VMware On the March

Research Note

Thomas Deane
Gordon Haff
Jonathan Enuice

There are few organizations these days that are not hugely interested in optimizing and simplifying their sprawling IT infrastructures. Vendors have responded by rolling out hundreds of products promising automation, virtualization, and utility-like computing. Some are part of all-encompassing strategic plans, from Adaptive Enterprise (HP) to Agile IT (EDS) to N1 (Sun) to On Demand (IBM) to TRIOLE (Fujitsu), to name a few. Others are more tactical and focused on a single type of problem or task. Some deliver on their promises today. Others are better described as "a work in progress."

Amidst this helter-skelter background, a few companies stand out. When it comes to virtualizing the nearly omnipresent x86 platform, VMware[1] has distinguished itself over the past several years. Now owned by EMC, it started as a small provider of tools catering primarily to software developers and testers, but has marched steadily towards and into production datacenters. The past three quarters have seen a blitzkrieg. Not only has VMware refreshed its product line at all levels, it has significantly expanded its flagship ESX Server to include multi-system management (VirtualCenter) and the unique ability to migrate live applications from one system to another (VMotion). With these, VMware has become an entirely strategic enterprise provider.

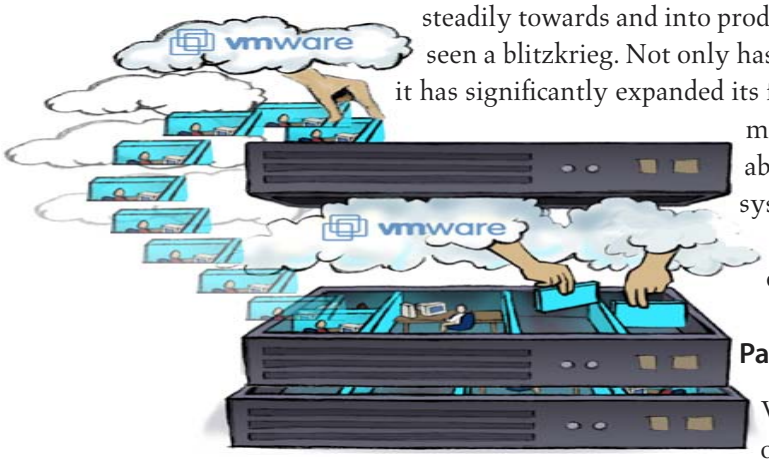

Papa Bear, Momma Bear, Baby Bear

VMware's vPlatform suite is the historical core of a product set that now also includes a vManage group (management software such as VirtualCenter and VMotion) and a vTools group (tools such as its Physical-to-Virtual (P2V) Assistant conversion tool).

The three vPlatform products are ESX Server, GSX Server, and Workstation—call them papa, momma, and baby bear. Each does essentially the same thing: create virtual machines (VMs) on x86 architecture systems. That is, they carve a single physical system into compartments, each running its own copy of an operating system and applications and content in the illusion that it is a complete, independent machine. Partitioning a system this way fools software into thinking that the virtual machine abstraction is a real, physical system. But, in reality, the VM has

1. VMware only runs on x86 architecture hardware.

access to only a portion of the system's physical resources. Thus, multiple VMs—each with its own independent operating system (OS) and applications—can co-exist on a single system.[2] What's more, these VMs can be created, resized, even destroyed on the fly.

VMs isolate application and OS faults. For example, if one VM gets a virus and "blue screens" the OS, none of the other VMs on the same physical system would be affected. Also, since physical resources are allocated on a per VM basis, no single VM can hog the entire system because of a runaway or poorly performing application.

Virtual machine technology is often considered a form of partitioning. It goes, however, considerably deeper. All versions of VMware, for instance, can create "snapshots" of the "state" of an application; the application and OS, configuration files, and local data can be saved as a disk file. VMs can be stopped, then later restarted exactly where they left off. Indeed, VMs can be moved to another system before being restarted—even one with quite a different configuration, as long as it's an x86 system that supports VMware. This motion can also be temporal: VMs can be saved to disk or backup media, and restarted at some later date. In short, VMware's VMs provide substantial opportunities beyond partitioning for server consolidation.

