Enterprise Deployment - Virtualization Considerations for Laserfiche

White Paper

February 2011

Table of Contents

Introduction

In this paper, we will explore running Laserfiche components in virtual environments. It is written for IT decision-makers and Laserfiche administrators interested in, but generally unfamiliar with, the benefits of investing in virtualization technology. After reading the paper, you should know if and how virtualization may be a valuable endeavor for your organization, and you should be equipped with a fundamental knowledge of the technology that will enable you to ask the correct questions of a virtualization vendor. The paper is not intended to be an exhaustive report on virtualization, but rather a summary of the virtualization benefits available that are relevant to a Laserfiche system.

Virtualization refers to physical abstraction of computing resources. A **virtual machine** is run on top of one or more physical machines, from which it assigned virtual components (virtual disks, virtual CPUs, virtual memory, etc.). A virtual machine can host an operating system (e.g., Windows Server 2008 R2), which can host applications (e.g., Microsoft SQL Server 2008 R2, Internet Information Services, Laserfiche Server).

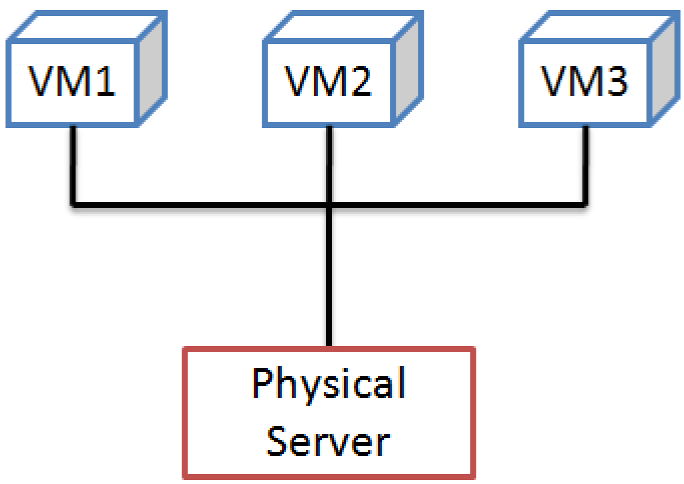

As seen in this diagram, one physical machine can host multiple virtual machines.

Vendors

In this paper, we will discuss VMware, Inc. products, as they are a leading provider of virtualization software and are used internally at Laserfiche. We encourage you to research other options, such as:

- Microsoft Hyper-V
- XEN Hypervisor (open source)

- Citrix XenServer (which is built on XEN Hypervisor)

Benefits

There are a number of obvious benefits to virtualization. You can:

- Minimize hardware costs by reducing the number of physical machines you need to purchase and maintain.
- Maximize physical office space availability, due to server consolidation.
- Reduce energy requirements. In the diagram above, only one physical machine is being powered, even though three virtual machines are being used.

  Note: VMware also offers additional power savings methods for physical machines. For more information, search vmware.com for "energy efficiency."

- Save on licensing costs. For example, a customer licensed with Windows Server 2008 R2 Enterprise may run one instance of the server software in a physical operating system environment (POSE) and up to four instances of the server software in virtual operating system environments (VOSE).
- Run a Windows or Linux application on a Mac operating system. For example, using VMware Fusion, you could run Laserfiche software within Mac OS X.

In addition, there are a number of less obvious benefits, which will be explained in detail throughout the rest of the paper.

VMware Architecture

Before exploring advanced virtualization benefits, it's important to understand the basic VMware architecture and terminology.

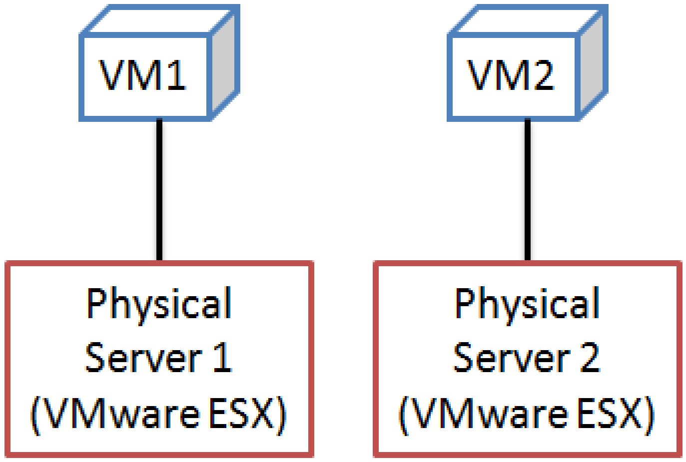

In this example, there are two physical servers. Each server is running **VMware ESX**, which is VMware's enterprise-level virtualization management operating system. VMware refers to the hypervisor used by ESX as **vmkernel.**

> **Note:** A hypervisor, also referred to as a virtual machine monitor (VMM), is the component that allows one or more operating systems to run on a host computer.
>
> **Note:** ESX is not itself a virtual machine and it does not run on top of any other operating system. It is said to be installed on "bare metal," as it contains its own kernel, which is the central component of most computer operating systems.
>
> **Note:** In mid-2010, VMware released ESXi 4.1, which is intended to replace ESX. For all practical purposes, in this paper, we will use ESX and ESXi interchangeably.

Vmkernel offers an interface to guest virtual machines that simulates hardware. A virtual machine file system (VMFS), which contains virtual machine disks (VMDKs), runs on ESX. In other words, instead of having an actual hard drive, virtual machines have VMDKs, which can be stored on a VMFS, which runs on ESX.

A VMFS is a clustered file system which can be accessed simultaneously by multiple physical servers (i.e., a VMFS can be hosted across multiple physical servers, instead of relying on a single physical server to host the entire VMFS). This allows you to take advantage of load-balancing the work, plus easily migrating virtual machines from one physical server to another without interrupting the virtual machine's operations.

