


Fault Tolerance for Virtual Environments – Part 1

March 2008

Not only do businesses today depend upon information technology (IT) for their very existences, but IT costs have become a major part of an enterprise's budget. As corporate data centers become bigger and bigger, often supporting thousands of servers, their costs for hardware, space, administration, and energy are rapidly increasing. In fact, their energy requirements sometimes bypass the available energy in their areas – remember the California brownouts of 2000?



Data Center Consolidation with Virtualization

However, a fortunate trend is evolving. Servers are becoming ever more powerful. Moore's Law states that server capacities will double every eighteen months,¹ and this trend not only has held for decades but is projected to hold well into the future. The result is that data-center servers are carrying less and less of their rated capacity. In fact, recent studies have shown that typical servers in a data-center environment that is governed by a one-application, one-server policy are running at only 10% to 15% of capacity.

If only we could harness this excess capacity, we could significantly reduce the number of servers in a data center by a factor of two, three, or even more. This would result in less hardware, less maintenance, less administration, less space, and less energy – in short, less cost by a large factor. This is the promise of virtualization.

¹ This is the common quote. Gordon Moore actually said that transistor density would double every two years.

Server Consolidation via Virtualization

Virtualization lets one physical server do the work of many. It does so by creating *virtual machines (VMs)*. A single physical server can host several virtual machines. We call this a *virtualized server*.

To an application, a virtual machine looks exactly like the physical server on which the application expects to run. However, the application does not have a physical server dedicated to it. Rather, it is running in its own virtual machine. Its virtual machine is sharing an underlying physical server with several other virtual machines. As a result, the utilization of physical servers in a data center can be increased from today's average of 15% or less to 70% or more. This workload consolidation can significantly reduce the number of servers required in the data center. Server consolidation is what provides the significant cost savings suggested above in hardware, maintenance, administration, space, energy, and so on.

Virtualization and Availability

But virtualization comes with a price, and that price is availability. If a physical server fails, it takes down only the application that is running on it. If the application is not mission-critical to the enterprise, this may be acceptable. However, if a virtualized server fails, it takes down the equivalent of several servers since each virtual machine hosted on the virtualized physical server fails. Thus, the failure of a virtualized physical server will take down many applications and is far more painful to the enterprise, especially if some of these applications are mission-critical.

Consequently, redundancy of physical servers in a virtual environment is necessary. Should a server fail, there must be a failover mechanism in place to rapidly move the failed virtual machines to functioning servers.

In fact, the failure consequences of a virtualized physical server argue strongly for physical servers that simply will not fail – at least not very often. This is the realm of fault-tolerant servers, which can survive any single fault as well as many cases of multiple faults.

In this multipart article, we describe today's virtualization techniques. We then look at the redundancy mechanisms that are available today to provide fault tolerance. Finally, we briefly review products providing these capabilities.


What is Virtualization?

Virtualization is an architecture in which access to a single underlying piece of hardware, like a server, is coordinated so that multiple guest operating systems (virtual machines) can share that single piece of hardware with no guest operating system being aware that it is sharing anything at all.² Simply put, virtualization allows a single physical server to be partitioned into multiple *virtual machines (VMs)* that can independently be used by *guest operating systems*.

An important characteristic of a virtual machine is that it is independent. It is totally isolated from the other virtual machines just as if it were running in its own separate physical processor. Any fault in an application or guest operating system in one virtual machine is completely transparent to the other virtual machines running on that physical processor and can have no impact on them.

² Bernard Golden, *Virtualization for Dummies*, Wiley Publishing Inc.; 2007.

This implies that there must be some kind of adjudicator that controls the access by the various virtual machines to the resources of the physical server - the processor, its memory, its data storage devices, and its I/O channels. This adjudicator is known as the *hypervisor*. The hypervisor traps guest operating system calls to the processor, memory, data storage devices, and network connections and allows only one virtual machine at a time to execute these calls. In effect, it is multiplexing the access of the various virtual machines to the underlying physical processor, thereby ensuring that each gets the resources it needs.



A Virtualized Server

The hypervisor typically provides a system management console that allows a system administrator to monitor the virtual machines and to start, pause, stop, delete, and otherwise control the VMs.

Virtualization is not new – it has been used in the mainframe world for over thirty years and was introduced by IBM for its System 360 in 1972. What is new is that virtualization is now becoming available to the world of industry-standard servers through products such as those from VMware and from Citrix (formerly XenSource).