While VMware's three vPlatform products all work in basically the same way, they differ considerably in implementation, capabilities, target markets, and how they are installed and managed.Running Atop: Workstation and GSX Server

The Workstation product continues VMware's original appeal to developers who need to set up a variety of different OS environments—and who often need to roll back a system to a known good state. Enterprises have also used VMware Worksta-

2. See Illuminata report "VMware: Virtual Partitions for the Server Masses" (October 2002) for more background on virtual machines and how VMware's software works.

Key VMware Product Features

Product	Workstation	GSX Server	ESX Server
Current version	4.5	3.1	2.1.1
Host OS	Windows, Linux	Windows, Linux	None (VMkernel built in)
Guest OS*	Windows, Linux, NetWare, FreeBSD, Solaris x86	Windows, Linux, NetWare, FreeBSD, Solaris x86	Windows, Linux, NetWare
Max # of CPUs per System	1	32	16
Max # of CPU's per VM**	1	1	2**
Max Memory per System	4GB	64GB	64GB
Max Memory per VM	3.6GB	3.6GB	3.6GB
64 Bit Support	Experimental support for 64-bit hosts today, full support for hosts and guest coming	Experimental support for 64-bit hosts today, full support for hosts and guest coming	Operates on EM64T and AMD64-based hardware in legacy mode. Full support coming
VMotion Support	No	No	Yes
VirtualCenter Support	No	Yes	Yes

* VMware qualifies specific versions of specific OSs. In many cases, other versions and other OSs will run; they just aren't supported by VMware if there are compatibility or other issues.

**Using VMware's vSMP, VMware's Virtual SMP add-on, which allows a single VM to span multiple CPUs.

tion as a way to run legacy applications and OSs alongside newer environments on the same desktop.

GSX Server, like Workstation, runs on top of an OS—either Windows or Linux—running on the host physical system. While GSX Server, unlike Workstation, is appropriate for production servers, it similarly leverages the device drivers, installation mechanisms, and support procedures of the off-the-shelf operating system on which it runs. That tends to make it a good match for consolidating workloads at the departmental or remote location level, where datacenter-level support skills are lacking.

Beyond server consolidation, GSX Server also provides unique application upgrade possibilities. Legacy applications can run on a Windows NT virtual machine, while other applications run within Windows 2000 or Windows Server 2003 VMs; yet others can run Red Hat Linux VMs.

GSX Server is also particularly useful for test and development environments. It lets a single server run multiple versions of Windows, Linux, Novell Netware, FreeBSD, or Solaris x86 simultaneously. Multiple versions of an application—whether that's development, test, and production or multiple development and test instances—can run simultaneously on the same gear, without interfering with each other. The ability to mix and match many OS and app permutations on a single system streamlines development and testing operations. Replacing numerous development/test machines with a virtual infrastructure also has a substantial return on investment, by avoiding the need to buy and administer a large number of dedicated physical hosts.[3]

This strategy works with many, but not all, testing situations. It is particularly good for functional tests. But if, for example, you are measuring network collisions and bottlenecks, emulating multiple systems using VMs sharing a single Ethernet network interface card (NIC) is not going to work.[4] Also, virtualizing and abstracting away differences in underlying hardware is not a benefit if dependencies and interactions with different hardware platforms are the object of the testing. The virtual environment has to truly emulate the real world—or at least get close enough for the purposes of the task at hand.

The newly released GSX Server 3 offers improved management features making VM deployment and configuration quite easy to perform. This version also improves performance, including a 10 to 20 percent bump in virtual disk and networking performance, and increased virtual memory size for each virtual machine.[5]

Industrial Strength: ESX Server

ESX Server is the papa bear of the VMware line—the industrial-strength foundation for partitioning, consolidating, and managing systems in mission-critical environments. In combination with VirtualCenter, VMware's multi-system command and control console, ESX Server provides the advanced workload management and provisioning to manage VMs in a large distributed data center. The VMotion option steps up capabilities even further, letting VMs migrate from one physical system to another with essentially zero downtime—more on this later. While GSX Server and ESX Server environments can now be managed by VirtualCenter, VMotion remains an ESX Server exclusive.

