


Fast Application Deployment *with* VMware *in the* Cloud

by Silas McCullough and Patrick Sklodowski



Imagine that your firm's top rainmaker calls you into her office and insists that she needs special software in order to entertain new mobile communications business from her most important client. You quickly ascertain that this must-have application will require a substantial backend, including a couple of dedicated servers, a database and lots of storage. Your current environment doesn't have the capacity for these new systems, and several weeks will be needed to test the software and put it into production. In the meantime, a competing law firm managed to get the same system up and running, and they snagged your firm's client.

The big question is how did the competing firm manage to launch the application before you? They leveraged their private Cloud offerings, which are built upon server virtualization technology from VMware. The new services were provisioned quickly and delivered to the client through the Internet.

Cloud Computing and Virtualization

Cloud computing holds the promise of providing application services in a more flexible and timely manner, and without much hassle. Clouds can be public — where the services are available to anyone over the Internet — or they can be private — where the Cloud is restricted to certain users and typically delivered in a more secure manner. In either case, the behind-the-scenes establishment of Cloud services requires a highly adaptable, responsive infrastructure. Server virtualization fits the bill and stands at the heart of providing Cloud-based application services.

VMware server virtualization, in particular, has become a key building block in establishing Cloud services. Applications run on hardware-independent servers called “virtual servers,” which exist within a set of physical servers that provide the memory, processing and storage resources for the virtual servers. In short, a physical server may host several guest servers, providing diverse services such as e-mail, database and litigation support. Since the application server is liberated from direct dependency upon physical hardware, there is no longer the one-to-one relationship between the application running on a server and the server box. Instead, using VMware virtualization technology, the physical server becomes a pool of resources available for any number of application servers.

How Does VMware Virtualization Work?

VMware virtualization begins with software called a “hypervisor” being installed onto a physical server. It establishes a base environment that allows multiple application servers to be loaded onto the physical server. More importantly, it permits those multiple servers to run concurrently; thus, sharing the resources of the physical server. Since these servers do not directly interact with and control the physical hardware, they are considered virtualized.

A virtual server provides the same services as the more traditional approach so the end-user sees no difference. From the system administrator's point of view, a virtualized server comes with many operational benefits, and for the business-minded, there are cost savings since fewer physical servers need to be purchased and maintained.

The process of establishing a virtualization infrastructure entails first specifying the application servers that are required for the production systems. These specifications account for the number of processors, memory sizing and disk space — just as you would currently specify for a physical server. With these specifications in hand, the next step is to size the virtualization environment. Sizing combines all the parameters to arrive at a number of physical servers, along with their own specifications that accommodate the application servers, while aiming for acceptable performance goals. These physical servers become the host systems, each running more than one

virtualized application server; therefore, the hosts should have similar hardware specifications, be powerful and contain ample memory.

It is especially important to correctly specify disk space since space must be allocated not only for the application data, but also for the virtual servers themselves. Disk performance must also be taken into account because the speed of disk access will impact performance while running application servers concurrently. Each virtual server competes for the underlying physical disk, so this disk must be sized appropriately to be responsive to the needs of all of the virtual servers. Under-specifying the disk performance will lead to application slowness, which is especially hazardous for disk-intensive applications such as e-mail and database systems.

“Using VMware virtualization technology, the physical server becomes a pool of resources available for any number of application servers.”



“Private VMware Clouds provide real disaster recovery solutions.”

Conceptually, a virtual server can be thought of as a single file, albeit a very large one, that VMware interacts with to provide an operating application server. Just like a physical server, a virtual server is powered on, but it's through the VMware software interface. During power-on, VMware reads the file and, from its contents, starts up a running application server within the physical host's resource environment. The entire process looks very similar to starting up a physical application server, complete with booting and startup dialogs — but viewed through the VMware interface.

What Are the Benefits of Virtual Servers?

Modularity and portability are key benefits of a virtual server. The virtual server's program files and parameters make up this file, which resides on storage that the VMware system can access. Optionally, application data may also be stored within this file or elsewhere, making an application server's disk structure highly modular. Because of the sizes of the virtual server files, they are often stored on external storage provided through a storage area network (SAN). It is possible to store them on local disk drives, but space limitations quickly restrict capabilities and options. Determining data storage is an important consideration in designing a private Cloud virtualization infrastructure.

