

Enterprise Deployment: Laserfiche 8 in a Virtual Environment

White Paper

August 2008

Laserfiche®

The information contained in this document represents the current view of Compulink Management Center, Inc on the issues discussed as of the date of publication. Because Compulink must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Compulink, and Compulink cannot guarantee the accuracy of any information presented after the date of publication.

This chapter is for informational purposes only. COMPULINK MAKES NO WARRANTIES, EXPRESS OR IMPLIED, AS TO THE INFORMATION IN THIS DOCUMENT.

Table of Contents

Introduction	3
Benefits	3
Drawbacks.....	4
Hardware Considerations	4
Best Practices	4
Case Study: Virtual Clustering for Marina Medical Billing Company	6

Introduction

Virtualization means simulating the interface of a server, a desktop computer, a storage device, an operating system, or even an application. Virtualizations are run on physical servers, computers, and storage devices, but the effect of virtualization is to hide the physical characteristics of the host machine.

There are many benefits to the practice of virtualization, but the fundamental difference between running an application in a physical environment (e.g., on a server machine) and running in a virtual environment is that in a virtual environment, you can have multiple virtual resources on one physical resource. In addition, you can also pool multiple physical resources to make them appear as a single virtual resource.

In practice, this means that if you estimate you require 20 servers to handle the needs of an organization, instead of purchasing 20 separate, physical server machines, you could purchase two or three larger servers and set up 20 or more virtual machines. What the virtual approach also gives you is high availability, scalability, portability, automatic failover, and increased flexibility for hardware upgrades, desktop management, and prioritizing resource allocation.

Some of the software products available for virtualization are VMware (used and recommended by Laserfiche), Microsoft Virtual PC, and XEN hypervisor (opensource).

Benefits

- Failover: virtualization management software can move virtual machines from one physical server to another (“hot standby”) in the event of failure
- Resource efficiency: resources tend to be underutilized by up to 80%; with a VM you can be sure you’re using the full capacity
- Easy hardware upgrades: move the VM to new physical machine with no reinstall of software needed (depends on implementation)
- Disaster recovery: in the event of complete physical server loss, with VM backup offsite, all you need to do is replace the physical server, reinstall the VM software, and restore the VM from backup; no need to reinstall multiple applications and configure them
- Increased security management: have users work on VMs on top of a secure base, which gives you the ability to manage potential security breaches—to “wipe” and rebuild the VM for the user in minutes, without having to rebuild all the software on the physical machine
- Dynamic allocation of resources: allocate more memory on physical machine to VM that will take more load

- Licensing savings: Windows Server 2003 license for 1 physical machine or 3 virtual servers; Microsoft Data Center Edition server license for unlimited virtual servers
- Less expensive hardware and savings of physical space: three large machines instead of as many as 70 smaller ones
- High availability with failover and dynamic resource allocation
- Portability: move VM from one physical machine to another with **no interruption of service**
- Scalability: can resize hard disks on the fly and use more of the SAN resources
- Save energy: save power by not keeping low-use physical machines running all the time
- Easier to manage: only working from one physical machine with the management software
- Run PC-only software on a Macintosh through a virtual machine
- Software not competing for resources

Drawbacks

- Performance: software running on a VM will always be slower than if installed directly on a physical machine
- Cost: need to invest in infrastructure, extra hardware, license fees, and VM software
- Consolidation and centralization: all your eggs in one basket if all servers fail

Hardware Considerations

In general, the more capacity you can provide for your virtual environment, the better off you will be. If you're going to run more than one VM per physical machine (and even if you're not), you will want more than one processor per physical server, and the more the better. The same is true for RAM.

Best Practices

While the needs and resources available at every installation will vary, Laserfiche offers a few general recommendations as the best practices for running Laserfiche in a virtualized environment. We recommend starting with hardware clustering of server machines with multi-core processors and lots of RAM. On those we recommend running VMware software, which has proven capable through extensive use both internally and in customer configurations.

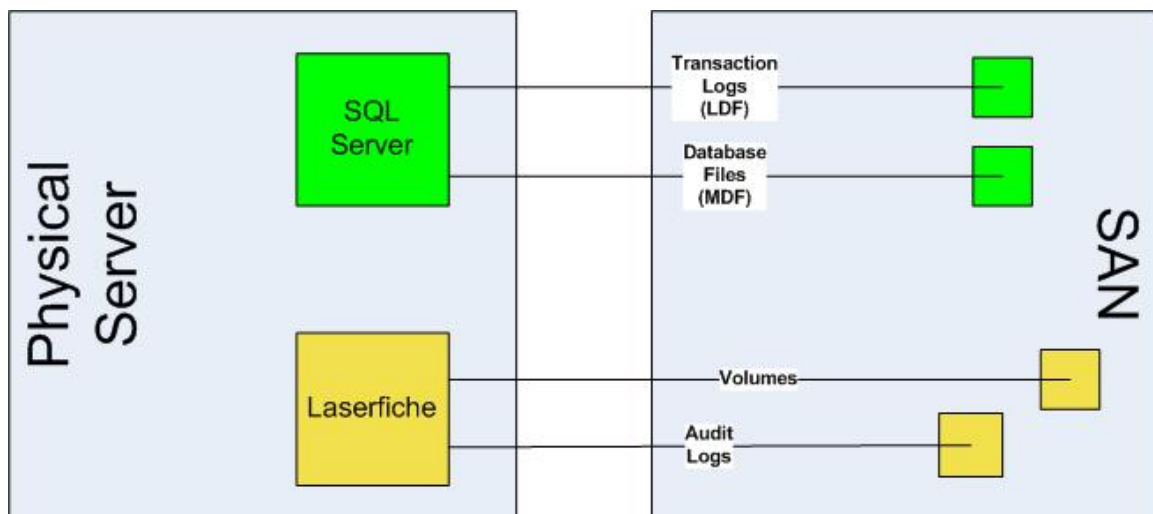
In addition, a SAN (Storage Area Network) is essential for the ability to move VMs while they are running. The SAN serves as a hard drive for everything attached to it and allows you to centrally manage disk usage.

