Implementing a Fully-Virtualized Data Center

By Ken Won

Most IT organizations have implemented server, storage and/or network virtualization to some degree, but to fully realize the benefits of virtualization, IT organizations need to consider adopting all three in an integrated manner. Server virtualization for x86 servers has been around for about a decade, and is relatively mature. Storage virtualization is not quite as mature, but numerous storage virtualization products are available today. And while network virtualization in the form of VLANs and VPN has been around for many years, new types of network virtualization are just now being introduced to the industry.

In this article, we'll take a look at virtualized servers, storage, and networking, and see how automated network switching helps unify these environments into a cohesive whole.

Server virtualization

Virtualization was initially developed for mainframe computers by IBM in the early 1960s as a way to maximize utilization of expensive hardware. The availability of commodity servers in the 80s and 90s drove hardware prices down to the point where virtualization was unnecessary, resulting in the practice of putting one application on one server. Unfortunately, this led to very inefficient use of servers, with average utilization at 15 percent or less as servers were heavily over-provisioned to support peak usage and future growth.

VMWare introduced server virtualization software for Windows machines in 1999, making it possible to easily combine multiple applications on individual servers and increasing server utilization rates to as much as 80 percent. With solutions like VMWare, servers can be divided into virtual machines (VMs) through the use of hypervisors, each of which functions as an independent operating environment within the server. Today, there are a number of companies that offer server virtualization products including Citrix, Microsoft, Red Hat and Oracle.

Virtualization has significantly reduced server counts and thus saved on space, power, and cooling resources. Today, roughly 80 percent of large data centers have implemented server virtualization. It has enabled on-the-fly provisioning of resources for specific applications, leading to such trends as software as a service (SAAS) and cloud computing.

However, server virtualization also brings about a quantum increase in complexity – rather than a few dozen servers, administrators must now deal with hundreds or thousands of VMs. Additionally, applications require the complete IT stack to be virtualized, not just the server. Storage and the network must be virtualized as well. While hypervisors have become more sophisticated over time, the challenge of deploying and managing a complete virtual environment remains.

Storage

Virtual server environments are extremely flexible – it's much easier to deploy new applications with VMs – but the data center also requires an equally flexible storage infrastructure to make this workable. As VMs replaced individual servers tied to specific storage resources, storage virtualization became necessary in order to avoid stranding expensive storage space on

underutilized storage servers. Storage virtualization can triple the rate of storage utilization so companies get more for their storage dollars.

Storage virtualization involves techniques such as data migration, caching, snapshots, and thin provisioning. Each vendor enables storage virtualization in a different way, using host-based software, storage devices, or network hardware. Virtualization can be symmetric (in-band), in which the virtualization function caches data for access by servers, or asymmetric (out-of-band) in which device drivers lookup the metadata about a file and then allow the server to access it directly from its location. Either way, the result is a flexible pool of storage that can easily be managed and made available to VMs and their applications.

Storage virtualization enables easy data migration and replication among storage locations while significantly increasing the rates of utilization for storage devices by enabling many servers to use the storage on any given storage device. The main drawbacks to storage virtualization are a lack of industry standards (each vendor rolls its own approach) and the inability to easily change from one virtualized environment to another – it is much easier to implement a virtualized storage environment than to de-implement it.

Network virtualization

The link between servers and storage is the network. What we are discovering today is that virtualizing the network eliminates manual processes, costs, errors, and inefficiencies that are caused by the demands of virtualized servers and storage. To support virtualized environments, the network needs to change from being a static data interconnect to a dynamic resource that can be automatically provisioned along with VMs and virtualized storage. New products are being brought to market to support this capability.

Network virtualization has been around a long time, just like server virtualization. VLANs and VPNs have allowed network administrators to create virtual networks within the physical network with different connectivity, security and QoS (quality of service) profiles. As customers are now demanding more flexible IT environments, virtualized networks need to support dynamic server environments that can be dynamically created, resized and removed in conjunction with the server and storage environment.

For example, new network virtualization capabilities introduce a new degree of automation into the process of managing network connectivity, security and QoS. Whenever a VM is moved, it is necessary to move the VLANs and port profiles associated with that VM. Since today's data centers are managed via separate server, storage, and network administration groups, a request must be made to the network group to move the VLAN and port profile to support a VM migration. However, this can take hours to days in a large organization, which really throws a wrench into the idea of dynamic moves and changes promised by server virtualization. Every organization is trying to do more with less, but without a dynamic, virtualized network linking servers and storage, doing more takes much more time and costs a lot more money. Requiring network administrators to manually monitor the network for changes or to manually migrate VLANs and port profiles places a huge additional burden on teams that are already up to their necks in work.

A virtualized environment changes the requirements for the network infrastructure. Virtualized environments have more traffic on the network and greater peaks and valleys in the traffic load. This means the network must be designed for higher capacity and the ability to manage through

larger peaks and valleys (with bigger buffers on switches). It also means that a high degree of network automation should be implemented in order to ensure that dynamic server and storage changes are not slowed down by the network itself.

Since the network provides connectivity to storage and servers, it is important for the virtualized network to be able to interoperate with multiple server and storage virtualization products and technologies. For example, the network should support hypervisors from multiple vendors, support management connectivity to multiple virtualization management tools and support networking industry standards. This relieves IT from being being locked into a particular vendor or architecture.

Implementing an automated, dynamic network layer in the data center stack is a new idea, but several vendors are leading the way toward this practice with automated switches. Automated switches can dynamically move VLANs and port profiles, notify administrators of looming problems, and otherwise handle a lot of the heavy lifting required in virtualized server and storage environments that would otherwise fall to network administrators.

Organizational Issues

The advent of virtualization technology is causing organizational changes. Traditionally, companies have separate server, storage and networking administration groups. Each group is responsible for their part of the infrastructure and the groups often have minimal interaction. This model has worked in the past because the IT infrastructure has been static with well-defined interfaces between servers, storage and networking.

However, with virtualization, this is changing. The migration of a VM from one server to another impacts not only the server environment, but also the storage and network. Under the traditional management model, the server administrator would work with the storage and networking administrators to facilitate this change. In today's virtualized environment, VMs need to be moved immediately, requiring integration and automation between the server, storage and networking. Virtualization administrators must be able to migrate resources on demand and not wait for the other administration groups to manually change resources

Virtualization technology has also moved the traditional boundaries of technology, causing problems with administration. For example, many to today's hypervisors include virtual network switches, or vSwitches. vSwitches are typically managed through the Hypervisor management console, usually by the virtualization manager. This means the virtual part of the network is managed by the virtualization administrator while the physical part of the network is managed by the network administrator. This raises the potential for inconsistent network policies, which could cause problems with network security or performance. New standards such as EVB (Ethernet Virtual Bridge) are being created to help address these issues, but the standards are not yet complete, requiring administrators to rethink how they manage virtual environments.

Conclusion

While the drive toward a fully dynamic data center is causing organizational issues and growing pains as administrators adopt new technologies, the end result is well worth it. A truly dynamic data center will, in the end, be easier to manage and far more response to the demands of end users. Network automation ensures that server and storage management changes flow through the data center automatically, making the network itself as virtual and dynamic as the servers and storage attached to it.