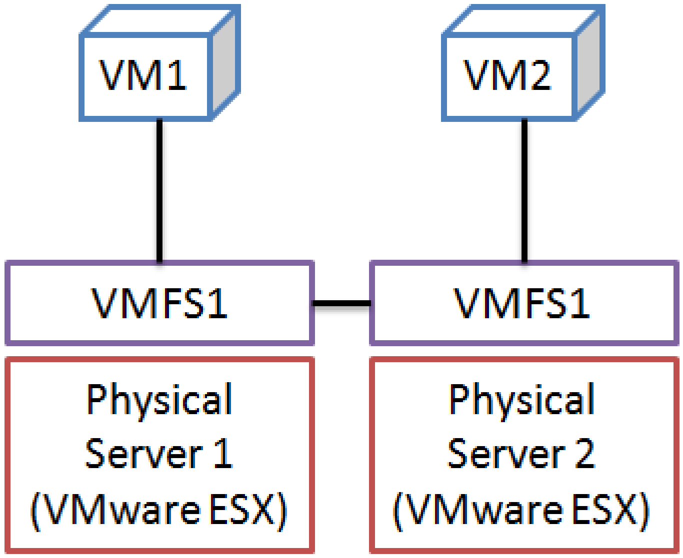

In the diagram above, VMFS1 is accessed simultaneously by Physical Server 1 and Physical Server 2.

> **Laserfiche Tie-in:** In the diagram above, VM1 could host a Laserfiche Server, while VM2 could host a Microsoft SQL Server instance that, in turn, hosts a Laserfiche repository database.

VMware vMotion

One of the key benefits to virtualization is the ability to perform hardware maintenance without causing any interruption in application service. VMware's vMotion technology allows you to migrate a live virtual machine to a new physical host with no impact on end users.

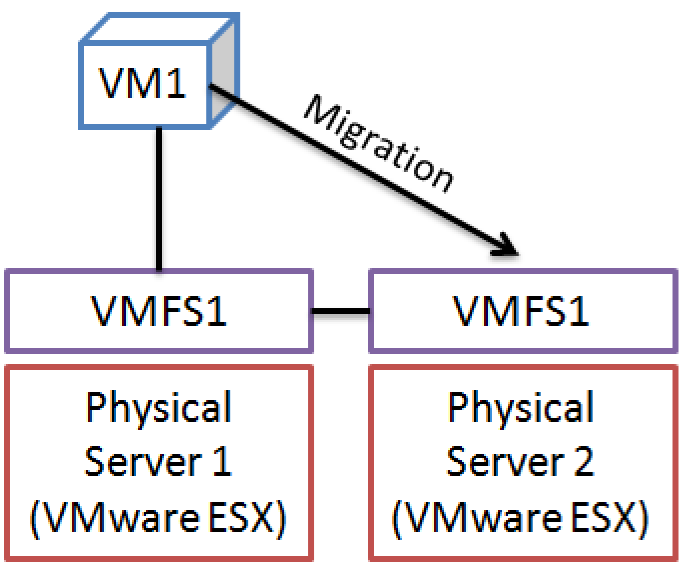

Laserfiche Tie-in: In the diagram above, VM1, which in this case is running a Laserfiche Server, is originally hosted on Physical Server 1. When an administrator needs to perform hardware maintenance on Physical Server 1, which will likely involve rebooting it, he/she can use vMotion to migrate VM1 to Physical Server 2. End users interacting with the Laserfiche Server hosted on VM1 will be unable to detect any difference in service—before, during, or after the migration.

Note: vMotion plays a role in many of the VMware technologies discussed below. For more information, search vmware.com for "VMotion."

High Availability

If your Laserfiche repository is mission-critical and must be available at all times, to the extent that downtime causes significant loss in revenue and/or productivity, we recommend investing in a **high-availability** (otherwise known as **failover**) solution. Virtualization offers a number of unique high-availability strategies, many of which are more robust than their traditional (i.e., physical) equivalents. In this paper, we will explore VMware High Availability and VMware Fault Tolerance.

VMware High Availability

If the physical machine hosting a virtual machine fails, VMware High Availability (HA) automatically moves the virtual machine to a new physical host and restarts it. For example, in the diagram below, Physical Server 2 has encountered a motherboard failure, so HA moves VM2 to Physical Server 1 and reboots it.

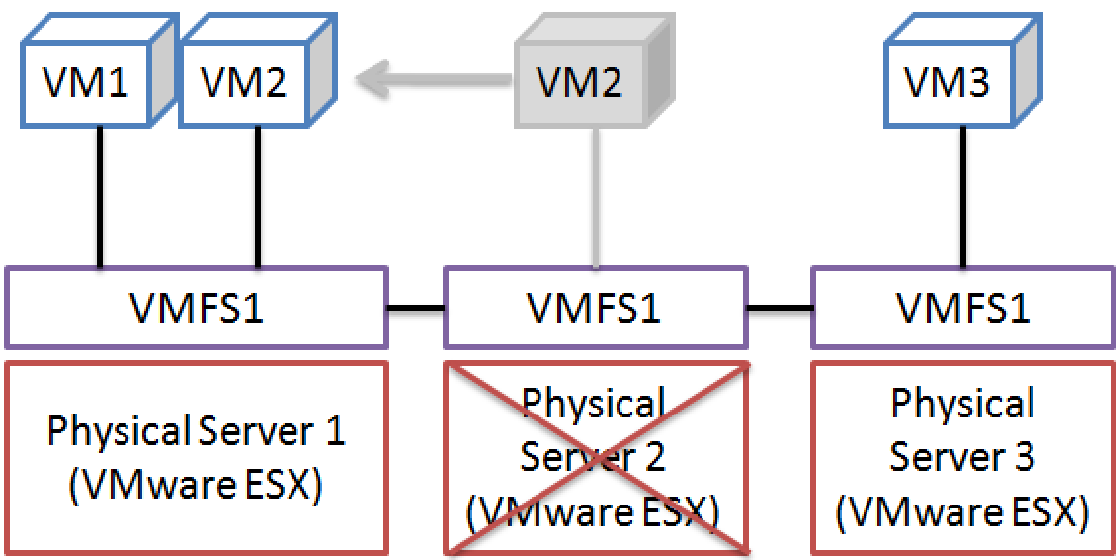

In addition, if a virtual machine's operating system fails, but its physical host does not, HA restarts the virtual machine on the same physical host.