Rather than layering on top of a pre-existing, pre-installed OS, ESX Server is self-contained. ESX Server is built on top of a microkernel-based OS (the VMkernel, which looks like a specialized Unix/Linux variant) that includes its own memory management, resource scheduling, and device drivers. ESX Server uses that level of control to support up to two processors[6] within a single VM, and to reduce the performance overhead of virtualization.

3. Also of note, VMware provides "non-persistent" and "undoable" virtual disk drives: useful abstractions that substantially ease and speed many configuration and testing tasks.

4. However, VMware does support up to 4 NICs per VM

5. 3.6GB can now be allocated to each VM in the latest versions of all vPlatform products.

6. Using the optional, add-on VMware Virtual SMP.

The latest 2.1 and 2.1.1 versions support AMD's popular Opteron processor, as well as Hyper-Threading[7] on Intel processors. They also increase support for storage area networks (SANs). However, the greatest advances are easier installation and management. There's now a graphical installer and redesigned network connections management, including support for the Virtual Local Area Networks (VLANs) that operate within an ESX server.[8]

Beyond Slice and Dice

VMware began as a way of portioning x86 servers. That continues to be a valuable capability. But as both an organization and a product suite, VMware is increasingly about much more than cutting servers up into little virtual pieces. VMware is increasingly focused on creating systematically virtualized infrastructures that span entire datacenters, and beyond.

vPlatform products provide considerable abstraction and encapsulation in the course of making x86 systems partitionable. But that abstraction and encapsulation has value well beyond partitioning. In particular, VMs are largely hardware-independent. Thanks to the "virtualization layer," VMs can be built once and then quickly deployed to essentially any system running VMware—even if the underlying hardware platform is different from the hardware platform on which the VMs were created. VMs can also be easily moved, copied, backed up, and cloned. That hardware independence also lets new systems be easily provisioned. Deploying a new VM can be as quick as copying a set of VM container files to a new system and starting it up.

This hardware independence is not, however, free in terms of either performance or cost. VMware imposes a modest but real CPU overhead.[9] However, whether you peg it at 2% or 20% is not particularly relevant. Insufficient CPU horsepower is not really a pressing issue in most of today's datacenters. Thanks to Moore's Law, most servers have plenty to spare. The issue is pairing that plentiful capacity with the workloads that need it in an efficient and dynamic way. The vPlatform suite tackles this problem on a server-by-server basis. However, VMware is now taking advantage of the virtualized servers to tackle the management of the entire virtualized infrastructure using vManage.

vManage and vTools Suites

vManage is at the core of where VMware is headed, and why EMC acquired the company.[10] To be sure, many of the details about how EMC will leverage VMware's assets remain fuzzy—with more vision than nuts-and-bolts specifics. This is partly the inevitable consequence of needing time to merge EMC's storage-centric sensibilities with VMware's server-centric ones. It's also a result of letting the VMware subsidiary maintain an arms-length relationship with the parent in some respects—the better to allay the concerns of key partners like HP and IBM that VMware might start to favor EMC storage over their own. But, that caveat aside, it's clear that compute virtualization dovetails with EMC's other storage virtualization assets as part of its ambitious Information Lifecycle Management (ILM) mission.

