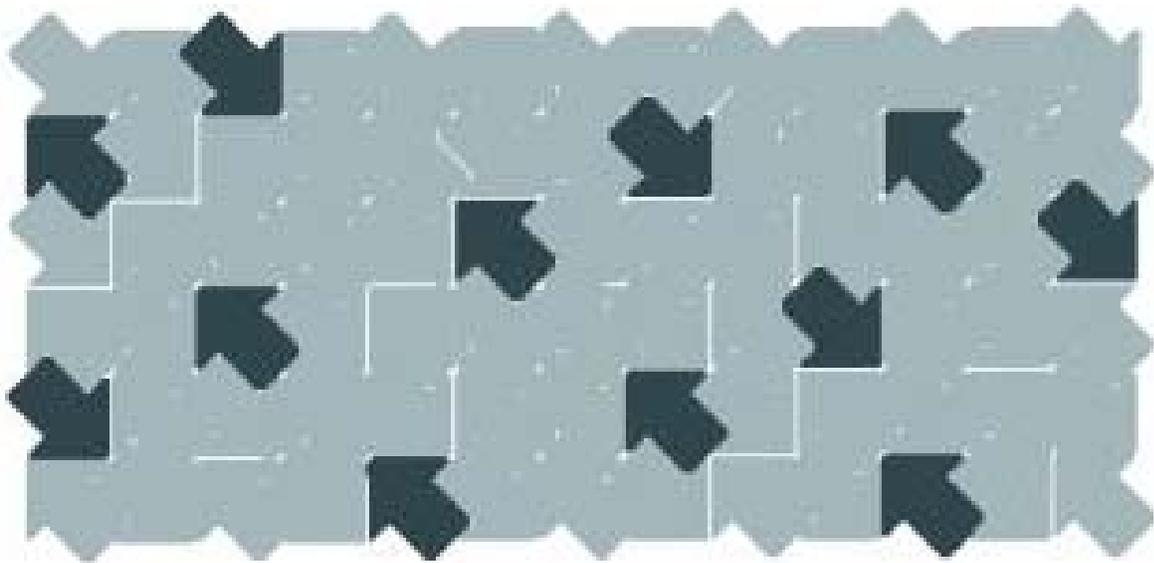


VMware Virtual Machine File System: Technical Overview and Best Practices

A VMware Technical White Paper

Version 1.0



VMware Virtual Machine File System: Technical Overview and Best Practices

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

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

© 2007 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,397,242, 6,496,847, 6,704,925, 6,711,672, 6,725,289, 6,735,601, 6,785,886, 6,789,156, 6,795,966, 6,880,022, 6,961,941, 6,961,806 and 6,944,699; patents pending.

VMware, the VMware “boxes” logo and design, Virtual SMP and VMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

VMware, Inc.

3401 Hillview Ave

Palo Alto, CA 94304

www.vmware.com

Contents

Introduction	4
Background	4
VMFS Technical Overview	5
Features of VMFS	6
Benefits of VMFS	6
Automated Cluster File System Capability	6
Optimizes VM Access	6
Encapsulates the Entire VM State in a Single Directory	6
Simplifies Provisioning and Administration of VMs	6
Provides Distributed Infrastructure Services for Multi-ESX Servers	7
Facilitates Dynamic Growth	7
Intelligent Cluster Volume Management	7
Enables HA with Lower Management Overhead.....	7
Simplifies Disaster Recovery.....	7
Comparing VMFS to Conventional File Systems.....	8
Best Practices for Deployment and Use of VMFS	8
How Large a LUN?.....	9
Isolation or Consolidation?	9
Isolated storage resources	10
Consolidated pools of storage	11
Best practice: Mix consolidation with some isolation!	11
RDMs or VMFS?.....	11
About RDMs.....	11
Why Use VMFS?	12
Why Use RDMs?	13
When and How to Use Disk Spanning?	15
Gaining Additional Throughput as Well as Storage Capacity	15
Suggestions for Disk Spanning	16
How to Use VMFS as a Clustered Volume Manager?	16
Conclusion	17

VMware Virtual Machine File System: *Technical Overview and Best Practices*

Introduction

The VMware Virtual Machine File System (VMFS) is a high performance cluster file system that allows 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 with shared access to a consolidated pool of clustered storage. And VMFS provides the foundation for virtualization spanning multiple servers, enabling services such as VMware Vmotion™, Distributed Resource Scheduler, and VMware High Availability.