> **Laserfiche Tie-In:** In the diagram above, a Laserfiche Server is hosted on VM2. Using HA, the organization can ensure repository access is not dependent on any one physical machine.

There will be a short amount of downtime (as VMware migrates and/or restarts the virtual machine) and any data that has not been saved (e.g., data that is stored in memory) will be lost. These limitations are overcome by VMware Fault Tolerance.

> **Laserfiche Tie-In:** If a Laserfiche Server is hosted on VM2, users logged into the Server's repositories will lose access during the migration and reboot (any unsaved work will be lost) and will have to log back in after the VM2 comes back online.

VMware Fault Tolerance

Fault Tolerance works similarly to VMware High Availability, but offers a more robust solution, as it provides virtually zero downtime and no loss of data.

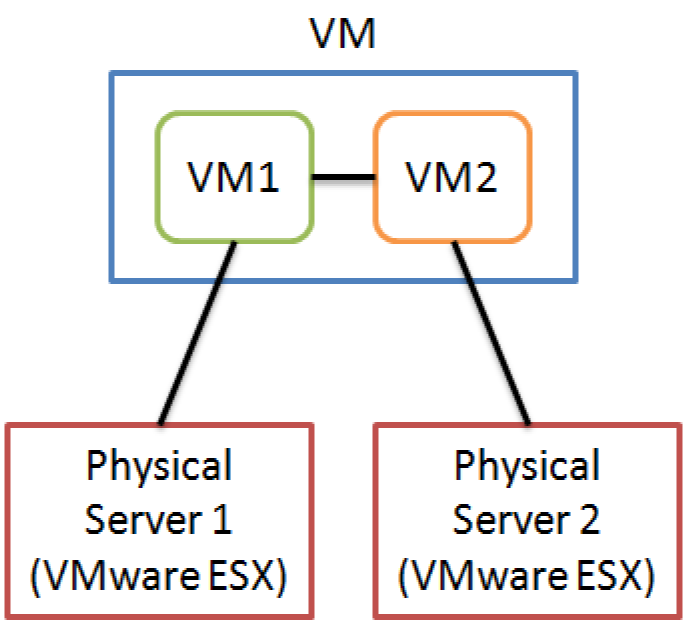

In the diagram above, there are two virtual machines (VM1 and VM2), each managed as a single unit (VM), but hosted on different physical servers (both VM1 and VM2 share the same virtual machine files). Another way to think about this technology is as a single virtual machine (VM) that is hosted concurrently on two separate physical servers.

> **Laserfiche Tie-In:** Assume that both VM1 and VM2 host Windows 2008 R2, and that this operating system hosts a Laserfiche Server.

Also assume that our goal is for VM1 to act as the primary virtual machine and for VM2 to act as the secondary (i.e., we want VM1 to handle all the work, unless it goes down, in which case VM2 should take over).

Since VMware Fault Tolerance is built directly into ESX, it is relatively simple to build failover logic into this scenario. ESX can be configured so that VM1 and VM2 run in "virtual lockstep" (i.e., vLockstep) with each other, meaning that both machines:

- **Receive all client requests as inputs** (e.g., a request from a Laserfiche Client to populate a field).
- **Process each request** (e.g., both Laserfiche Servers identify the field and determine if the user in question has sufficient security rights to perform this action).
- **Output data** (e.g., assuming the user has sufficient security, both Laserfiche Servers attempt to write the new field value to the database).

All output by VM2 is suppressed by ESX (e.g., only VM1 gets to write to the database). However, by keeping both virtual machines in virtual lockstep, both are identical at all times.

In the event that either VM1 or Physical Server 1 becomes non-operational (e.g., Physical Server 1's motherboard dies), ESX immediately makes VM2 the primary virtual machine. In addition, ESX can automatically create VM3 on a new physical server (e.g., Physical Server 3), and initiate virtual lockstep between it and VM2.

If VM1 dies, users interacting with the Laserfiche Server hosted on this machine would not lose their repository connection (or any saved work) when VM2 became the primary virtual machine.

Note: VMware Fault Tolerance can only be used on virtual machine's with 1vCPU (one virtual CPU). While the vCPU is usually not a bottleneck for Laserfiche components, it can be for Microsoft SQL Server and Oracle Database components. Keep this in mind when you consider implementing VMware Fault Tolerance. For more information on this issue, search vmware.com for "VMware Fault Tolerance, single vCPU workloads, and performance on modern hardware." For more information on VMware's Fault Tolerance solution, search vmware.com for "Fault Tolerance."

Scalability

It is often difficult for organizations to determine how much storage capacity to allocate to Laserfiche Server and repository components. On the one hand, it is vital that users be able to grow the repository as necessary, without worrying about space limitations as they scan documents, import files, add metadata, etc. On the other hand, to ensure your return on investment (ROI) is high, it is important that you do not have significant amounts of "white space" (unused hard drive space). An elegant solution to this problem is found by using virtualization and a storage area network. Both types of technology can use **thin provisioning**: a mechanism that allows space to be easily allocated on a "just-enough" and "just-in-time" basis.

Storage Area Networks (SAN)

Before explaining thin provisioning, it's important to understand the basics of SAN technology.

Few organizations with large amounts of data primarily rely on the hard disks built into server machines, as doing so, among other things, severely limits your ability to increase storage capacity quickly, easily, and without downtime.