VMware's VirtualCenter provides a central point of control for remotely monitoring, creating, deleting,

7. HyperThreading is Intel's implementation of Simultaneous Multithreading, a technique to increase processor utilization by making one physical processor appear as multiple logical processors to the OS and applications. See Illuminata report "Breaking Up The Microprocessor Monolith" (July 2003).
8. VMware VLANs aren't quite the same thing as Ethernet VLANs even if the name implies it. Instead, they let VMs communicate securely among themselves as if they were connected to a common isolated physical network. VMs can be configured to be on different VLAN segments of the datacenter.
9. Numerical estimates of VMware overhead range widely, and they are very dependent on the specific application and vPlatform product. VMware claims overheads in the 2-15% range. A joint VMware/IBM evaluation found a 12-15% overhead for ESX Server on the Trade3 benchmark, which models an online stock brokerage application. An earlier review of Workstation found 4%, 13%, and 16% overheads for various tests. Other anecdotal estimates indicate a 20% or so overhead, especially under high I/O loads.
10. See Illuminata report "EMC Takes its Virtualization Well Beyond Storage" (December 2003).

and moving VMs and their associated applications on ESX Server systems across the physical infrastructure of a data center. For example, VirtualCenter can deploy new servers using a template-based wizard—taking advantage of the fact that virtualized hardware abstracts out many of the specifics and dependencies that often cause problems with cloning systems.[11]

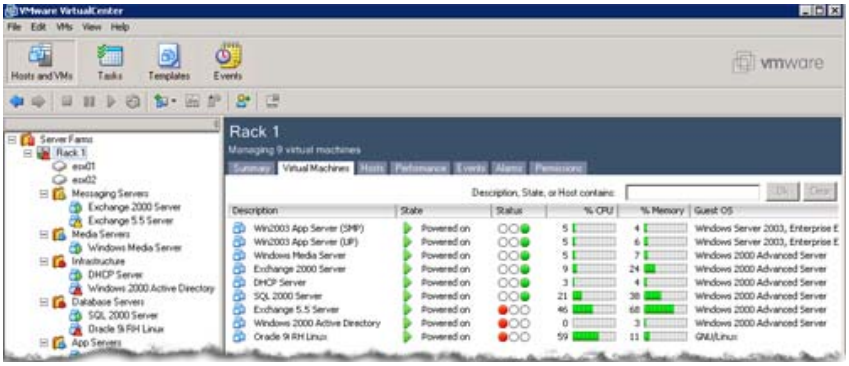

VMware's VMs can also be deployed using a number of third-party deployment tools. These imaging and scripting tools include, for example, HP's Rapid Deployment Pack (RDP), Symantec Ghost, Microsoft ADS, and Veritas OpForce (Jareva). After all, a VM looks like any other type of server to this type of software. For more intimate interaction with the VMs, VMware provides a Virtual Infrastructure SDK to help third-party products integrate more tightly with the VMware virtual infrastructure.

VirtualCenter treats a server's VMs running within an ESX or GSX Server, as well as the physical servers themselves, as a single pool of resources. It unifies the management of these diverse environments, which can include a variety of different OSs, onto a GUI dashboard on a single pane of glass.[12] This can be a significant benefit in today's typical complex data center environment in which a manager has to deal with multiple screens of information. VirtualCenter also monitors system availability and performance and provides automated email alerts.

VMotion augments VirtualCenter to dynamically add or move VMs as needed to compensate for changing workload conditions, take care of scheduled hardware maintenance, and even handle hardware failures. This ESX Server-specific tool allows running VMs to be moved "on the fly" from an ESX Server running on one physical machine to another with essentially zero downtime and *without impacting users*. The key operating phrase here is "on the fly." Not quiescent or stopped users, processes, applications, and OS's—but live and kicking ones. One can use VMotion to do hardware or software upgrades without users' ever noticing it. Way cool!

In order for VMotion to work, the ESX Servers have to be connected to a common high-speed network (LAN) and storage area network (SAN), and they must use the VMware file system (VMFS). The context of a VM is contained in VMFS container files. In order for any VM to be moved, these VMFS files need to be visible to both source and target ESX Servers. The key here is SANs, SANs, and more SANs (no NAS at present, and certainly no direct-attached standalone disk). And this certainly is a good fit with the EMC storage centralization religion. Is EMC's plan becoming clearer yet? Given a SAN infrastructure in place, applications can be transported with barely a sub-second "hiccup."