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

This paper gives a technology overview of VMFS, including a discussion of features and their benefits. The paper highlights how VMFS capabilities enable greater scalability and decreased management overhead. It also provides 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 at the same time decreasing the management overhead required to provision and monitor resources. Virtualization provides the answer to this challenge. VMware Virtual Infrastructure 3 (VI3) enables administrators to run more workloads on a single server, and it facilitates virtual machine mobility without downtime.

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

VMware has addressed these needs by developing the Virtual Machine File System (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. VMFS provides the foundation for storage access to virtual machines by making available an automated clustered file system along with clustered volume management capabilities for the virtual environment.

VMFS Technical Overview

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

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

Figure 1 shows how multiple ESX servers with several virtual machines running on them can use VMFS to share a common clustered pool of storage.

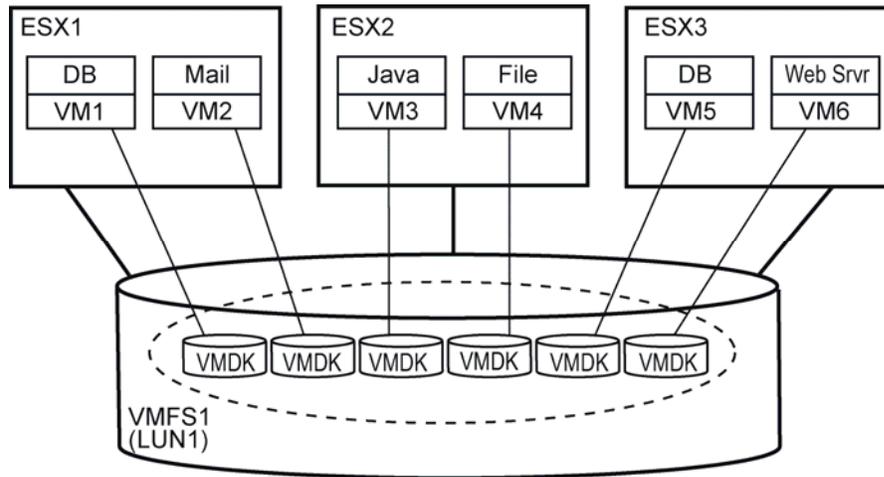


Figure 1. VMS allows multiple ESX servers to share storage

Each of the three ESX servers has two virtual machines (VMs) 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 sees 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 the VM cannot be opened by more than one ESX server in the cluster.

Each of the three ESX servers sees the entire LUN. The LUN is a clustered volume, and VMFS provides the distributed lock management that arbitrates access, allowing each VM and ESX server to share the

clustered pool of storage. Thus, 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
- Clustered file system with journal logging for fast recovery
- Encapsulation of the entire virtual machine state in a single directory

Benefits of VMFS

As an intelligent and automated storage interface for virtual machine (VM) environments, VMFS provides both an automated cluster file system capability and intelligent cluster volume management functions. VMFS has a number of benefits that make it particularly well suited as a cluster file system for the virtual environment. It is included with VI3 at no additional cost and is tailored to VM performance patterns.

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, and virtual machines can be dynamically and automatically migrated between instances of ESX Server.

Optimizes VM 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 sub-block 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 VM State in a Single Directory

VMFS stores all the files that make up the virtual machine in a single directory, and automatically creates a new subdirectory for each new virtual machine.

Simplifies Provisioning and Administration of VMs

VMFS reduces the number of steps required to provision storage for a VM. It also reduces the number of interactions required between virtualization administration (VI admin) and the storage administration team to allocate storage to a new VM. VMware administrators appreciate the automated file naming and

directory creation, and the user friendly hierarchical file system structure that eases navigation through the files making up the virtual machine environment.

Provides Distributed Infrastructure Services for Multi-ESX Servers

VMFS provides on-disk locking that allows 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 allows them to read and write to the same storage pool at the same time. It also provides the means by which VMware's Vmotion and Distributed Resource Schedule (DRS) can dynamically move an entire VM from one ESX server to another in the same cluster without having to restart the VM.

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.

Intelligent Cluster Volume Management

VMFS simplifies administration with an intelligent interface that makes it easy to manage allocation and access of virtual disk resources.