For example, the fictional ACME Corporation has one Laserfiche volume that contains 500GB of data, all of which is physically stored on the server that hosts the volume. When the server reaches capacity, their only options are to migrate the volume to another machine with more storage capacity or to expand the storage capacity of the current server, both of which will likely require downtime and, compared to other options, more work than is necessary. Instead, the organization should invest in a storage solution that is independent of the server machine, such as a SAN.

SANs, which provide an entire network for file storage, are physical devices that contain a pool of drives representing their storage capacity. One of the primary benefits of a SAN is that its capacity can be expanded on-the-fly (by adding more drives), and many operating systems (including those from Microsoft) can recognize the capacity change without any interruption in service.

Multiple servers can use the same SAN simultaneously, and you can assign servers specific portions of the SAN, which are called **logical unit numbers** (LUNs). A LUN can be thought of as a virtual partition that defines which part of the SAN a particular server has access to. LUNs can be allocated space (x% of the total SAN capacity), which can also be expanded on-the-fly, usually without downtime. You can directly mount a LUN to a machine, enabling it to appear as a local volume (e.g., D:\).

Most SANs use RAID (redundant array of independent disks) hard drives, which offer significant resilience to disk failures and aggregation of disk capacities. While this paper will not cover RAID in detail, it is important to have a firm grasp of the following RAID features:

- **Data Mirroring:** Each drive in a RAID array has a twin drive, which contains the exact same data at all times, ensuring data redundancy.

  Note: An alternative to mirroring is **parity**.

- **Data Striping:** Servers request data from drives, and attempt to write data to them, more quickly than the drives can respond, which leads to latency. Data striping is a technique of storing a single file, for example, across multiple drives in pieces, so that a server can read/write data from multiple drives simultaneously, which results in faster performance.

  Note: Data mirroring and striping are explained here at the node level, where a group of drives sharing the same hardware components are a single **node** or **drive set**. In the Network RAID section below, these concepts are discussed in the context of mirroring and striping data across multiple nodes, which offers similar and potentially compounded advantages.

In addition, most SANs use either iSCSI or Fibre Channel to connect data storage devices over an IP network.

> **Note:** Though only briefly covered in this paper, it is important to understand the role both RAID and iSCSI/Fibre Channel play in a SAN. If you are unfamiliar with these concepts, we suggest you familiarize yourself with them.

SANs are often used with virtual machines, to store the VMFSs and VMDKs, which allows you to combine virtualization and SAN advantages. Consider the following diagram:

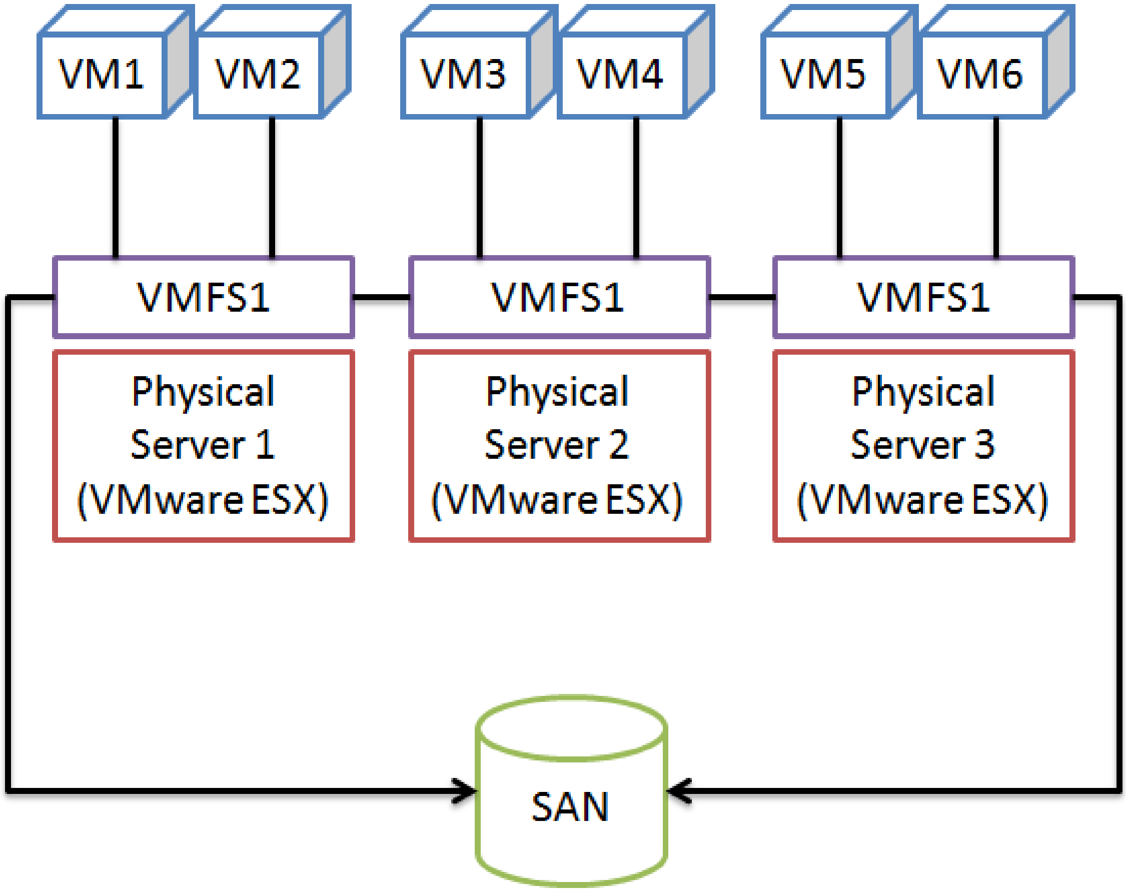

In this diagram, three physical servers are running ESX, and each is accessing the same VMFS. However, the physical servers only handle the processing actions, and the SAN is used to store all of the virtual machine's application data.

Laserfiche Tie-in: While VM1 could host the input/output actions for a Laserfiche volume, this virtual machine could reference the actual volume files off of a direct-mounted LUN that belongs to SAN.

