



TECHNICAL WHITEPAPER

Taming Virtual Machine Sprawl: How to Get Your Virtualization Under Control

Taming Virtual Machine Sprawl: How to Get Your Virtualization Under Control

Virtualization technology has tremendous potential to help control the spiraling cost of IT, but virtualization projects often don't deliver the expected cost savings. The problem is not the concept of virtualization, but the lack of a management strategy that addresses all the costs that result from increasing the number of applications that share each server. A comprehensive virtualization management strategy that addresses hardware, software, power, real estate, and administration costs is vital to realizing the cost-saving potential of virtualization technology.

Virtualization Potential

Under-utilized servers are an enormous factor in today's IT budget crisis, with system utilization inside large corporate data centers rarely exceeding 10 percent. Virtualization has the potential to slash IT costs by letting organizations accomplish more work with fewer servers.

Virtualization technology is a framework for separating the behavior of a system from its physical implementation. Using virtualization software, a physical machine can run multiple execution environments that appear as private, self-contained systems to users and to applications. The software operating environment of each virtual machine (VM) is isolated from other VMs in the same physical machine. This allows multiple, virtual instances of a server to run on a single, physical system. Theoretically, running multiple VMs on one server will increase the machine's utilization and reduce costs by letting data centers perform the same amount of work with fewer systems.

The VM concept has been used successfully in the mainframe environment since the late 1960s, when IBM introduced the IBM System/360/67. Since then, virtualization on mainframes has become routine and robust. Virtualization on microcomputers appeared much later, and that technology is still evolving.

Virtualization in the current era of server technology first became popular in the 1990s when a special breed of VM was created to host the Java programming language. These Java Virtual Machines (JVMs) let users move applications from one architecture to another without recoding. Millions of JVMs are now installed on PCs and servers across the globe. JVMs have revolutionized software development practices, but have not significantly reduced IT infrastructure costs.

Virtualization on today's servers is accomplished with software known as a Virtual Machine Monitor (VMM) or Hypervisor. This software isolates an application from the underlying hardware so programs intended for one piece of hardware can execute on a different piece of hardware. Using a VMM, a single physical server can simultaneously host multiple operating systems (OSs) such as Windows and Linux, or multiple copies of a single OS. VMMs are available from companies such as VMware and Microsoft, as well as from open source projects such as Xen.

The Real Cost of Virtual Machines: Management

Running multiple VMs in a single server can reduce the number of physical machines required to perform a given amount of work. This leads users to believe that implementing VMs will reduce costs. What many users fail to realize is that each VM presents the same software, maintenance, and administration expenses as a physical server. Research indicates that less than one-fifth of the total cost of a server is related to hardware, while administrative costs make up more than twice the hardware cost. With

between four and eight VMs commonly running on a single physical server, these costs can add up fast. Users often must purchase a copy of the OS and relevant applications and software licenses for each VM. All the OS copies and the applications need maintenance and patching. Administrative costs might remain static or even increase in such a configuration, but are unlikely to go down because each VM must be managed just like a physical system. A typical VMM-based server consolidation model will have several VMMs on each physical server (see Figure 1). If each of 4 physical servers is divided into 4 VMs, an IT administrator must manage 4 physical systems plus 16 virtual systems.

In addition to increasing software and administration costs, improperly managed virtualization implementations often reduce hardware and application performance. Hardware performance in virtualized systems is frequently impacted by uncontrolled resource-sharing within each physical server. As an application running on one VM grows and takes an increasing percentage of the available memory, for example, other applications running on different VMs in the same physical machine can run short of memory. Those applications then perform more slowly, which can drain CPU cycles from the most important applications.




Figure 1 Virtual machines spawned by physical servers

Application performance is also at risk because VMs have a single-system mentality that does not account for the network or the complexities of service delivery for large-scale applications. Large applications such as an ERP system operate across many systems and run many services. These applications often have elements with very low utilization, such as an order entry system that is used heavily for only a few days each year at month- or quarter-end periods. A typical organization's basic VMM strategy today does not define how to link the resources to support that application during peak usage periods and still use those resources for other tasks the rest of the time.

Basic VMM tactics also do not account for the availability risks of running more applications on a single piece of hardware. If that hardware fails, up to eight times as many applications are likely to be at risk. In effect, without a mature virtualization management strategy, the increase in availability risk becomes directly proportional to any increases in utilization seen by the user.

