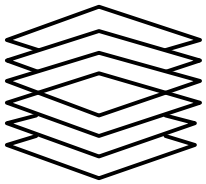

VMware: Tool for Server Consolidation

EXECUTIVE SUMMARY

Server consolidation holds out a promise of major improvements, which range from greater availability and expanded memory to more current and more efficient ways to break up work by task or by group. VMware™ GSX and ESX Server address many of these issues and offer an exciting new tool for all those who are planning consolidation projects in the IA-32 environment.

The "virtual machine" notion offers a number of advantages to server consolidation, which are described in more detail in the body of this report. Greater availability is a major goal of consolidation. VMware offers a direct benefit. Through isolation of applications into their own virtual machines, VMware allows the consolidated system to create the division of tasks that is so attractive to server consolidation teams. While this sounds like fault tolerance or partitioning, it actually means that "guest operating systems" cannot bring down the "host operating system," the other virtual machines or the VMware Server software itself. VMware also supports encapsulation in which the state of a world can be saved to disk and restarted at another time.

These benefits and many others are opening considerable interest in the virtual-machine concept and VMware in particular. Throughout this Server Consolidation Series, D.H. Brown Associates, Inc. (DHBA) and its research partners, who in this case include Aurema, Compaq, IBM, Intel, Microsoft, Sun, Unisys, and some large user companies, recommend a comprehensive plan that identifies all elements of consolidation. With the right planning, VMware can play a central role in achieving the cost and service goals of server consolidation.

INTRODUCTION

The virtual-machine concept is not new. IBM mainframes have delivered a product – the VM operating system – that has embodied this technology for several decades. Recently, VM started enjoying a renaissance because it offers a way to run multiple Linux systems on the mainframe. In other words, a large number of Linux machines can be consolidated on the mainframe using VM.

VMware, Inc. has adapted this concept for IA-32 architecture systems. VMware started with a workstation product that gained wide acceptance. It recently released a sever version of its technology. Future server versions are also planned.

TABLE OF CONTENTS

SERVER PRODUCT ARCHITECTURE

A few definitions provide a good starting point. Three terms are critical to VMware: Virtual Machine; Guest OS; and Host OS.

Most central is the "virtual machine." A virtual machine is a way of representing a real machine using software to control the execution of an operating system. The operating system that is being controlled is called the "guest operating system (OS)." What makes a virtual machine different from a simulator is that the guest OS is allowed to run most of its code directly on the hardware avoiding the severe performance penalties associated with simulation. The virtual-machine environment represents very precisely a real-machine environment. In the case of VMware, the virtual processor is an Intel Pentium or later, depending on the host processor.

GSX Server runs as an application on the "host operating system (OS)" that is running on the host machine. Various versions of Windows or Linux can serve as a host OS. The guest OS runs in a virtual machine. Many operating systems available on Intel chips can serve as guest operating systems. In fact, even DOS is supported in this mode.

ESX Server is designed for up to eight processor departmental and enterprise servers, and runs natively on the hardware with its own kernel. VMware GSX Server is designed for up to four processor workgroup and departmental servers. It runs as an application on top of Windows or Linux.

The major differences between ESX and GSX Server are:

- ESX Server is installed directly on hardware with a bundled console operating system.
- GSX Server runs on a wider range of hardware, on top of Windows or Linux.
- ESX Server offers extremely high performance for disk and Ethernet I/O.
- ESX Server has advanced workload management and resource governing.

VMware products use a patent-pending technology to allow the guest OS to run on the hardware. While simulation could allow execution of a guest OS, the slowdown might be in the 10 to 20 times range. By letting the guest OS run on the IA-32 hardware directly, VMware can offer much better performance than any simulator. VMware's software needs to take control so that it can execute certain I/O and supervisory instructions. VMware does not permit the guest to execute any instruction that could compromise the system's integrity.

The table on the next page summarizes some of the key features of the VMware products. More details are available on the company's website.