Enables HA with Lower Management Overhead

VMFS enables portability of VMs across ESX servers to provide high availability while lowering management overhead. As a clustered file system (CFS) and clustered 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. Distributed lock management and the use of SCSI reservations provide the foundation needed for the multi-server virtualization that enables Vmotion. In fact, without VMFS, Vmotion, VMware DRS and VMware HA would not be possible.

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 DR solution. VMFS is one of the key features of VI3 that lead to its

being awarded TechTarget's 2006 DR product of the year. (See [DR products of the Year](#).) The VMFS features especially useful in DR include:

- Hardware independence between primary and secondary sites
- Encapsulation—all files for a VM in one place
- Robust journal file system capability
- Integration of raw disk maps (RDMs) in the VMFS structure
- Re-signature option for snapshots

Comparing VMFS to Conventional File Systems

Conventional file systems allow only one server to have read-write access to the same file at a given time. In contrast, VMFS is a cluster file system that leverages shared storage to allow multiple instances of ESX Server concurrent read and write access to the same storage resources. VMFS also has distributed journaling to allow fast and resilient recovery across these multi-server resource pools.

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 the VM snapshot capabilities so necessary for disaster recovery, and is the interface which VMware Consolidated Backup (VCB) leverages 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/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 allow VMs to join a 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, for there is often no one single right answer. The intent here is to discuss what the tradeoffs and considerations are, and offer some insights in choosing the answer that best fits a specific configuration.

The topics addressed are:

- How big to make LUNs 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 and if so, any concerns or suggestions?
- How much of a clustered volume manager is VMFS in terms of discovery?

How Large a LUN?

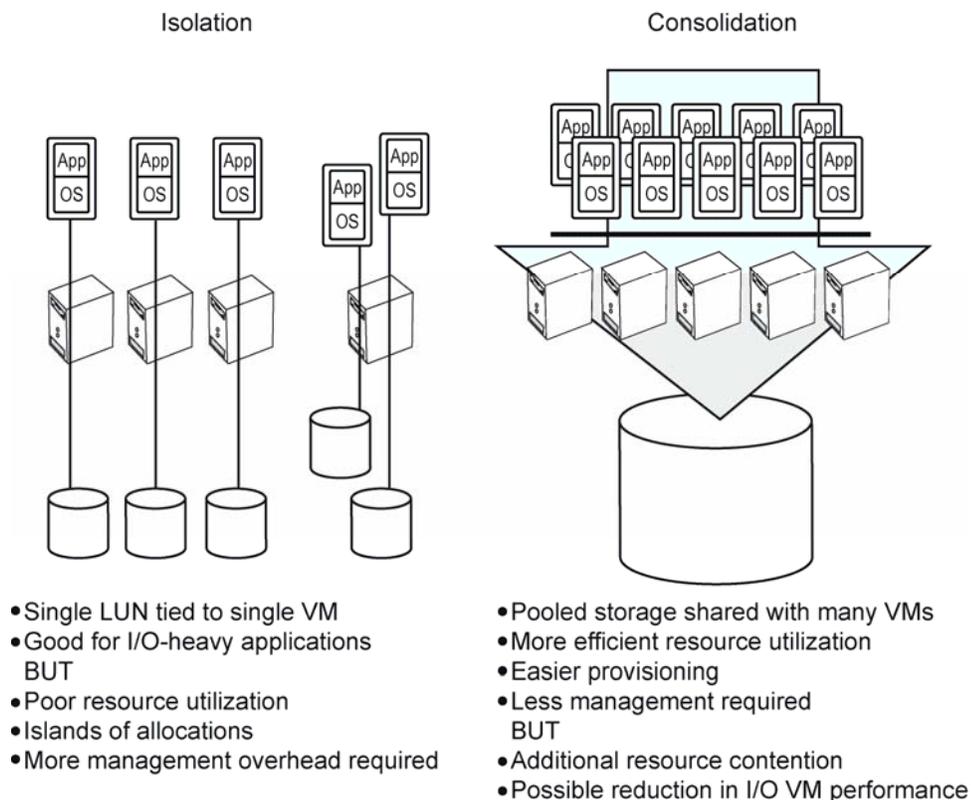
The best way to configure a LUN for a given VMFS volume is to size for throughput first and capacity second. That is, you should aggregate the total I/O throughput for all applications/VMs that might run on a given shared pool of storage, then make sure you have provisioned enough back-end disk spindles (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 consolidating a number of workloads onto a single ESX server, or onto a collection of ESX servers that are addressing a shared pool of storage.