VMware rounds out its product suite with vTools. These are handy utilities for the virtualized datacenter, or those headed in that direction. For example, VMs can be statically moved from Workstation to GSX Server to ESX Server—and between different versions of these products. Some of these moves require a straightforward but mandatory conversion; VMware provides tools where needed for this exact purpose. There are also Physical-to-Virtual (P2V) Assistant tools to convert physical OS environments to virtual environments—but not

11. See Illuminata report "What is Automated Provisioning?" (May 2004).
12. The VirtualCenter console runs on a Windows PC. While potentially a negative in a Linux/UNIX-centric environment, it is an unusual datacenter that doesn't already use Windows clients for other purposes.

the other way.[13] However, VMware provides tech docs to help customers with that process.

Good, But Not Cheap

VMware revenue has been growing rapidly. In EMC's second quarter of 2004, the VMware subsidiary contributed revenue of $47.2 million, up from $39.3 million in the first quarter, and a 200 percent rise from the same quarter of 2003. They don't post those sorts of results by giving their product away.

The chart shows the initial cost of GSX Server and ESX Server, including a mandatory yearly support contract.[14] For a typical 2U, 2 CPU rack-mountable server that might cost $7K-8K,[15] the additional $4.5K for ESX Server (with its yearly service contract) adds about 60 percent to the initial purchase price. In general, the cost of VMware is often a significant portion of the base server cost. So VMware virtualization doesn't come particularly cheap.

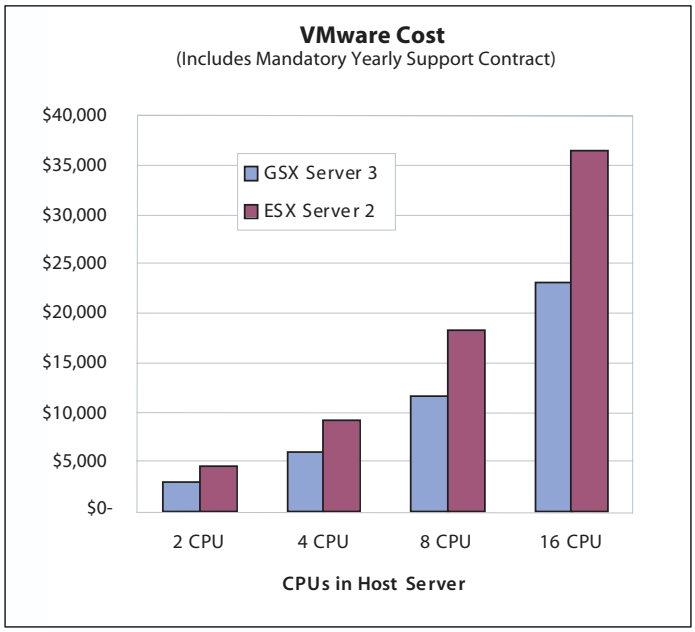

The good news on the cost front, however, is the ROI. VMware allows operators to avoid the purchase of two, four, or more additional servers—as well as their associated power, cooling, floor space, networking, and administrative costs. So if you've got the headroom and can consolidate a few servers, it's a winning proposition. One user is, for example, consolidating numerous Windows Domain Controllers and Active Directories—running across a range of Windows NT, Windows 2000, and Windows 2003.[16] In Boston, a government agency needs 6-7 new Windows servers to test just *a single instance* of a new application they are developing. Situations like these show why, even if VMware isn't cheap in an absolute sense, it would often be a good *value* at twice the price. And this is before one considers the additional things one can do with a VMware-based virtual infrastructure that were simply impossible before—such as the ability to move workloads around in real-time.

