

Virtualization Technologies

By

Genesis Multimedia Solutions
info@thecreation.com

June 9, 2007

V1.0

Purpose

The purpose of this whitepaper is to provide a general overview of virtualization technologies for someone who is new to virtualization, or for someone who is interested in the benefits and disadvantages of virtualization.

What is virtualization?

Virtualization is the practice of creating an abstraction layer between a computing resource and another computing resource or user.

An abstraction layer allows, for instance, a user to see multiple *logical* representations of a computing resource even though there is only a single *physical* resource.

This paper concentrates on platform virtualization, where an entire *physical* machine is partitioned into multiple logical *virtual* machines.

Wikipedia has a great explanation of other virtualization concepts here:
<http://en.wikipedia.org/wiki/Virtualization>

Physical Machine Virtualization

A physical machine consisting of CPUs, memory, disk, and network computing resources can be virtualized into logical *virtual* machines consisting of their own virtual CPUs, virtual memory, virtual disk, and virtual networking resources. This practice has been used in larger mainframe and super-computing systems for many decades, but is only now starting to show maturity in smaller computing systems such as x86-based systems using AMD and Intel processors.

Multiple virtual machines share the resources of the physical hardware that they run on. The abstraction layer that provides the physical resource scheduling is called a hypervisor. A hypervisor can be software or hardware. AMD and Intel are now providing hypervisor *assistance* in hardware.

The basic purpose of a hypervisor is to schedule the use of single or multiple physical computing resources by single or multiple logical computing resources. The complexity of this operation is dependent on the capabilities of the hypervisor, including:

- logical resource priority assessment
- guaranteed resource service levels
- dynamic addition/removal of logical resources
- moving of logical resources without interruption of service to the logical resource
- grouping of logical resources for prioritization
- de-duplication of resources

Physical Machine Virtualization Components

CPU virtualization abstracts physical CPUs into logical CPUs. For example, 10 logical CPUs could be created on a physical machine where only 2 CPUs exist. When a logical CPU is idle, the physical CPU can be assigned to other logical CPUs that have work to do.

Memory virtualization abstracts physical memory into logical memory. The concept of virtual memory inside an operating system, where more logical memory can be created than physical memory, uses disk storage as a temporary location of memory when physical memory is exceeded. This concept can be used at a higher layer, with more capabilities when used in a virtual machine environment.

For example, one virtual machine's memory can be completely isolated from other virtual machines' memory, even though they share the same pool of memory. De-duplication of memory can be performed, allowing more efficient use of memory. If the same operating system is installed in multiple virtual machines running within the same memory pool, significant memory savings can be made by de-duplicating the portions of memory where the operating system resides.

Network virtualization abstracts the physical network into logical networks. This concept has been used for many years where multiple VLANs (Virtual Local Area Networks) are transported over a single physical medium. A physical network cable can be used to transport packets of information from one end to another, but with VLANs being configured, each packet is assigned a VLAN ID. The hardware or software on each end of the cable routes each packet based on the VLAN ID, completely segregating each VLAN's traffic.

Virtualization of physical machines requires the virtualization of network resources, where each virtual machine is assigned a virtual network interface, an abstraction of the physical network interface. VLAN IDs are commonly assigned to all traffic coming in and out of a physical network interface, allowing the assignment of specific VLAN ID(s) to a single virtual machine's virtual network interface.

Disk virtualization abstracts physical disks into logical disks. This concept has been used for many years in RAID controllers, where a physical machine sees a logical volume; however, multiple physical disks are being abstracted to *appear* as a single physical disk.

Disk virtualization can also abstract a physical disk into logical file-based disks. A single file can be used as a virtual machine's virtual disk. This provides a user with adding/resizing/deleting entire virtual disks with the ease of file-based operations.

Virtual disks can perform operations similar to a SAN (Storage Area Network) by abstracting physical disks into logical volumes while providing additional services such as prioritization, dynamic online resizing, expansion, moving, and snapshots

Where virtualization should be used?

With all of its great features, should virtualization be used everywhere? In the near-future, the answer to this question will be a confident yes, but for now, there are currently some performance disadvantages.

The combination of the effort of the hypervisor as well as the management functions can reduce the overall performance of a virtual machine compared to a physical machine. Manufacturers are continuing to provide hardware support for various hypervisor and management functions for these efforts, but there will always be some performance reduction compared to a physical machine.

The question of where virtualization should be used is dependent on the performance *requirements* of the CPU, memory, disk, and network resources. Virtualization technology is extremely efficient, and can provide great benefit to both low-load and high-load environments. As more logical resources are scheduled by the hypervisor, the more complex the scheduling gets, and thus the slower the performance of each resource. If the number of logical resources is controlled, a virtualized machine can perform very close to a physical machine.

Currently, I/O operations such as disk and network activity are the common bottlenecks in virtualized environments. This is why database and firewall functions are commonly assigned to physical hardware in high-load environments. However, if the demands of a database and firewall are very high, they are usually running in redundant and load-balanced infrastructure, which could be virtualized redundant and load-balanced infrastructure. If loads are high, it is important to have a plan in-place to scale out since scaling up is only as successful as the fastest hardware available.

Benefits of virtualization

Efficiency is a common reason to deploy a virtualization solution. Using idle resources to perform *work* is the goal. By performing the same work on fewer physical resources, fewer energy resources are required, lowering energy bills for both running unnecessary servers as well as the cooling required. Physical space is also conserved, lowering the number of data center racks required.

