What is Virtualization?

Virtualization allows multiple operating system instances to run concurrently on a single computer; it is a means of separating hardware from a single operating system. Each "guest" OS is managed by a Virtual Machine Monitor (VMM), also known as a hypervisor. Because the virtualization system sits between the guest and the hardware, it can control the guests' use of CPU, memory, and storage, even allowing a guest OS to migrate from one machine to another.

Background

Over the last 10 years, the trend in the data center has been towards decentralization, also known as horizontal scaling. Centralized servers were seen as too expensive to purchase and maintain. Due to this expense, applications were moved from a large shared server to their own physical machine, often using commodity hardware. Decentralization helped with the ongoing maintenance of each application, since patches and upgrades could be applied without interfering with other running systems. For the same reason, decentralization improves security since a compromised system is isolated from other systems on the network.

However, decentralization's application sandboxes come at the expense of more power consumption, less physical space, and a greater management effort which, together, account for up to $10,000 in annual maintenance costs per machine[1]. In addition to this maintenance overhead, decentralization decreases the efficiency of each machine, leaving the average server idle 85% of the time[2]. Together, these inefficiencies often eliminate any potential cost or labor savings promised by decentralization.

Virtualization is a modified solution between centralized and decentralized deployments. Instead of purchasing and maintaining an entire computer for one application, each application can be given its own operating system, and all those operating systems can reside on a single piece of hardware. This provides the benefits of decentralization, like security and stability, while making the most of a machine's resources.

1 http://www.xensource.com/files/xensource_wp2.pdf

2 Ibid.

Why It Matters

As virtualization disentangles the operating system from the hardware, a number of very useful new tools become available. Virtualization allows an operator to control a guest operating system's use of CPU, memory, storage, and other resources, so each guest receives only the resources that it needs. This distribution eliminates the danger of a single runaway process consuming all available memory or CPU. It also helps IT staff to satisfy service level requirements for specific applications.

Since the guest is not bound to the hardware, it also becomes possible to dynamically move an operating system from one physical machine to another. As a particular guest OS begins to consume more resources during a peak period, operators can move the offending guest to another server with less demand. This kind of flexibility changes traditional notions of server provisioning and capacity planning. With virtualized deployments, it is possible to treat computing resources like CPU, memory, and storage as a hangar of resources and applications can easily relocate to receive the resources they need at that time.

Three Approaches

Virtualization comes in a variety of implementations. In its basic form known as "full virtualization" the hypervisor provides a fully emulated machine in which an operating system can run. VMWare® is a good example. The biggest advantage to this approach is its flexibility: one could run a RISC-based OS as a guest on an Intel-based host. While this is an obvious approach, there are significant performance problems in trying to emulate a complete set of hardware in software. Even with painstaking optimization, it is very difficult to get useful performance from a fully virtualized environment.

At the other end of the spectrum is the Single Kernel Image (SKI), in which the host OS spawns additional copies of itself. This kind of virtualization can be found in Swsoft Virtuozzo and Sun® Solaris® Zones. SKI can be thought of as "lightweight" virtualization. While this approach avoids the performance problems with pure emulation, it does so at the expense of flexibility. It is not possible, for instance, to run different versions or even different patch levels of a particular operating system on the same machine. Whatever versions exist in the host, that same software will be provided in the guest. SKI also sacrifices the security and reliability provided by other virtualization methods. If the kernel is exploited, all OS instances resident on the system will be compromised.

"Paravirtualization," found in the XenSource® open source Xen product, attempts to reconcile these two approaches. Instead of emulating hardware, paravirtualization uses slightly altered versions of the operating system which allows access to the hardware resources directly as managed by the hypervisor. This is known as hardware-assisted virtualization, and improves performance significantly. In order to retain flexibility, the guest OS is not tied to its host OS. Drastically different operating systems can be running in a hypervisor at the same time, just as they can under full virtualization. In this way, paravirtualization can be thought of as a low-overhead full virtualization.

Xen

With the release of Xen 3.0, virtualization reaches maturity. Xen is the first virtualization solution to support Intel's VT technology which permits each guest OS to run at full processor speed, with only 0.5% to 3.5% overhead typically incurred by the virtualization process. Guests can be migrated from one machine to another in less than 100ms. Through the hypervisor, operators can control the use of CPU, memory, block, and I/O devices dynamically.

Possibilities of Virtualization

Once Xen's basic virtualization features are understood, many exciting applications become possible. In the same way that virtual LANs allow administrators to ignore the physical layout of their networks, virtualized operating systems allow administrators to ignore their physical installation. The hardware in the data center becomes truly commoditized – hardware upgrades can occur seamlessly, without the OS or application realizing that its host machine has been changed. Downtime can be reduced dramatically.

Administrators no longer have to wait for every application to be certified on a new operating system release before an upgrade: just migrate the guest OS, and everything works as before. During regression tests, a testbed can be created or copied easily, eliminating the need for dedicated testing hardware or redundant development servers.

In security, virtualization is an elegant solution to many common problems. In environments where security policies now require systems separated by a firewall, those two systems could safely reside on the same physical box. In a development environment, each developer can have their own sandbox, immune from another developer's rogue or runaway code.

If virtualization is tied to a system monitoring solution, or to a provisioning and management tool like Red Hat® Satellite Server, systems could be migrated automatically to better-suited hardware during periods of peak use or maintenance. Imagine a farm of servers which can be re-tasked in seconds, according to workload and time of day.

Xen Adoption and Red Hat

Given the current level of performance and maturity, and the possibilities it provides, Xen is the undisputed leader in open-source virtualization. Dozens of corporations and universities are involved in the project, including Red Hat, IBM®, Oracle®, Intel®, AMD®, Cisco®, and Veritas®. Red Hat has been an early adopter of Xen, an active contributor, and has already incorporated the code into its Fedora distribution. Xen is also a crucial component of Red Hat Enterprise Linux® 5.

Virtualization is only one part of Red Hat's larger strategy to commoditize each major computing component, and make it simple for administrators to bring computing resources to the application that needs them. With the advent of the Global File System (GFS), storage is an easily allocated resource and still delivers a high level of performance. Red Hat Network (RHN) allows administrators to treat systems as generic resources that can be easily installed, upgraded, retasked, and reallocated. Together, GFS, RHN and virtualization free an application from the technical and logistic constraints that are so familiar to the data center, realizing Red Hat's vision for scalable, flexible, reliable, and manageable enterprise computing.

Related Articles, Sites of Interest

XenSource: http://www.xensource.com/
Fedora + Xen Kickstart: http://www.fedoraproject.org/wiki/FedoraXenQuickstart
Xen Whitepaper: http://www.xensource.com/files/xensource_wp2.pdf
Virtualization 101, Brian Stein, Red Hat:
http://www.computerworld.com/softwaretopics/software/story/0,10801,108907,00.html
http://www.computerworld.com/databasetopics/data/datacenter/story/0,10801,108907,00.html