The storage vendor likely has its own recommendation for the size of a provisioned LUN, so it's best to check with them. However, if the vendor's stated optimum LUN capacity is backed with a single disk that has no storage array write cache, the configuration could 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. Remember, there is no single correct answer to the question of how large your LUNs should be for a VMFS volume.

Isolation or Consolidation?

The decision 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 really 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. Both these 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.



- Single LUN tied to single VM
 - Good for I/O-heavy applications
 - BUT
 - Poor resource utilization
 - Islands of allocations
 - More management overhead required
- Pooled storage shared with many VMs
 - More efficient resource utilization
 - Easier provisioning
 - Less management required
 - BUT
 - Additional resource contention
 - Possible reduction in I/O VM performance

Figure 2. Consolidation has more advantages than isolation

The arguments regarding isolation and consolidation are detailed below.

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 255 LUNs. Plus, you 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 need to 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 may 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, and like to have the system administration or applications teams 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 allows 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 out of 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.”

Because workloads will vary, when looking at the number of VMs per LUN there is no exact rule to determine the limits of performance and scalability. These limits also depend on the number of ESX servers sharing concurrent access to a given VMFS volume. The key is to recognize the upper limit of 255 LUNs and understand 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 worth the increase in disk utilization and improvements in management efficiency.

RDMs or VMFS?

Another question is when to use VMFS and when to use raw disk maps (RDMs). This section explains the tradeoffs.

About RDMs

First, a little more detail about raw disk maps. 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.

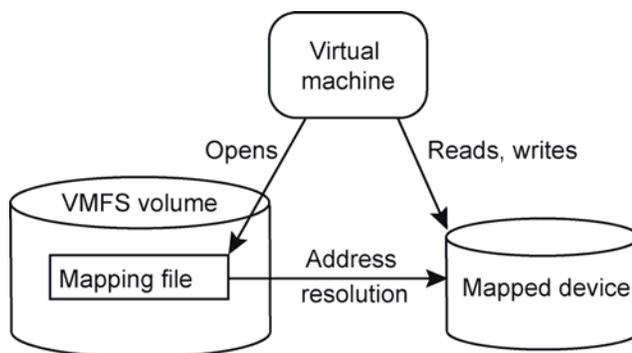


Figure 3: The RDM file in a VMFS volume provides a symbolic link to the raw mapped device.

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 in both RDM physical and virtual compatibility modes.

Why Use VMFS?

For most applications, VMFS is the clear choice. VMFS 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 raw disk maps. 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. VMFS 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 255, you can reach this number rather quickly if mapping a set of LUNs to each and every VM running on the ESX cluster.

When using RDMs, each LUN must be sized for the needs of each specific VM to which is it 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 allows the partitioning of a larger VMFS volume (or a single LUN) into several smaller virtual disks, which enables a centralized management utility (Virtual Center) to be used as a control point. The control point resides at the ESX server level, between the storage array and the virtual machine.

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, allowed 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 and every VM in much the same 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 Services (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 load the OS and application data into the new VM.

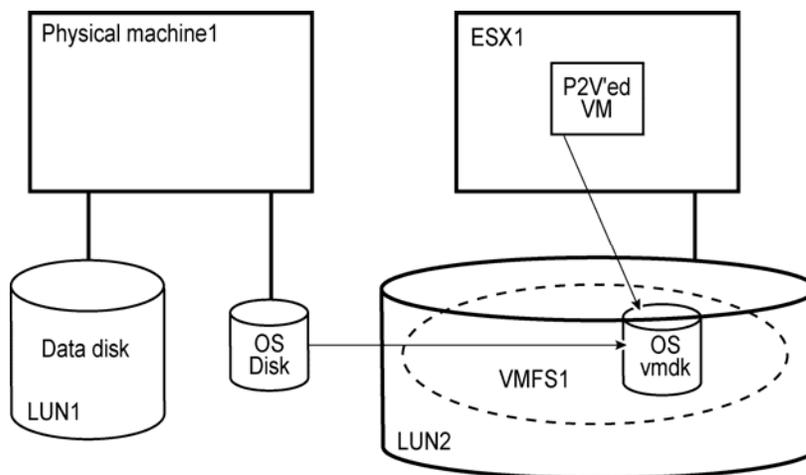


Figure 4: Use P2V to build the new VM and load the OS

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.