TABLE 1: Key Features of VMware

	VMware Workstation	GSX Server	ESX Server
Environment	Desktop	Server	Server
Server Hardware	• Standard PC with Pentium instruction set (PII, PIII, or 4, or AMD with minimum of 128 MB of memory).	• 400 MHz or faster, Pentium II , III, or 4, or AMD Athlon with 256 MB memory, maximum 4 GB.	• 500 MHz or faster Pentium II, III, or 4, or AMD Athlon or Duron with at least 256 MB (512 MB recommended), maximum 4 GB.
SMP Capability[1]	Yes	• Up to 4 processors, with typically 4 virtual machines per processor.	• Up to 8 processors with up to 20 virtual machines per processor.
Host OS	Windows or Linux	Windows or Linux	No host needed
Guest OS Supported	• MS-DOS; • Windows 3.1, 3.11, 95, 98, and XP; • NT and Windows 2000; • FreeBSD; or • Linux.	• MS-DOS 6.0; • Red Hat Linux; • Linux-Mandrake; • SuSe Linux; • TurboLinux; • Most versions of Windows including 95 and 98; and • NT 4.0. Windows 2000 Professional, Server, and Advanced Server.	• Windows 2000 Server and Advance Server; • NT 4.0; and • Red Hat Linux 6.x or 7.x.
Availability	Now	Now	Now
Evaluation Copy	Download	Request from company	Request from company
Price	$299	$2,499	Request from company

[1] VMware products are capable of running on SMP machines. Guest operating systems in virtual machines cannot utilize more than one processor.

SERVER CONSOLIDATION CONSIDERATIONS – SCALABILITY

Server consolidation faces unique challenges. Consolidation does not mean cramming as many applications into the largest computer that one can find or afford. This would create more problems than it would solve. The goal of consolidation is to create a group of systems located so that they can be managed and maintained more efficiently. In almost every company we surveyed, lower costs and increased security were almost as critical as increased availability. These benefits result from the professional management that can be brought to bear on the new system as well as more robust hardware and fewer points of failure.

However, one of the major barriers to consolidation in the Intel space has been the lack of scalability of the platform,[2] the operating systems, and the applications. After all, the roots of the Intel environment lay in the PC. With some exceptions, PCs are single-processor systems. The operating systems derived from this environment were not designed to take advantage of multiple processors; nor were the applications. For example, it is widely known that NT does not scale very well beyond four processors. Thus, consolidation of work on an NT system presents a special limitation, bounded by four processors.

Far greater scalability in the hardware is now available for consolidation projects. Unisys has announced a 32-processor SMP IA-32 machine. Four- and eight-way SMP machines are now common. IA-64 systems promise to match the scalability of UNIX. Windows 2000 and the newer versions of Linux are not far behind. DHBA expects that Microsoft will make any necessary adjustments[3] to Exchange, SQLServer, and other key parts of Windows 2000 to improve scalability. In addition, we expect that Oracle and DB2 will also make whatever changes are needed for improved scalability. Thus, the key hardware and operating-system barriers to scalability in the IA-32 environment are on their way to being solved.

A long-term problem may well be applications written by customers or small development shops that were developed for single-processor systems. They have never been tested for scalability. Sheer economics indicates that many of these applications will never be rewritten to take advantage of the more modern environments.

Therefore, DHBA believes that many companies will be maintaining applications that will not scale as far as future requirements call for. The greater throughput required will exceed the maximum platform they can take advantage of, making it advantageous to have multiple images of the application in production.

[2] In fairness, the microprocessor itself has been able to scale for a long time. Sequent (and NCR in the past) has been shipping scalable Intel systems for many years. However, the mainstream vendors shipping Intel products have been limited mostly to four-processor systems with some eight-processor systems available.

[3] In many cases, reasonably good scalability up to eight processors and beyond may already be possible. Matching the 64 or greater processors sported by the UNIX vendors will almost certainly require investments and development.

Companies will have to manage multiple copies of these applications in an efficient way. Hardware partitioning, which is available on enterprise servers, offers this capability. But software solutions will also have a place and here we believe that VMware with its virtual machine capability will play a major role.

Generally, software partitioning is ultimately more flexible than hardware partitioning – hardware partitions are fixed-size[4] so they are less able to take full advantage of available processing resources. VMware Server products allow users to turn their computers into a "pool of logical computing resources," which can then be dynamically allocated to any operating system or application at the user's choice of granularity.