By placing virtual server files on external SAN storage, the firm can take advantage of VMware's more advanced features such as the ability to dynamically distribute resource utilization across servers, and providing local high-availability and application continuity to all virtual servers. VMware provides the ability to move the running virtual machine to another physical host with no effect on the running application server. This process can be performed manually or automatically by VMware should it detect high utilization occurring on one physical server. Should a physical server suffer a hardware issue and crash, VMware will automatically boot the affected virtual machines on other physical servers. VMware, in the firm's private Cloud, extends these abilities to all servers in the environment without the need for complex and expensive clustering technologies required for physical servers.

The virtual server is fundamentally a file, so it can be moved between storage locations. VMware allows this move to occur with no downtime required of the server. The ability to move the running virtual machine to different host servers, and to move the underlying file to another location, has the primary benefit of extracting the virtual machines from the physical environment. Hardware and storage upgrades to the private Cloud can occur for most virtual machines in the background with no downtime or impact to application services.

Copying a virtual server provides even more operational, as well as strategic, advantages. A copy can:

- **Be used to test an upgrade without affecting the production environment**
- **Become a clone for additional virtual servers**
- **Act as an offline server for disaster recovery purposes**

In addition to an outright virtual server copy, VMware provides “snapshot” copy capabilities as a powerful virtual server file management tool. This is a point-in-time copy of the virtual server that can be reloaded later if desired. It is quick and can be managed directly through the VMware software. Snapshots are useful in reducing the time required to try out new options or install new software since changes made to the virtual server are thrown out when returning to the snapshot. For example, a snapshot of the virtual server is typically taken before applying a service pack. If the service pack causes problems, then you simply go back to the snapshot. It does not have to be uninstalled, and there is not a long wait to retrieve a backup.

Private VMware Clouds provide real disaster recovery solutions. Another benefit of the independence between the virtual machine and physical hardware is the ability to move and power on the virtual machine in any VMware Cloud, without the difficulties of restoring a physical server. In its most simple form, during a disaster, the firm's virtual machine files could be restored from backup into another VMware Cloud hosted at a separate location or by a service provider. Unlike a restore to physical hardware, the machines will always boot

Fast Application Deployment with VMware in the Cloud

with data intact. Upon configuration of the network properties and connectivity to the Internet or wide area network (WAN), all of the firm's servers will be accessible to users. With more advanced technologies, like SAN to SAN replication and VMware Site Recovery Manager, disaster recovery and application continuity procedures can be heavily automated, allowing full data center recovery in a matter of hours.

Utilizing VMware virtualization and the Cloud can help your firm drastically reduce the amount of time it takes to get a new application up and running for a client. This gives your firm an edge over the competition and, by leveraging the capabilities presented by VMware, not only can you provide tools quicker, you can do it with the ease of mind that it's easy to return to the previous functioning level when a problem is encountered. **ILTA**

This article was first published in ILTA's March 2011 issue of *Peer to Peer* titled "Communications Technologies" and is reprinted here with permission. For more information about ILTA, visit their website at www.iltanet.org.



Silas McCullough is the General Manager of mindSHIFT Technologies' Legal Professional Services Group, which is committed to delivering high-quality technology services to mid-size and large law firms. During his 14 years in the legal IT industry, Silas has been instrumental in developing new IT solutions for law firms ranging from incorporating complex technologies to implementing cloud computing services. He can be reached at Silas.McCullough@mindSHIFT.com.



Patrick Sklodowski is the Managing Director of mindSHIFT Technologies' Philadelphia, Pennsylvania office and mindSHIFT's Messaging Practice Group Leader. With more than 15 years of industry experience, Patrick oversees a veteran team of multi-disciplined integration specialists. In addition, Patrick is responsible for maintaining the company's legal IT messaging expertise, while also identifying and developing the skills and practices required to provide law firms with the latest technology innovations. He can be reached at Patrick.Sklodowski@mindSHIFT.com.