I/O density) when adding to an existing VMFS volume.

The current size limit for a VMFS extent is 2TB. For large VMFS volumes, spanning is required to concatenate multiple 2TB extents. As there is currently a limit of 32 extents in a VMFS volume, the size limit of any VMFS volume is 64TB. Spanning of multiple volumes (LUNs) is required to reach that upper limit and is needed for any VMFS volume that is greater than 2TB in size.

Suggestions for Rescanning

In prior versions of vSphere, it was recommended that before adding a new VMFS extent to a VMFS volume, you make sure rescan is executed for all nodes in the cluster that share the common pool of storage. However, in vSphere there is an automatic rescan that is triggered so that each ESX Server updates its shared storage information when a change is made on that shared storage resource. This auto rescan is the default setting in vSphere. It can be turned off for some environments where many changes might need to be made at one time. Delaying the rescan until after those changes have been completed prevents unneeded processing time.

Conclusion

VMware vStorage VMFS provides the foundation for virtualization to span multiple systems. It enables optimized I/O handling of multiple virtual machines sharing a common clustered pool of storage. It also provides more efficient storage management for both the initial provisioning and the ongoing management of the shared storage resources. Leveraging vStorage Thin Provisioning can dramatically increase the utilization rates of your storage resources in a virtual environment. VMFS insight and monitoring are built into VMware vCenter as an intelligent interface that enables efficiency and easy management.

VMFS is the leading cluster file system for virtualization. It is unique in the market today and provides the capabilities that empower virtual machines to scale beyond the limits of a single server without downtime. VMFS enables Storage VMotion to migrate virtual machines from one storage resource to another without downtime. It also improves resource utilization and lowers the management overhead for storage administration in a virtual environment. VMFS is a key reason why VMware vSphere is today's leading virtualization solution, one that is more scalable and reliable than any other offering currently available.

About the Author:

Paul Manning is a storage architect in the Technical Marketing group at VMware and is focused on virtual storage management. Previously, Paul worked at both EMC and Oracle, where he had more than 10 years' experience designing and developing storage infrastructure and deployed best practices. He has also developed and delivered training courses on best practices for highly available storage infrastructure to a variety of customers and partners in the United States and abroad. He is the author of numerous publications and has presented many talks on the topic of best practices for storage deployments and performance optimization.

Paper Number: WP-022-PRD-01-02

You can find the most up-to-date technical documentation on our Web site at <http://www.vmware.com/support/>. The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to: docfeedback@vmware.com.

Glossary

CFS – Cluster file system

CVM – Cluster volume manager

Datastore – A formatted file system that an ESX Server mounts and uses as a shared storage pool. It is built upon either a VMFS volume for block-based storage or a mount point for NFS shared storage

DRS – Distributed resource scheduler

ESX Server – The hardware on which the vSphere software is loaded and running

Extent – Part of a larger allocation of a resource, or an appendage to an existing resource

Fibre Channel (FC) – An ANSI-standard, gigabit-speed network technology used to build storage area networks and transmit data. FC components include HBAs, switches and cabling

HA – High availability

iSCSI – Internet small computer serial interface

LUN – Logical unit number — what the server detects as a single disk resource

MSCS – Microsoft Cluster Service

NFS – Network file system (provided by network attached storage (NAS))

P2V – Physical-to-virtual conversion utility

RAID – Redundant array of independent disks

RDM – Raw disk map

SAN – Storage area network

SCSI – Small computer serial interface

Storage VMotion – A live migration of a virtual machine on disk files from one datastore to another (The VMhome changes location, but the virtual machine remains running on the same VMware ESX Server)

VCB – VMware consolidated backup

vCenter – Centralized management utility

vCenter Client – Virtual infrastructure client, a browser-based management access point

VM – Virtual machine

VMDK – Virtual machine disk, a file in the VMFS that backs the storage object for the disk for a virtual machine

VMFS-3 – Version 3 of the Virtual Machine File System

VMFS extent – The portion of a VMFS volume that resides on a single LUN

VMFS volume – The aggregation of block-based storage that backs a datastore. Normally a single LUN, but can be up to 32 LUNs when a VMFS volume spans across multiple VMFS extents

VMhome – The subdirectory in the VMFS that contains the virtual machine configuration file (.vmx) and, in most cases, all of the remaining components that form the entire state of the virtual machine

VMkernel – Core management component of the VMware ESX Server system

VMotion – A means by which the virtual machine can be moved from one VMware ESX Server to another without any downtime (VMhome will remain on the same datastore)