[4] There are forms of hardware partitioning such as the LPARs found on the IBM mainframe that are more flexible than this. Such flexibility is coming to UNIX systems.

ADVANTAGES OF VMWARE FOR SERVER CONSOLIDATION

Based on conversations with IT professionals, we have identified seven key advantages of VMware in server-consolidation projects. Each of these will be discussed further. (Note that the treatment here of VMware features is not comprehensive.)

1. Isolation of the Virtual Machines
2. Encapsulation
3. Hardware Independence of the Virtual Machines
4. Compatibility with the Workstation Version of VMware
5. Remote Management
6. Disk Management
7. Virtual Networking and File Sharing

ISOLATION OF THE VIRTUAL MACHINES

This property is key in consolidation. Neither applications running in a virtual machine, nor the guest OS, can do anything that will bring down VMware, another virtual machine, or the host OS. Note that isolation does not mean fault tolerance. The host hardware or software can still fail and bring the system down. Isolation means that multiple virtual machines can exist on the same platform without any possibility of failure in one causing failures in others. The guest operating systems may still fail but the failures will not spread.

ENCAPSULATION

Encapsulation means that the state of a virtual machine can be saved to disk and then the virtual machine can be restarted in the time that it takes to reload the data from the disk. Restart times can be reduced to a matter of a few seconds. In the past, failures that have occurred in an application or the operating system may have required lengthy restart times. With proper planning, this function can significantly improve the service that the end user sees. DHBA has discussed consolidations with customers where a significant financial impact of downtime has been quantified. In many of these cases this factor alone contributed to the justification of the project.[5]

[5] Beware, you will need the buy in of the financial people in the company to get such a happy result. They need to agree to a cost for downtime to the business units. This is not easy to get.

HARDWARE INDEPENDENCE OF THE VIRTUAL MACHINES

This feature should make life easier for testing and development. Virtual machines can be constructed on any VMware platform and then executed on any platform where VMware is installed. In other words, it is relatively easy to create a virtual machine and then move it to a test machine. When testing is complete, it can be moved into production without any requirement to rebuild it for the production machine. This has many utilities. In consolidation projects, for example, finishing the project on schedule is vital. It speeds the process – reducing the number of more expensive machine hours required – and saves money.

COMPATIBILITY WITH THE WORKSTATION VERSION OF VMWARE

All versions of VMware use the same core technology. Learning the workstation version helps to understand the server versions of the product. This eases development and testing and relies on current in-house skills gained through the more widely available workstation. It may also mean a reduction in the amount of time that expensive hardware is tied up for development and testing.

REMOTE MANAGEMENT

VMware bundles useful software that allows remote management of the virtual machines. For example, desktop software allows an administrator to have a remote console for each virtual machine. The remote console is always available and does not require that the guest OS be running. If the OS crashes, the administrator can debug the situation from the remote location. Administrators can also use a web browser, Internet Explorer or Netscape, to manage the virtual machines. This management facility is especially important in a consolidation environment since it can allow business units to manage their virtual machines via the Internet.

DISK MANAGEMENT

VMware provides options on how virtual or real disks are handled. The first option is to have disks handled conventionally. Any changes are permanent. VMware calls this a “persistent” disk. In the second option, called a “nonpersistent” disk, all changes that are made are thrown away at the end of each session. If one has ever struggled with the disk setup for a test session, for example, they will realize what an advantage a nonpersistent disk can be. One can set up disk files once, and at the end of each test, have them reset automatically. This can save a huge amount of setup time and greatly facilitate testing. Testing, of course, is critical in consolidation projects. There is a third option as well, called an “undoable disk.” After the test, one can choose whether to keep or undo the changes. This will be valuable when examining the changes before resetting for another test.

VIRTUAL NETWORKING AND FILE SHARING

VMware can connect different virtual machines running in the same physical system with TCP/IP or other network protocols even though no physical connections are used. Virtual machines can also connect to separate virtual machines running in another system or to existing networks. This feature is useful in testing and may find a place in production. IT managers report that handling the network is one of their biggest problems in a consolidation project. Stress testing the network proves particularly troublesome. With some planning, these VMware features can help considerably. (Other protocols such as Netware's IPX or NFS are also supported.)

