Virtualization as an Enterprise Computing Infrastructure

As the IT industry continues to refine the notion of using commercial off-the-shelf parts to create enterprise-class data processing systems, the problem of transcending the tight coupling of software stacks to hardware platforms continues to be an obstacle. Hardware virtualization tools, such as VMware® ESX Server™ software and Microsoft® Virtual Server software, represent an evolving genre of products that can dissociate software from hardware, providing several impressive systems management benefits.

BY J. CRAIG LOWERY, PH.D.

In recent years, open standards–based hardware and software products have proliferated. Many enterprise IT administrators have embraced this trend by migrating to standards-based systems in their data centers. However, issues of manageability still arise. More and more, IT administrators are turning to hardware virtualization tools, such as VMware® ESX Server™ software and Microsoft® Virtual Server software, to help manage these standards-based platforms.

Understanding the problem with scaling out

Sheer computing power is inexpensive. In an enterprise data center where demand for computing resources continues to rise, a natural tendency has been to acquire more systems to meet the needs of users. At first, this proliferation was thought to be an economical solution, but it soon became apparent that scalability was not simply a matter of finding rack space and power for new servers.

On the surface, using several inexpensive, standards-based components—such as one- and two-processor Intel® architecture–based servers with 1U form factors—can be an appealing option for the typical enterprise data center. As the need for more processing power arises, administrators increase capacity by adding more of these small components to the aggregate system. Because the components are standards-based, they are largely interchangeable and even disposable; if a system malfunctions, it can be easily replaced. This method, commonly referred to as *scaling out,* has merit but is not without caveats.

Manageability has become the chief obstacle for the commercial off-the-shelf (COTS) approach to enterprise computing. Certainly, many standard systems can be colocated in a data center, loaded with software, and placed into service. However, little time passes before administrators must apply operating system (OS) patches, repurpose a server, perform hardware maintenance, or adjust hardware allocation because of changes in demand. These challenges may foil the scale-out proposition for early adopters, even though great improvements are being made in software change-management tools. In fact, management issues in a scale-out environment can be so problematic that many IT departments have returned to consolidation of hardware resources. Instead of scaling out, they choose to *scale up*—that is, replace their many small, inexpensive systems with a few large, proprietary, and very costly systems in an effort to make the data center easier to manage.

Unfortunately, those who have consolidated into a scale-up configuration are treating a symptom, not the root

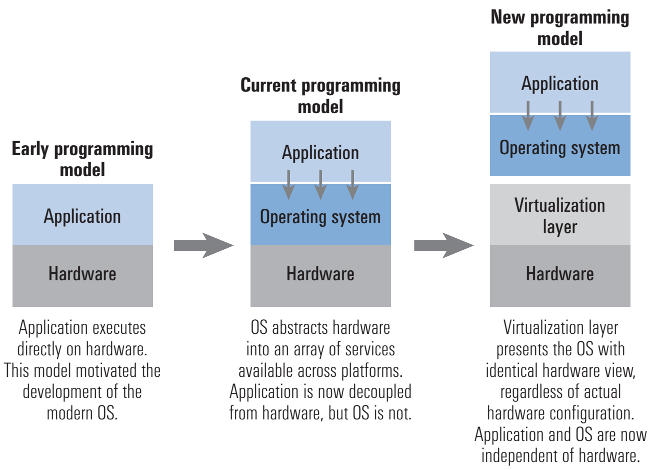

Figure 1. Evolving programming models

cause, of management issues in a COTS-based data center. The root cause is not the proliferation of the servers themselves, but the tight coupling of software to hardware. Each management issue mentioned previously—software patching, hardware repurposing, maintenance downtime, and workload balancing—is difficult to solve even with the most advanced systems management tools, largely because software is closely bound to hardware.