Thin Provisioning

Thin provisioning can be implemented at both the LUN and the VMDK level.

When thin provisioning a LUN, you allocate it a certain amount of space. The LUN only uses the amount of space it currently needs, not the total allotment.

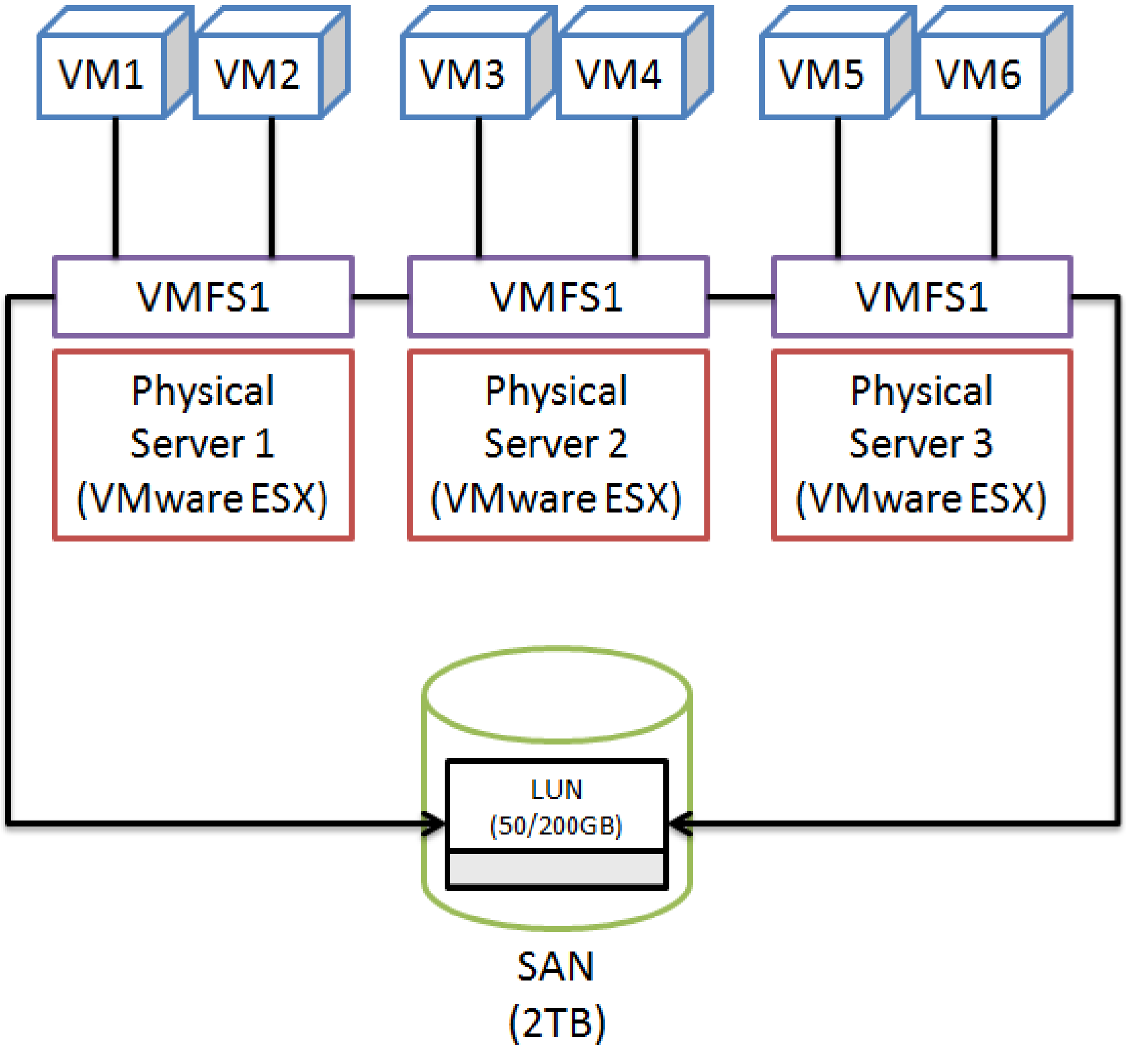

For example, in the diagram above, assume each physical server stores its portion of the VMFS on the same LUN. The LUN is thinly provisioned 200GB of the SANs 2TB of storage capacity. Immediately after standing up the system, however, only 50GB of the LUNs 200GB allotment is being used.

Does this mean the LUN has 150GB of "white space" that the SAN can't use? No. The SAN only gives the LUN the amount of space it currently needs, until it reaches its allotment, in which case an administrator must decide to allocate more space or make other changes to reduce disk use. In this scenario, since the LUN currently only needs 50GB, the SAN is free to use the other space elsewhere, as necessary.

Similarly, when thin provisioning a VMDK, you allocate it a certain amount of space, as a percentage of the storage capacity it has access to (e.g., a LUN). Similar to LUN thin provisioning, a thinly provisioned VMDK only uses the amount of space it currently needs.

It is both VMware's and our recommendation to utilize VMDK and SAN thin provisioning simultaneously, to maximize your storage efficiency. According to VMware's performance studies, and the behavior Laserfiche has observed, there are virtually no performance ramifications for doing so. Keep in mind, however, that the storage administrator must be vigilant about disk usage, as over-provisioning can become a problem. In other words, when thinly provisioning at two different layers (LUN and VMDK), there are two hard limits to be concerned with, and "insufficient space" errors, which may lead to major errors and downtime, are more likely to occur if administrators do not actively track disk usage.

Laserfiche Tie-in: For Laserfiche systems, thin provisioning can significantly help storage administrators deal with growth of Laserfiche volumes and databases. Without thin provisioning, administrators must estimate the amount of space to grant a volume/database, and also estimate how long until more space will be needed. If these estimates are incorrect, the system may either have a significant amount of unused disk space (if the estimate is too generous) or may not be able to grow as necessary until more space is purchased and added (if the estimate is too conservative). The thin provisioning approach, however, allows administrators to estimate how much space to grant, but pay no consequence if the estimate is too generous, and to easily grant more space if the estimate is too conservative.