HOW VMWARE ADDRESSES TECHNICAL BARRIERS TO SERVER CONSOLIDATION

Consolidating applications onto a single hardware platform means that they must run on the same version of the operating system.[6] Suppose that applications A and B are consolidated on a single system. All may go well until a new version of application A is installed, which either needs a new level of the operating system or perhaps a service pack for the operating system. What happens if application B will not work with the new operating system? Suppose that the users of application A need immediate implementation of the new release because customers are demanding it. Cases of this sort can force a de-consolidation. In other words, applications A and B will be split apart to run on separate machines.

VMware products can address this situation. Each application can run in its own virtual machine. In addition, with its ESX Server, VMware offers resource-management facilities to allocate CPU, memory, and I/O resources among the virtual machines.

Another potential problem with consolidation concerns applications that may not run together very "happily." Such problems are hard to diagnose and even harder to correct. If the problem takes the new system down, the consolidation of these applications may have to be abandoned. VMware – with the isolation that it provides for the virtual machines – solves such problems without the need to diagnose or correct them.

[6] On mainframes or larger UNIX boxes like Sun's E10000, a hardware-partitioning scheme may allow different versions of the operating system on a single machine. Partitioning is also available in the IA-32 world with the Unisys ES7000. In the future, IA-64 systems will probably match the facilities in mainframes and UNIX.

LIMITATIONS OF VMWARE

There are many advantages of VMware. There are also limitations:

1. Support only for IA-32 Architecture Systems
2. Overhead of Learning a New Application
3. No Support for Windows 2000 Datacenter Server
4 VM Performance
5 Small Company Growth Pains

SUPPORT ONLY FOR IA-32 ARCHITECTURE SYSTEMS

VMware products currently only support IA-32 hardware. One can imagine IA-64 support in the future, although the changes required for the EPIC architecture may not come easily for VMware. In addition, competitors may enter the market, since the patent protecting VMware's technology probably applies only to the IA-32 architecture. Other architectures like SPARC may seem tempting, but in DHBA's view, the company has been wise to resist defocusing its efforts. IA-32 and IA-64 should keep the company occupied for the immediate future.

OVERHEAD OF LEARNING A NEW APPLICATION

Server consolidation is complicated on its own, requiring technical and managerial skills, effective public relations, and project management. Adding the need to understand another application [ie VMware] increases the cost, length, and complexity of the consolidation project. However, VMware seems to have a short learning curve. Thousands of workstation users have mastered the product with relative ease. The server versions are more complicated with scripting capabilities and resource management controls. In addition, the ESX Server will require more knowledge, since it does not depend on a host OS. On balance, learning VMware should not be a major burden on a project. Of course, it would be wise to learn it[7] *before* the start of the consolidation project. (The workstation version of the product should offer an easy way to gain familiarity with the concept.)

NO SUPPORT FOR WINDOWS 2000 DATACENTER SERVER

Currently, no version of VMware supports Windows 2000 Datacenter Server as a guest OS – a limitation from a server-consolidation view. Windows 2000 Datacenter Server is likely to be the target for many consolidations and VMware would be quite useful in these cases. One can only hope that the ESX Server may add this capability in the future.

VMware GSX Server is a complementary tool to Windows 2000 Datacenter Server. When GSX Server is installed on a system running Datacenter, it enables

[7] Please see: Jason Compton, *VMware2 for Linux* , Prima Publishing, August 2000.

IT organizations to logically partition the workload of the enterprise-class server, permitting many applications to share the system in secure, isolated virtual machine environments.

IT organizations that wish to consolidate applications with incompatible operating systems or service pack requirements will find VMware ESX Server to be a powerful complement to Datacenter in their enterprises. Datacenter hosts the environments that coexist well on the same platform and ESX Server supports the more problematic combinations.