Why is this technology just now becoming widely available? It is because of the standardization of industry-standard servers on a relatively cheap common chip architecture – the x86 class of microprocessors. This has allowed the development of hypervisors to virtualize a common hardware architecture rather than having to support multiple such architectures.

Why Virtualization?

We have alluded earlier to the many advantages of virtualization. They include:

- **Better Hardware Utilization:** As Moore's Law continues to predict, server capacity is increasing at a rapid rate. Yesterday's servers that were running at full capacity are today running only at a fraction of that. Virtualization allows a data center to use this excess capacity to run the load of many servers as virtual machines on a single physical server.
- **Server Consolidation:** Since virtualization allows the functions currently being performed by several physical servers to be consolidated onto one server, the data center requires fewer physical servers. If, for instance, the data center can consolidate on average the workloads of four current servers onto a single server, it can reduce its server count by 75%.
- **Less Hardware Maintenance:** The fewer the servers, the less is the maintenance workload. Fewer maintenance personnel are needed. Fewer spare parts need to be stocked. Fewer upgrades need to be made.

- Reduced System Administration: True, the administration of applications remains. However, fewer physical systems need to be administered; therefore, fewer administrators need to be employed. Typical industry experience is that data-center administration costs can be cut by 30% to 50%.
- Reduced Space Requirements: Large server farms can take up a lot of expensive space. By significantly reducing the size of the server farm, the space required to house the server farm is correspondingly reduced.
- Reduced Emergency Power Needs: UPS (uninterruptible power supply) requirements are proportionately less. This affects both the size of the diesel generators required and the capacity of the battery backup system needed to power the data center until the diesel generators kick in. In addition, diesel fuel costs are reduced, which may be significant for extended outages.
- Reduced HVAC Requirements: The amount of heating, ventilation and air conditioning for the data center is proportionately reduced. The same goes for lighting of the data center space. This affects not only the initial cost for these systems but also reduces the ongoing high-energy costs associated with powering these systems.
- Reduced Energy Costs: With fewer servers and reduced HVAC and lighting requirements, the amount of energy demanded by the data center is dramatically reduced. With energy prices rapidly escalating, this can be a significant operational cost savings. More important to some companies is the positive environmental impact that reducing energy consumption can have. Virtualization is green!
- Reduced Capital Costs: Less money needs to be invested in server hardware, data-center space, and HVAC and lighting infrastructure.
- Reduced Operating Costs: Less hardware maintenance and system administration, along with the significant savings in energy costs, result in dramatically reduced operating costs and a reduction in the total cost of ownership (TCO) for the data center.
- Improved Availability: If data storage is provided by network attached storage (NAS) or by a storage area network (SAN), today's virtualization products provide seamless failover of virtual machines from one physical server to another, thus eliminating planned downtime for hardware and software upgrades. Failover to another physical server in the event of a server crash is also supported to minimize unplanned downtime. Though these are capabilities similar to those for clusters, virtualization provides failover with little if any of the significant administrative attention required of clusters.

With all of these advantages, what are the downsides of virtualization? One is the hesitancy to learn a new technology. The other is convincing the manager who has been running his precious application for years on his own server to now move it to a server shared by other (perhaps less well-behaved) applications.

Virtualization Architectures

Virtualization requires a virtualization layer between the virtual machines and the physical machine. The virtualization layer is responsible for multiplexing the access requests of the VMs to the resources of the physical machine. There are two primary ways in which virtualization is implemented in currently available products:

- Operating System Virtualization, in which the virtualization layer sits on top of a *host operating system* that is installed on the physical server. The host operating system

provides the interfaces between the virtual environments and the physical processor and its I/O devices.

- Bare-Metal Virtualization, in which the virtualization layer (the hypervisor) sits directly on top of the hardware with no intervening host operating system. The hypervisor in this case provides the common device drivers.

We will explore these architectures in detail in our next part in this series.

Summary

Virtualization has the potential to significantly reduce the size of server farms required by data centers. This turn results in significant reduction in capital and operating costs – less equipment, less space, less power, less everything.

In the following parts of this series, we will describe the various virtualization architectures. We will then explore the fault-tolerant capabilities of these architectures, and we will review various products that provide these capabilities.