Not All Quiet on the Virtualization Front

VMware is making good progress on its vision of Virtual Infrastructure. It is not, however, alone. Veritas, for example, is clearly charging into provisioning and virtualization with its purchases of Jareva and Ejacent.[17] IBM's POWER5 servers are bringing out fine-grained "micro-partitions," as well as higher level software tools under its Virtualization Engine banner.[18] Sun—and not incidentally, also its beta users—tout the N1 Grid Containers

13. While it's conceivable that an admin might want a virtual-to-physical tool to eliminate the VM overhead on a highly utilized server, many datacenters, in practice, are preferring to virtualize all their servers for administrative ease—even if there's just one VM per physical server.
14. Based on a "Gold" support contact that includes 9-to-5, Monday-to-Friday service. Datacenters running mission-critical workloads would often opt for pricier 24/7 support.
15. See Illuminata report "Servers: Who's Cheap?" (May 2004).
16. As one user in the VMware forum stated: "We have over 25 domain controllers (Windows 2000 and Windows NT 4.0) running on VMware. It was the best thing we could have done. We usually allocate 384MB of memory and 10GB of disk to Windows 2000 DC's, and 192MB of memory and 4GB of disk to Windows NT 4.0 domain controllers. We are supporting about 30,000 domain based users, and all the DC's are running at about 5-10% utilization. We've found zero downsides..."
17. Many believe that at least one reason EMC snapped up VMware was to keep it out of Veritas' hands, and thus deny Veritas a particularly strategic asset and opportunity.
18. See Illuminata report "POWER By The Piece" (July 2004).

coming in Solaris 10 as a more flexible and lower overhead mechanism than VMs.[19] HP has talked about doing its own virtual machine technology for Itanium. Third-party software providers such as Aurema, Ensim, SWsoft, and Trigence provide related OS-extending resource management and application container technologies. Intel is optimizing virtualization performance in projects such as "Vanderpool" and "Silvervale" that, in addition to accelerating virtual machines, may lower barriers to market entry. Finally, add open source projects such as User Mode Linux (UML) and coLinux to the mix. The virtualization front is red hot.

Few if any of these alternatives truly compete directly with VMware. Virtualizing x86 CPUs isn't easy, and VMware has special (patented) sauce for doing so. So while the alternatives don't do precisely what VMware does, and certainly don't do all of what VMware does, they do compete for the same pool of customer attention, commitment, and dollars. And they enjoy some advantages of their own. Some are more efficient with system resources. Some are more application-, rather than platform-, centered. Some just cost less.

Perhaps the most important competitor, at least going forward, is Microsoft. The Virtual Server 2005 product that will (eventually) emerge from Microsoft's acquisition of Virtual PC technology from Connectix[20] could begin to commoditize VM technology. And Microsoft's brand and marketing presence is hard to over-rate in the Windows arena—especially with the opportunity to integrate it into future platforms such as "Longhorn."

Yet despite Microsoft and the others, VMware clearly retains its lead. It has production products for x86, today—and they are field- and marketplace-proven. It's moved forward with higher-level capabilities such as VirtualCenter and VMotion, and seems likely to do even more within EMC's ILM and disaster-tolerance strategies in the quarters and years to come. Moreover, VMware is one of very few who have the deep experience needed to create, harden, and optimize complex, sophisticated virtualization products. There's an enormous difference between having things that run in the labs or that can be cobbled together to play around with on a PC and things that are genuinely production-ready. In this regard, VMware could still have a year or even multi-year lead on many competitors.

Beyond base product, VMware has a substantial lead in breadth and diversity. VMware today runs Windows (several versions), Linux (many versions), and other OS's (like NetWare). Other solutions tend to be Windows-only, Linux-only, or Unix-only. Microsoft Virtual Server, for instance, could in theory run something other than Windows. But while Linux or NetWare atop Microsoft Virtual Server is in theory doable, it won't be supported. Even if it were, would you trust Microsoft to support it in a production context? We know we would. *NOT!*

Whoo-Whoo! Where Goes the Virtual Train?

Thanks to a combination of systems vendors, Microsoft, and open source, in a few years charging for basic partitioning will make as much sense as charging for a scheduler or a file system would today. Possible—but it would be at most a specialized niche market.