Example: The Laserfiche Information Technology department uses three servers and a SAN, each with dual quad-core processors and 32GB RAM, running VMware ESX server virtualization software. On that setup we are running some 66 virtual machines, everything from low usage test VMs to key infrastructure components, such as the enterprise mail server.

In addition, we recommend a modification to the standard file storage architecture of the Laserfiche Server and SQL Server. In Laserfiche 8, supporting files—such as audit logs—can be stored separately from the Laserfiche Server. 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 data files should then be stored outside of the virtual servers on a network storage device. Specifically, these files are

- Audit logs (Laserfiche Server)
- Volume files (Laserfiche Server)
- Database (MDF) files (SQL Server)
- Transaction (LDF) logs (SQL Server)

Separating the files from the system components will maximize the processing resources available to the software and improve performance. The following shows a configuration with the virtual servers installed on one physical server, and all supporting files stored on the storage area network.



Case Study: Virtual Clustering for Marina Medical Billing Company

Marina Medical Billing Company came to Laserfiche looking for a document management solution, and asking for a rigorous stress test to prove capability. Laserfiche assembled a three-company project test team and used a test bed of hardware and software that is available to the average Laserfiche customer.

Key features of the test bed architecture included:

- Multiple physical servers, each running VMware software and hosting VMs (for failover capability)
- VMware Distributed Resource Scheduler (DRS) software (to monitor and dynamically reallocate computing resources across cluster)
- SAN (to add storage devices to cluster)

Hardware

Application hosting	<ul style="list-style-type: none"> • Three HP® DL380 G5 servers, each with <ul style="list-style-type: none"> ○ 16GB RAM ○ Two Intel® Xeon® 2.7GHz processors with hyperthreading
Storage	<ul style="list-style-type: none"> • HP modular storage array 1500 SAN (storage area network) with two drive arrays • Set of fault-tolerant HP 4/8 SAN switches
Test bed	<ul style="list-style-type: none"> • 12 Dell® Inspiron® 1300 notebooks, each with <ul style="list-style-type: none"> ○ Intel Pentium® M 1.7GHz processor ○ 1GB RAM

Software

Server 1	<ul style="list-style-type: none"> • VMware® ESX server® virtualization software • Laserfiche Server 7.2 • Microsoft® Windows® server 2003 R2 (3.75GB RAM) • VMware VMotion® software for failover
Server 2	<ul style="list-style-type: none"> • VMware ESX server virtualization software • Microsoft SQL Server™ 2005 • Microsoft Windows server 2003 R2 (3.75GB RAM) • VMware VMotion software for failover
Server 3	<ul style="list-style-type: none"> • Used for virtualization management
Test bed	<ul style="list-style-type: none"> • Windows XP SP2

Methodology

The test database was created in four days by replicating a single six-page document many times over many nodes of a folder tree to form a single 10GB volume. This volume was exported and copied 313 times, then each one was attached again to the Laserfiche database via automated batch files. When ready for testing, the database contained 9,610,672 documents with a total of 57,629,838 pages.

In the test bed, two notebooks were dedicated to the constant scanning of 1,500 images via Laserfiche's Quick Fields and Universal Capture software and the writing of those images to the database. The remaining 10 computers all ran a Visual Basic script that queried the database for a random string of characters and retrieved the matching document or documents. The script was set to make 500 queries or search requests, pause briefly, and then run again in a continuous loop. Twenty instances of the script ran on each notebook to simulate 200 users.

The demonstration for Marina Medical consisted of an hour of continuous full-speed and full-volume processing. To prepare for this demonstration, the project team ran the scripts in the test bed continuously for approximately 24 hours. Microsoft Windows Server 2003 performance counters were used to monitor the test and the demonstration.

In the end, the Laserfiche system performed flawlessly and far exceeded the test requirements established by the customer.



Laserfiche 8 in a Virtual Environment
August 2008

Author: Tammy Kaehler
Editors: Constance Anderson and Pieter Nauta

Compulink Management Center, Inc.
Global Headquarters
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.

Copyright © 2008 Compulink Management Center, Inc.
All rights reserved