In spite of the considerable achievements in hardware standards, significant differences in server systems still exist—usually caused by newer technologies that are still maturing through the standardization curve. Differences in components such as network interfaces, disk subsystems, processor families, and storage configurations can result in an OS configuration that is tailored to the hardware configuration. If an administrator removes the disks from one system and inserts them into another system, chances are the second system will not even start, much less function stably, unless its configuration is identical to that of the first system.

Decoupling the hardware/software stack

Consider that if all server hardware were identical, then a software stack—an OS and its hosted applications—could quickly move from system to system without alteration. In such a world, software and hardware are *decoupled,* and the hardware is truly interchangeable and disposable. Servers can be repurposed within minutes by simply copying, or *cloning,* an archived, pre-patched software stack. Decoupling also enables IT administrators to balance workloads more effectively by reassigning hardware to software dynamically.

However, standard hardware does not imply *identical* hardware. Standards help to manage the technology differences between systems in such a way that their benefits can be easily accessed by all parties. For computer components, standards often take the form of an interface specification. For example, TCP/IP—a standard for network communication—is an interface for accessing a network that takes on many different physical forms. Yet the format and function of TCP/IP datagrams and the protocol are the same, regardless of the physical network.

Virtualization applies the same notion of abstraction to computer hardware by creating a new interface, the *virtual machine,* which becomes the standard interface for software stacks to execute on standardized, but different, underlying physical hardware. This is not a new idea; however, it is being applied at a different point in the combined hardware and software stack (see Figure 1).

Early computers ran application software directly on the hardware—a tedious programming model that motivated the development of the modern OS. Operating systems shield applications from the underlying hardware by abstracting the hardware as an array of managed resources and services. Although applications have more freedom from hardware in this current model, they still depend on the OS for a file system that holds the application's persistent state—that is, the application's configuration and data. The new hardware-virtualization programming model bundles applications and operating systems together (including the file system), allowing them to move as a unit across various physical platforms.

Understanding the basics of virtualization

A virtual machine (VM) is a software-created construct that is functionally equivalent to physical hardware, at least from the perspective of its ability to execute software. A VM has the same features as a physical system: expansion slots, network interfaces, disk drives, and even a BIOS. Like physical systems, VMs can be powered up and powered down. When the VM is powered up, an OS runs on the virtual disk. Also called the VM's *image*, *stack*, or *container*, the virtual disk is actually a large file stored on the underlying physical system. Administrators can transfer a VM from one physical system to another by moving its image, or by storing its image in a shared location accessible by both physical systems.

Figure 2 depicts two options for implementing VMs. Both approaches divide the hardware/software stack into three layers: application software layer, virtualization layer, and physical

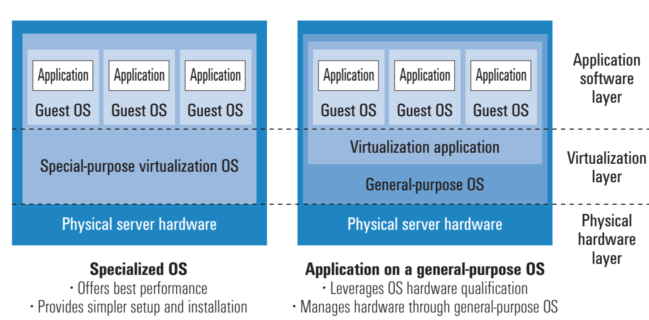

Figure 2. Two methods for implementing hardware virtualization

hardware layer. The two implementations differ only in how they implement the middle virtualization layer.

Specialized OS. The virtualization layer is a special-purpose OS that is installed directly onto the physical hardware, as one would install other operating systems. This OS can manage only the resources of the physical system, to share these resources among the VMs that it hosts. It does not, for example, have its own general-purpose file system or support user processes. Because no generic hosting OS exists, this approach offers the better performance of the two alternatives and has fewer setup and installation steps. Because it is not a mainstream OS, however, driver availability may be limited for some physical components. Products such as VMware ESX Server use this approach.