So with VMware's boilers blazing and its full head of steam, where will it head? In short, up-market and up-level. In the near term, virtualization is a high-value proposition, and has yet to be successfully commoditized by anyone. That means there's a fair amount of track remaining on the current business. And as that starts to be less attractive, there are other destinations—ones that require getting to VMware's current level of technology first. For example, consider disaster preparedness and change management.

Facilitating Disaster Recovery & Tolerance (DR/DT) - Disaster Recovery (DR), at a minimum, constitutes taking data backups regularly and storing backup media safely at a different site. It's what you do when you're worried about a system crash, an

19. Never mind that not long ago Sun was promoting physical partitions as the *only* safe way to carve up machines.

20. Virtual Server has had, um, quite a long gestation period. It was supposedly almost ready to ship when Microsoft acquired Connectix in early *2003*. See Illuminata report "Microsoft Assimilates Virtual Machines" (February 2003).

application error, or a virus corrupting data. Full-blown Disaster Tolerance (DT) involves setting up a working replica of your computing environment at a secondary site. It's for those thinking about a 9/11 event or 9.5 on the Richter scale.

VMware can help in both of these areas. VM snapshots are point-in-time copies of a system's state—a known good copy of user, application, and OS state—saved to disk. They allow a datacenter manager to revert back to a particular system state in case something goes wrong. External data and transaction sources might not be included, but synchronization with those can often also be arranged. At a minimum, VMware considerably simplifies the DR/DT problem. Moreover, it does so at a much more effective cost than can be done today, because it avoids the "idle standby" nature and reduces the 1:1 ratio of production to DR/DT equipment that makes current solutions so darn expensive.

VMware's corporate parent, EMC, makes this direction clear and substantial. EMC's product set (SRDF, Time Finder, etc.) combines nicely with those of its divisions (Legato and Documentum) and with VMware's ability to snapshot VMs. Taken together, further developed and integrated, you have the foundation for particularly powerful recovery plans featuring short downtimes and modest costs.

Managing Change – VMware has promoted itself primarily on the basis of server virtualization and consolidation. But IT costs aren't just about utilization. A great deal of today's IT spending goes toward grappling with change. Migrating from outdated to newer platforms; coping with seemingly gratuitous updates to operating systems and middleware; and coordinating the inconsistent and unsynchronized updates of numerous IT pieces. Change is a monster! VMs, however, help reduce the pain of system updates, by allowing older software to run happily alongside new, and to run on newer gear without going through disruptive upgrades.

VMware could even begin sell a *Lite* vPlatform option that could run just one VM at a time. This would be a "Change Management Facilitator" rather than a "Consolidator." Such an offering would provide insulation change, and begin to establish the VMware environment (rather than specific instances of Windows or Linux or Xeon whatever) as the working platform for datacenter deployments.

Conclusion

VMware remains the unique provider of system virtualization and consolidation for the ubiquitous x86 platform. It evolved and matured into an enterprise provider and platform at just the point in time when customers were avidly exploring ways to consolidate workloads and figure out how to virtualize and simplify their IT infrastructures. That mission will continue to be a high-value one for the next few years.

But eventually the combination of Microsoft, open source, and systems vendors' offering their own approaches will commoditize basic virtualization and consolidation capabilities. To that end, VMware is currently transitioning to be a provider of beyond-the-VM products like VirtualCenter and VMotion that operate at a higher pan—the datacenter level—in contrast to its traditional products that are more narrowly concerned with slicing and dicing individual boxes. It is also moving toward even more systematic outcomes such as facilitating IT auditability, disaster recovery, disaster tolerance, and change management. While perhaps more abstract and more difficult to simply explain than simple slice-and-dice, such high-level functions are no less valuable to large enterprises. Indeed, over time, they are more valuable.

This up-market evolution is critical to VMware's continuing success. It was these up-leveled products and opportunities for which EMC bought VMware. Whether for basic or more rarefied functions, enterprises are buying into the VMware platform eagerly—and for good reason. VMware is a company on the march.