High Uptimes can be achieved by moving logical resources to different physical resources while remaining online, or if a physical resource fails, starting all logical resources on a different *functioning* physical resource. Maintenance and upgrades of physical resources can be performed without downtime of the logical resources.

Scaling Out can be achieved with proper planning. Instead of waiting until the need arises, scaling out to two or more virtual servers from the beginning, all running on the same physical machine, will provide an easy transition from a single physical machine to multiple physical machines, all without downtime.

Easy Configuration of servers is provided by the many aspects of virtualization. The management of physical hardware is eased since all aspects can be configured with the click of a mouse. For example, to clone an existing virtual server, you simply *copy* the virtual disk file of an existing server. To add logical hardware, an interface can be provided, a user can simply choose what hardware they wish and create a new virtual machine. Depending on the environment, all logical hardware can be added or removed while the virtual machine is running, all without touching the physical hardware. Virtual machine templates can be made and easily duplicated for large load-balanced deployments or for a large number of copies of similar environments without the need to build each environment from scratch.

Snapshots are virtual copies of a logical disk. The logical disk can be reverted to a point-in-time where a snapshot was taken. If a snapshot is taken before software upgrades or additional software is added to a virtual machine, and a failure occurs, a simple click of a button will allow the virtual machine to be recovered back to the point-in-time when the snapshot was taken, allowing multiple tests to occur without the fear of damaging a virtual machine. Backups may also be performed out-of-band, similar to the way a SAN snapshot can be taken of a volume, allowing a backup to be made from the snapshot, and not the “live” data, yet without downtime to the virtual machine.

Replication is often used for either high-availability or disaster recovery scenarios. It allows an exact copy of all data to be copied to a local or remote location without affecting the online status of the source. If the source fails, the replicated copy can be brought online.

All replication environments require disk virtualization to accomplish this task, where either the disk controller hardware or operating system software performs the replication of a logical volume. Disk virtualization is often a native feature of a virtual machine, and thus replication is possible by intercepting the reads/writes of the disk activity and sending them to a replicated location.

Reporting on virtual technology performance statistics can be provided in-place of many of the available statistics provided by an operating system, but out-of-band, meaning that statistics are gathered from the virtual hardware, and not the operating system. This would allow, for instance, the ability to measure disk activity during the booting of an operating system. Everything from CPU, memory, disk, and network utilization can be performed outside of the operating system, providing an picture below the operating system layer.

Consistency is a very important feature of virtualization. Virtual hardware rarely changes, and so the constantly changing world of drivers for new hardware is hidden from the user. Drivers for the virtual hardware are provided by the virtualization hardware or software provider. Moving a virtual machine from one physical machine to another is a seamless process as long as the same virtualization layer is in-place, regardless of whether one physical machine is made of completely different hardware than the first. Most virtualization technology today requires similar CPUs when moving virtual machines, but this is mostly done due to the inefficiency of on-the-fly conversion of CPU instructions for different CPU types.

Disadvantages of virtualization

Performance is the most popular disadvantage in discussions of whether to use virtualization technologies. Performance is sometimes a major concern, especially in highly visible systems where high latency is widely noticed, and potentially financially harmful. Virtualization technologies are often used to consolidate existing infrastructure and not used in large-scale deployments of a single application that meets these requirements. Large high-availability and highly-scalable systems are generally cheaper and faster than large virtualized environments.

Server Sprawl is a common issue with virtualization technologies. This occurs when administrators feel that the separation of applications by complete operating system environments is more secure and reliable than separation of applications by the security within an operating system. By installing an application in its own dedicated operating system, upgrades can be isolated to either that single application or operating system without affecting other applications or operating systems. Availability is often improved by this technique, however, the management and cost of the operating systems produces a burden not predicted. With the ability to copy virtual machines so easily, this technique is made effortless, without thinking of the consequences of the management and cost burdens.

Complex Deployment and Maintenance is another common issue with virtualization technologies. When a large number of logical resources are consolidated onto a small number of physical resources, the consequences of downtime increases. To compensate, redundant physical hardware configurations are often used, increasing the complexity of deployment. This does not necessarily mean the physical hardware deployment is any more complex than a non-virtualized environment, but it does add a layer of complexity onto typical physical deployments. With the added features that virtualization provides, the possibilities increase, thus complicating the deployment options and management of the environment.

Data Center virtualization

Data centers typically consist of a large number of heterogeneous servers, separating customers physically by racks or cages and logically by network switch ports.

Within this environment, small customers typically have servers with large amount of unused local storage and wasted idle computing resources. Larger customers typically increase their efficiency by utilizing shared-disk systems (SAN or NAS systems); however, this does not provide any efficiency of CPU, memory, or network resources. Virtualization can provide help with the efficiency of wasted CPU, memory, disk, and network resources of both customers while at the same time providing improved performance of each of these resources to each customer.

If an entire data center was virtualized, all of the benefits described in this paper would apply to the entire data center, including added reliability to both types of customers.

Conclusion

As virtualization technologies continue to evolve, both in terms of the hypervisor as well as the management components, the utilization of computing resources will continue to come closer to 100% utilization with the ability to scale out and in on-demand as physical resources are added and removed.