Note: For more information, search vmware.com for "thin provisioning."

Resource Pooling

Virtualization offers robust options for resource pooling and automated workload balancing. VMware Distributed Resource Scheduler (DRS) enables you to pool different physical servers (each of which is running ESX) into a cluster. Within the cluster, you can create isolated resource pools, which combines all of the CPU and memory resources of the individual physical machines in the pool. DRS can automatically allocate pool resources across the virtual machines hosted within a pool based on need. In addition, it can automatically migrate virtual machines to different physical servers in order to improve performance.

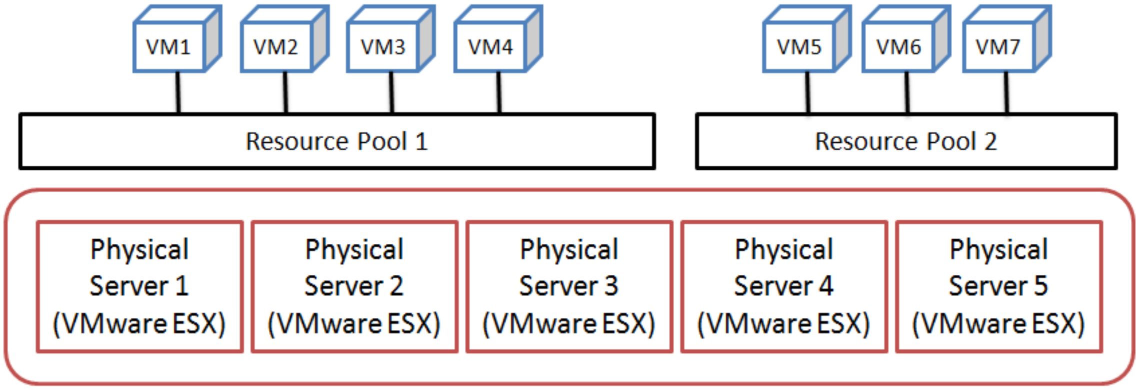

In the diagram above, two resource pools are hosted across five clustered physical servers, with each pool using a specified amount of the cluster's physical CPU and memory resources. DRS distributes each pool's allotted resources to its virtual machines based on the configuration that will lead to the highest level of performance (though not shown here, each virtual machine is still hosted on a physical machine). Throughout the day, DRS can dynamically alter this configuration as need be (e.g., granting a virtual machine more resources or moving a virtual machine to a new physical server). If you add new virtual machines or physical hosts to the cluster, DRS can dynamically migrate virtual machines and physical resources on-the-fly.

Note: If you prefer, DRS can be set to a manual mode, where it only suggests changes and relies on administrator approval before executing them (a semi-automatic mode can also be used, which works as hybrid between fully automatically and manual).

Administrators can define which physical machines belong to which resource pools (e.g., you might break up pools based on applications or departments). In addition, administrators can choose to grant a resource pool all of its physical machine's CPU and memory resources, or to reserve some for the cluster as whole that can be used as **floating capacity**. This allows a virtual machine to exceed its reserved capacity (e.g., during periods of heavy usage) by borrowing from the floating capacity, if those resources are not currently being used.

Laserfiche Tie-in: Resource pooling allows you to automate the process of ensuring your Laserfiche components have sufficient processing power to meet client's requests. It also ensures your Laserfiche system will always be running in the most efficient environment possible, based on the CPU and memory resources you've provided. For example, if VM1, which hosts a Laserfiche Server, suddenly begins to receive high traffic, DRS can migrate it to a new physical host and allocate it more processing power, to ensure it is able to handle the sudden spike in traffic.

Note: For more information, search vmware.com for "Distributed Resource Scheduler."

Network RAID

Data mirroring and striping were explained earlier at the node level, where a group of drives sharing the same hardware components are a single node or drive set. Some SANs, however, can mirror and stripe data across different nodes, which offers similar and potentially compounded advantages.

Performance Scaling

Some SANs that offer network-level RAID use striping to distribute a SAN's data across multiple nodes, instead of relying on a single node to handle all client requests. This allows different nodes to work in parallel to serve data. For example, if all of the data in a Laserfiche volume is striped across two nodes, each can work in parallel to serve the data to clients, which, in most cases, results in shorter response times.

In many cases, the performance advantages from network-level striping are compounded on top of the striping advantages that may be present at the individual RAID node level, as discussed earlier.

If a SAN vendor provides Windows Multipath I/O (MPIO) drivers, network RAID can provide read performance that scales roughly linearly with the number of nodes involved. In addition, write performance may also be improved, though this depends on the architecture of the particular SAN, and how the replication priorities are configured.

Note: For more information on MPIO, search Microsoft.com for "Microsoft Multipath I/O."

High Availability

If SAN data is mirrored across multiple nodes, the system can protect itself against single component failures in a particular node, since other nodes can provide the data.

Laserfiche Tie-in: When Laserfiche application data (e.g., volume and database files) are stored on a SAN that uses network RAID, users will see an improvement in both performance (especially in high-traffic situations) and overall system availability.

Disaster Recovery

Disaster recovery is primarily concerned with the continuation of critical applications after a disaster that leaves the primary place of business inoperable and, in some cases, causes the destruction of data. For example, if a fire destroyed an office, both the servers and their SAN could be permanently destroyed.

An offsite backup plan—which usually includes scheduled tape or DVD backups, after which the backup medium is stored offsite—are highly recommended for all Laserfiche deployments to ensure data protection. For mission-critical deployments of Laserfiche, however, protection of data alone is not sufficient, as a speedy return to service availability is also required. After a disaster occurs, waiting to stand up a new system, then restore from tape/DVD, may take much longer than the organization can afford.