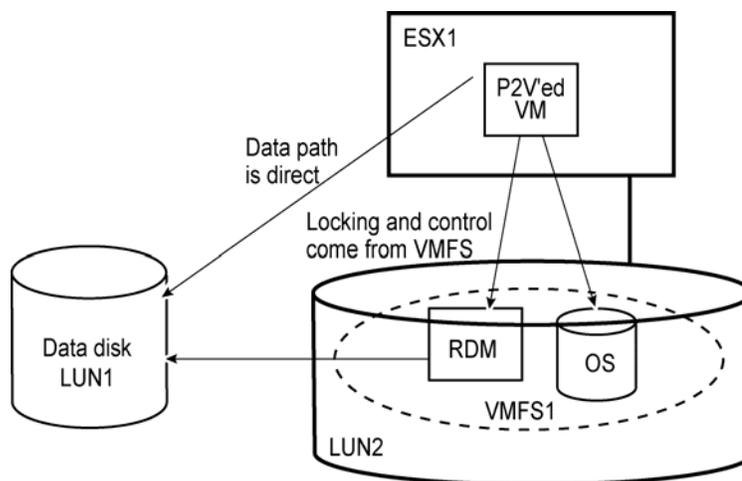


Figure 5: Then create an RDM pointing to the data disk from the VM

The path to the data disk (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 Services in a virtual environment: Another common use of RDMs is for Microsoft Cluster Service (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:

Table 1. Summary of MSCS Shared Storage

	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

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 infrastructure, refer to the VMware Setup for Microsoft Cluster Service documentation at: http://www.vmware.com/pdf/vi3_301_201_mscs.pdf)

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 add additional capacity by adding a new extent to the VMFS volume. This will add the additional LUN to an existing VMFS volume. However, until the initial capacity is filled that additional allocation of capacity is not yet put to use. Note that the VMFS does not stripe across LUNs when more than one LUN is allocated to a given VMFS volume.

Gaining Additional Throughput as Well as Storage Capacity

Additional capacity with disk spanning does not necessarily add additional 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 spindles other than the first LUN, and will result in additional throughput capacity as well. It is best to make sure 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, so 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 Disk Spanning

- When you use spanning to refresh the VMFS volume status, make sure rescan is executed for all nodes in the cluster that share the common pool of storage.
- When you want to span an additional disk, shut down the VMs that are running against the VMFS volume to which you are adding capacity. Although the datastore can be dynamically grown, best practice is to shut down the VMs because some operating systems will not handle the change in size of the underlying storage pool.
- Also, be sure to rescan after adding an extent to a VMFS volume on one ESX server. Otherwise, that same extent might be inadvertently added by another node in the cluster, which could potentially cause loss of data as well as significant confusion. Best practice is to add extents to an existing VMFS volume from a single node; then when done, rescan the storage resources from all ESX servers capable of accessing that shared storage resource.

How to Use VMFS as a Clustered Volume Manager?

When using VMFS in a large cluster it is important to understand that when a rescan is performed on a single node and a new LUN or VMFS volume is found, the other nodes are not immediately informed of its presence. This can lead to confusion and in some cases a loss of data if spanning is done for multiple nodes for the same disk.

Conclusion

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.

VMFS is the leading cluster file system for virtualization. It is unique in the market today and provides the capabilities that allow virtual machines to scale beyond the limits of a single server 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 Infrastructure 3 is today's leading multi-node virtualization, one that is more scalable and reliable than any other offering.

Glossary

CFS

Clustered file system

CVM

Clustered volume manager

DRS

Distributed Resource Scheduler

ESX server

The hardware on which the VI3 software is loaded and running.

Extent

Part of a larger allocation of a resource or an appendage to an existing resource.

Fibre Channel

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

HA

High Availability.

iSCSI

Internet Small Computer Serial Interface.

LUN

Logical Unit Number, what the server sees as a single disk resource.

MSCS

Microsoft Cluster Services.

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.

VCP

VMware Consolidated Backup.

VI Client

Virtual Infrastructure Client, a browser-based management access point.

VI3

Virtual Infrastructure 3.

Virtual Center

Centralized management utility.

VM

Virtual Machine.

VMDK

Virtual Machine Disk.

VMFS-3

Version 3 of the Virtual Machine File System.

vmkernel

Core management component of the ESX server system.

Vmotion

A means by which the virtual machine can be moved from one ESX server to another without any downtime.