Application on a general-purpose OS. The virtualization layer is the combination of a general-purpose OS, such as a Microsoft Windows® or Linux® OS, and a virtualization application. Because it leverages a preexisting, full-service OS, this approach offers broader compatibility with hardware and can take advantage of existing management infrastructures already created for these environments. However, because the virtualization layer depends on an intervening OS, performance can be noticeably degraded. Products such as Microsoft Virtual Server and VMware GSX Server™ software use this approach.

Realizing the benefits of virtualization

Hardware virtualization can offer enterprise computing environments benefits such as consolidation, normalization, isolation, replication, and relocation.

Consolidation. This practice involves converting many physical servers to VMs and hosting the VMs on a reduced number of physical systems. Consolidation saves space, but that is only one of many areas in which enterprises can benefit significantly from hardware virtualization.

Normalization. This ability shields software from hardware peculiarities and change. For example, consider a driver update in a heterogeneous OS environment including Windows, Linux, and Novell® NetWare® platforms. Without virtualization, three or more driver versions must be identified and installed. However, with virtualization, only one driver must be found and installed at the virtualization layer—guest operating systems and applications continue to use generic drivers with their virtualized generic hardware. Legacy operating systems that are no longer supported by the manufacturer can execute on newer physical hardware, which enables delayed migration of legacy systems.

Isolation. Some applications do not behave well in a shared, multi-programmed environment, either because of conflicting configuration requirements, resource domination, faulty code that can crash the entire system, or other side effects. Virtualization can be used to partition applications into their own private OS containers on a single physical system so that resource allocations can be enforced at the virtualization layer. If an application crashes its own VM, it does not affect other VMs.

Replication. Quick deployment of software stacks becomes a trivial task using a virtualization approach. First, administrators create reference software stacks for various applications—Web servers, mail servers, or file servers, for example—and then copy these images to a read-only archive, or a *gold master library*. Later, when a particular type of server is needed, administrators can clone the image from the archive and place it on a system that has spare capacity, usually within a matter of seconds.

Relocation. For the same reasons that virtualization facilitates replication, the decoupling of hardware and software enables the software stack to move freely between available virtualized platforms. At the very least, administrators can shut down a VM, move its image to another physical system, and then bring the VM back up. More sophisticated mechanisms allow for executing VMs to move in real time, without shutting them down and without affecting service to users. This mobility offers the greatest potential benefit to the scalable enterprise.

Moving mindsets

Hardware virtualization is not new to computing. However, it is a recent development in low-cost Intel architecture–based systems, for which it holds great promise for mitigating the management problems associated with scale-out architectures. The main challenges that virtualization technology poses are, surprisingly, not of a technical nature. Virtualization of Intel processor–based systems is best known as a tool for space consolidation, especially in test and development environments—not in production environments. Consequently, many IT administrators familiar with virtualization technology may have preconceived notions about its applicability in a production environment—and be unaware of the numerous benefits beyond simple consolidation.

Although it does not solve every systems management problem, virtualization is ready to move into production in select large-scale enterprise computing environments. Prime candidates include server farms hosting a diverse mix of lightly to moderately loaded services. Many early adopters of virtualization technology in enterprise environments are reporting great success. While traditional systems management tools are evolving to make it easier to manage non-virtualized platforms, they are also incorporating virtualized platforms into the aggregate data center view. Virtualization has matured to the point where it provides the required stability, and the benefits are extremely compelling. Although some may try to minimize the scope of this impact, most would agree that hardware virtualization has a definite role to fulfill in the scalable enterprise.

J. Craig Lowery, Ph.D. (craig_lowery@dell.com) is a senior engineering development manager in the Dell Product Group—Enterprise Solutions Engineering. His team is currently responsible for developing products that realize the Dell vision of the scalable enterprise. Craig has an M.S. and a Ph.D. in Computer Science from Vanderbilt University and a B.S. in Computing Science and Mathematics from Mississippi College.