While Datacenter is not currently supported as a guest OS on VMware Server products, this is not due to any technical incompatibility, but rather, to two specific considerations. The first is Microsoft's providing Datacenter only in conjunction with specific hardware systems; it is not provided as installable software such that it could be loaded into a VMware virtual machine. The second is the single-processor implementation of virtual machines in VMware's current products, which masks the added benefits of Datacenter over its stablemates, Windows 2000 Server and Advanced Server. When multi-processor virtual machines are available, assuming the distribution issue can be addressed, VMware virtual machines should provide an excellent platform for Datacenter as well as other versions of Windows 2000, Windows XP, and the forthcoming .NET Server products.

VM PERFORMANCE

Historically, the weakness of the VM concept has been performance. Although VMware can run the guest OS in a relatively efficient way, performance can still be a problem. Generally, it is fair to say that performance has not been a significant problem with VMware users. However, applications will use more memory and CPU in a virtual machine than they would in a real machine. It is difficult to come up with general rules because there are too many variables.[8] However, some rules of thumb might help. DHBA also suggests that as they gain more experience with a wide array of consolidations VMware may wish to develop formal performance guidelines.

- Begin by analyzing the resource usage of the workloads to be consolidated. (It is amazing how many server-consolidation projects get into trouble by ignoring this step.)

[8] In a March 2001 article, "Consolidate Servers with VMware GSX, ESX," by Moshe Bar, the author measures some test cases. He concludes that compiling the Linux 2.4.2 kernel in a world on his system took about twice as long as in a real machine. The real machine was a dual CPU and he thought that the world, which was constrained to a single machine, gave reasonable performance. When he started six worlds and ran the same compilation simultaneously in each they finished in about 14.5 minutes. The author seemed quite satisfied with the performance that he got from VMware. The entire article is well worth reading. You can do so at http://www.byte.com/column/BYT20010327S0009.

- Determine the resource needs of each workload when it is moved into a world. The elongation factor[9] can now be calculated between the real machine results and the virtual machine in the world.
- Be conservative. Consolidate to machines with higher megahertz speeds. For example, if the real machine is a 350 MHz processor, step it up to a 500 MHz machine, which is the minimum that VMware recommends for the ESX Server.
- Be generous with memory. Swapping to disk will hurt performance.
- Take care with the I/O requirements. You need to have enough disk I/O capability to avoid conflict among the worlds over the disks.
- Make sure there is a provision for multiple virtual machines hitting their peak requirements together. Ignore this rule only if absolutely certain that the worlds never peak together. For example, one world may peak at 10 a.m. and the other at 5 p.m.

SMALL COMPANY GROWTH PAINS

VMware is a new, rapidly growing company. Hiring and training a technical staff poses a challenge. The company is making efforts to spread knowledge of its technology to other players inside and outside the firm, but this will take time. Discussing support requirements with the company and getting specific commitments from VMware is advisable.[10] In other cases, plan the project so that responses from VMware are not on the critical path. As VMware moves into the datacenter, it is expected that it will provide the support that its customers will need.

[9] The elongation factor is how much longer it takes to run a given workload in a world than in a real machine. For example, if it runs in 10 minutes in a real machine and 11 minutes in a world the elongation factor is 1.1.

[10] 24x7 support is available from VMware for a fee.

PLANNING FOR SERVER CONSOLIDATION WITH VMWARE

These notes are not intended to discuss the totality of creating a server-consolidation plan.

DHBA is developing a complete Server-Consolidation Series that addresses broader issues. However, there are some unique considerations that apply to VMware that are addressed here:

1. Plan to build up experience with VMware. If at all possible do not wait for the start of the project. Keep the focus on applications that lend themselves to VMware use.
2. Develop a comprehensive strategy for the use of VMware. Include documentation of the worlds needed.
3. Address the performance needs of the worlds that will be created. Allow time for the tuning of the I/O of these worlds. Identify applications that are not suitable to run in a world.
4. Test all the applications with VMware.
5. Stress test all applications that will run with VMware. Research partners who have discussed consolidation with DHBA report universally that they wish that they had done more stress testing. This testing should also help assess what the host hardware can deliver.
6. Be sure to understand what VMware can and cannot do. For example, while it does provide a hardware level of communication between worlds, it does not address the issue of data interchange between different virtual machines. If such facilities are required, include the resources to develop them, or acquire them from the various vendors of the applications in question.