It may be necessary to invest in a remote failover location to bring online almost immediately after a disaster. In order for such remote sites to be effective, they must constantly and incrementally acquire new data from the primary site in an automated fashion. In most cases, this cannot be done synchronously because the rate at which data can travel between the primary and backup site cannot keep up with the rate at which data is added to the primary storage site. Many SAN systems, however, can be configured to replicate asynchronously in the background.

Due to bandwidth costs, the more data you need to replicate, the higher the cost of the remote failover location. Therefore, it is important that you strategically design your replication plan. We recommend the following:

- If you have virtualized the server operating systems that are hosting your Laserfiche components, do not store both the operating system and the application data (e.g., repository databases and volumes) on the VMDK. If you do, you will have to regularly replicate the entire VMFS the VMDK is stored on. Instead, only store the operating system on the VMDK, and store the application data on one or more separate LUNs. Using this strategy, you do not need to regularly replicate the VMFS, as a one-time (or once-a-month, to account for operating system updates/changes) backup will suffice. The LUN(s) hosting the application data should be regularly replicated. In the event of a disaster, the VMFS can be brought online from the failover location, which contains the VMDKs that should point to the replicated application data LUN(s). Using this setup, the entire failover Laserfiche system can be operational in a very short amount of time (often within minutes).

  Note: Data added between the last replication and the disaster will not be available once the failover location is brought online.

- Only replicate mission-critical applications.

Another disaster recovery configuration is to strategically place servers and storage nodes so that they are in different locations on the same campus. For example, using VMware Fault Tolerance, the secondary virtual machine could be placed in a different building than the primary virtual machine (both machines should be on the same network). If the primary building loses power, for example, the secondary takes over. Though this does not protect against campus-wide outage, it may be a more cost-effective alternative to an offsite disaster recovery plan.

Laserfiche Tie-in: You should always have an offsite backup plan for your Laserfiche data to protect against a campus-wide disaster. If your Laserfiche system is mission-critical—meaning significant repository downtime due to a disaster cannot be afforded—we recommend investing in a remote failover location, or a separate location on-campus.

Disk Snapshots

Many SANs offer the ability to snapshot the current state of a LUN, which provides many benefits:

- If a file was permanently deleted, you could roll back to a previous snapshot and retrieve it. This concept also applies to an application that stops functioning correctly due to an unknown reason, as you could rollback to a previous snapshot where the application worked correctly.
- Since most LUN snapshots can be mounted by any server, you could mount a snapshot to a non-production server and test an upgrade or high-risk change, or perform development work. These actions could be performed on an exact clone of the production environment and data, but without affecting the live primary system.
- In some cases, snapshots can be used as a backup strategy, but ensure you also store the snapshots on media not attached to the SAN, to protect against situations where the entire SAN is destroyed.

Note: Many SANs can create differential snapshots, which can significantly decrease the amount of disk-space a snapshot uses.

Laserfiche Tie-in: Below is a diagram that illustrates these benefits. VM1 hosts a Laserfiche Server, which references a Laserfiche repository database and volume on Original LUN. Two snapshots of this LUN have been created: Testing Snapshot and Backup Snapshot. The former is attached to a test server (which also hosts a Laserfiche Server), where development or testing work can be done on an exact clone of the repository's data.

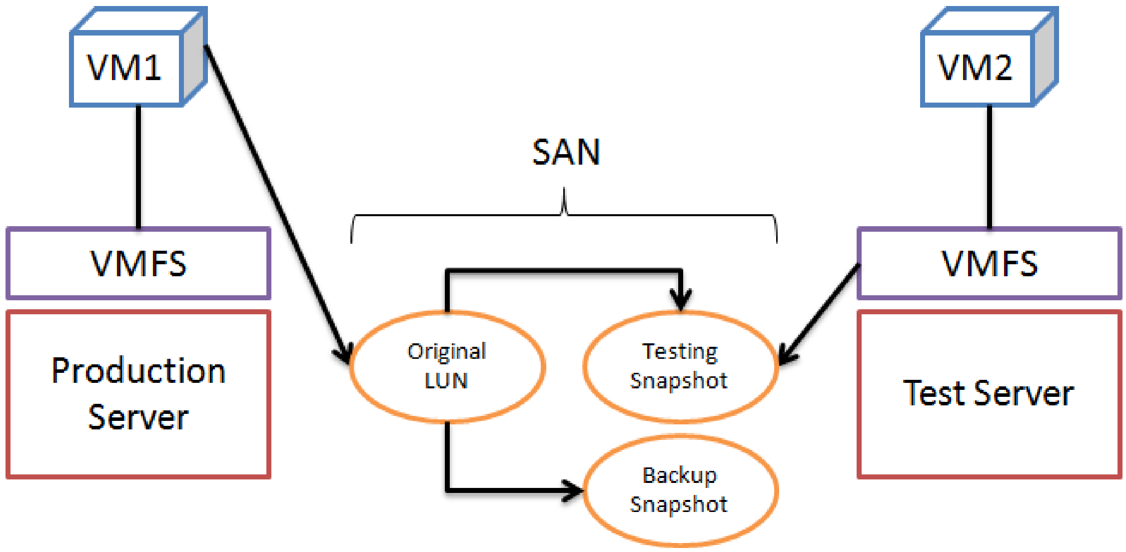

Note: In most cases, snapshots can be created manually, or you can schedule the SAN to create them.

Note: In addition to snapshoting a LUN using a SAN, VMware enables you to snapshot a virtual machine, which can also be helpful.

Drawbacks

When considering virtualization, it is important to keep in mind the following drawbacks:

- Performance: Software running on a virtual machine will always be slower than if installed directly on a physical machine. In many cases, this can be compensated for by purchasing more robust hardware for machine hosting the virtual machine.
- Training: Virtualization and SAN set up and maintenance often require extensive internal training.

Best Practices

While the needs and resources available at every organization will vary, Laserfiche offers a few general recommendations as best practices for running Laserfiche in a virtualized environment.

Hardware Considerations

In most cases, you should have more than one processor per physical machine hosting a virtual machine, and the more the better. The same is true for RAM. In general, the more powerful the host machine, the more robust its virtual machines will be. You should always meet or exceed the hardware requirements provided by the virtualization vendor.

> **Tip:** As explained earlier, you can cluster physical machines together and pool their resources to host virtual machines. If you take one cluster node offline (e.g., to perform maintenance), this will not cause any of the cluster's virtual machines to become unavailable, as the remaining cluster nodes will take over the work originally assigned to the offline node. In this situation, the nodes that remain online must have enough CPU and memory resources to run all of the virtual machines. As a result, when making hardware considerations, ensure the cluster will have enough resources to service all of the virtual machines in the event a node goes offline.

Software Considerations

As explained earlier, we recommend using VMware software, which has proven capable through extensive use both internally and in customer configurations. In addition, a SAN is essential.

> **Example:** At the time this paper was posted, the Laserfiche Information Technology department uses six physical servers and a SAN, each with dual quad-core processors and between 48 to 96 GB of RAM running ESXi 4.1. On that setup, we are running around 200 virtual machines—everything from low usage test virtual machines to key infrastructure components, such as an enterprise mail server.

Hosting Laserfiche Components

When working in a virtual environment, we recommend configuring the Laserfiche Server and SQL Server on separate virtual servers (located on the same or different physical servers). The application data (listed below) should be stored on a direct-mounted LUN:

- 1: Repository directory and audit logs (Laserfiche Server)

Note: These two components can be separated from each other, if necessary.

- 2: Volume files (Laserfiche Server)
- 3: Database (MDF) files (SQL Server)
- 4: Transaction (LDF) logs (SQL Server)

The following diagram shows a configuration with the two virtual servers (running a Laserfiche Server and a Microsoft SQL Server, respectively) installed on one physical server, and all application data stored on the SAN (each orange circle represents a different LUN; see the number scheme above).

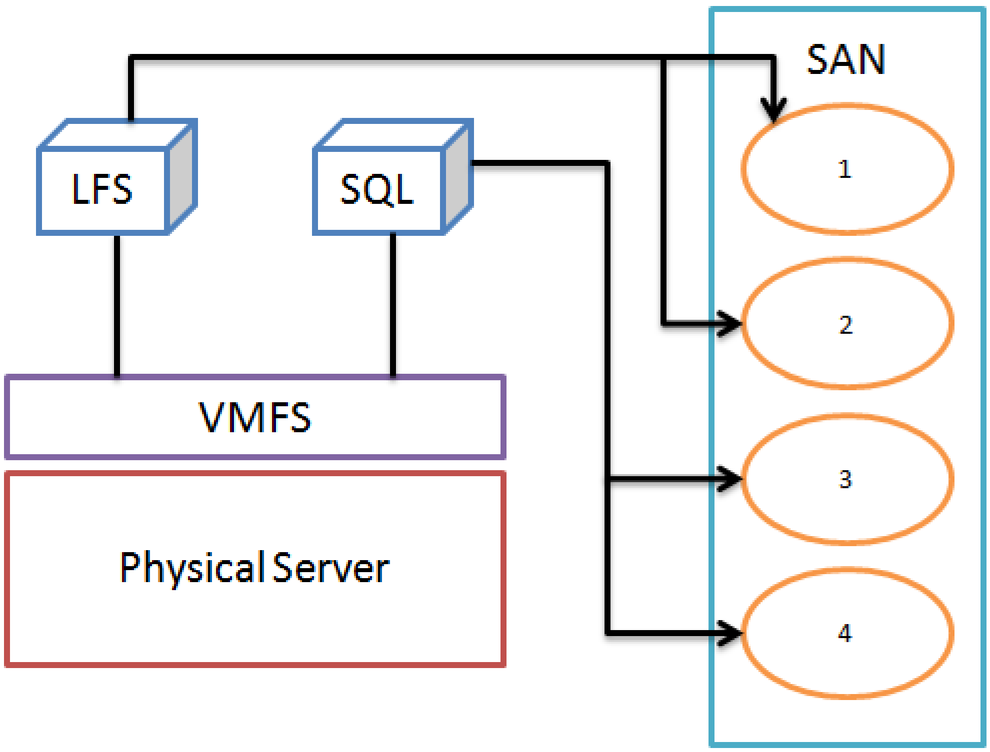

Note: In very high traffic situations, Microsoft SQL Server or Oracle Database can become a significant bottleneck. If you encounter this problem, we recommend the database server run in an un-virtualized environment on a dedicated high-performance input/output (I/O) subsystem.

Laserfiche®

Enterprise Deployment - Virtualization Considerations for Laserfiche
February 2011

Author: Jonathan Powers
Editor: Tammy Kaehler
Technical Editors: Jeff Huang, Masashi Ito, and Trent Strong

Description:
In this paper, we will explore running Laserfiche components in virtual environments. It is written for IT decision-makers and Laserfiche administrators interested in, but generally unfamiliar with, the benefits of investing in virtualization technology. After reading the paper, you should know if and how virtualization may be a valuable endeavor for your organization, and you should be equipped with a fundamental knowledge of the technology that will enable you to ask the correct questions of a virtualization vendor. The paper is not intended to be an exhaustive report on virtualization, but rather a summary of the virtualization benefits available that are relevant to a Laserfiche system.

Laserfiche
3545 Long Beach Blvd.
Long Beach, CA 90807
U.S.A

Phone: +1.562.988.1688
www.laserfiche.com

Laserfiche is a trademark of Compulink Management Center, Inc. Various product and service names references herein may be trademarks of Compulink Management Center, Inc. All other products and service names mentioned may be trademarks of their respective owners.

Laserfiche makes every effort to ensure the accuracy of these contents at the time of publication. They are for information purposes only and Laserfiche makes no warranties, express or implied, as to the information herein.

Copyright © 2011 Laserfiche
All rights reserved