Taming Virtual Machines: Automation and Control

VMware pioneered the field of x86 virtualization, and products with similar capabilities have recently emerged from Microsoft and Xen. Just as today's data centers run a collection of operating systems for different purposes, data center administrators are implementing a diverse collection of VMs from different vendors and hosting services inside virtual machines as well as directly on physical hardware. Within this diverse environment, a practical virtualization strategy must minimize the business risks related to system

failure, reduce software costs, maximize application performance, meet the multi-system needs of large-scale applications, and ease the management burden inherent in VM implementations. All of these goals can be accomplished by balancing resources between physical and virtual machines automatically across the data center.

Balancing resources is a critical component of using virtualization successfully for server consolidation because it addresses the primary cause of hardware under-utilization. Servers are not under-utilized because of a lack of work throughout the organization, but because there is no way to dynamically move applications on overloaded servers to under-utilized machines. The solution is a dynamic, real-time IT infrastructure that allows physical and virtual resources to be quickly repurposed to the most critical tasks and applications without an extensive labor investment that increases operational costs.

A Comprehensive Virtualization Management Solution

Cassatt® Corporation provides software that complements virtualization technology by dynamically allocating physical and virtual resources to applications across the data center on demand, without human intervention (see Figure 2a and Figure 2b). Cassatt can allocate resources between VMMs from VMware, Xen, and Microsoft, as well as physical hardware running Intel, AMD, and SPARC microprocessors. The software also provides for advanced management of Java Virtual Machines in conjunction with application servers such as BEA WebLogic.

Resource allocation is managed by a central optimization engine that constantly monitors all physical and virtual resources. The optimization engine distributes applications to the different servers based on current demand, using individual service level agreements (SLAs) for each application. IT administrators create the SLAs using simple, but fine-grained, policy-based automation tools.

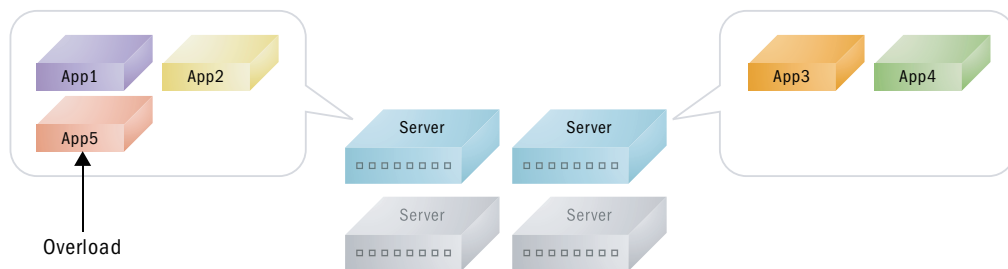


Figure 2a An application becomes overloaded

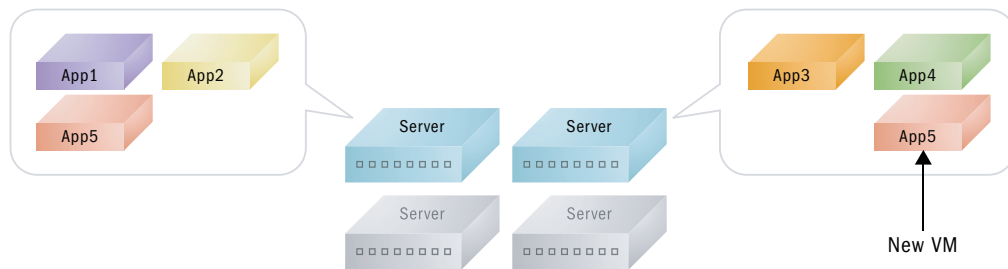


Figure 2b The Cassatt software automatically responds to the overload by creating a new VM

Dynamic application distribution increases server utilization while maintaining quality of service, even for complex, interconnected applications. Maximizing server utilization dramatically reduces the need for expensive server over-provisioning, which has traditionally been the only way to ensure acceptable service levels during peak usage.

The Cassatt software further reduces the need to over-provision for peak usage by automating the management of a global free pool of bare-metal servers. These servers remain powered-off unless a physical server fails or an application requires additional resources and none of the servers that are currently running have significant free compute power (see Figure 3a and Figure 3b). Under these conditions, the software will automatically use a remote management interface to power-up a server from the global free pool.




Figure 3a A physical server fails

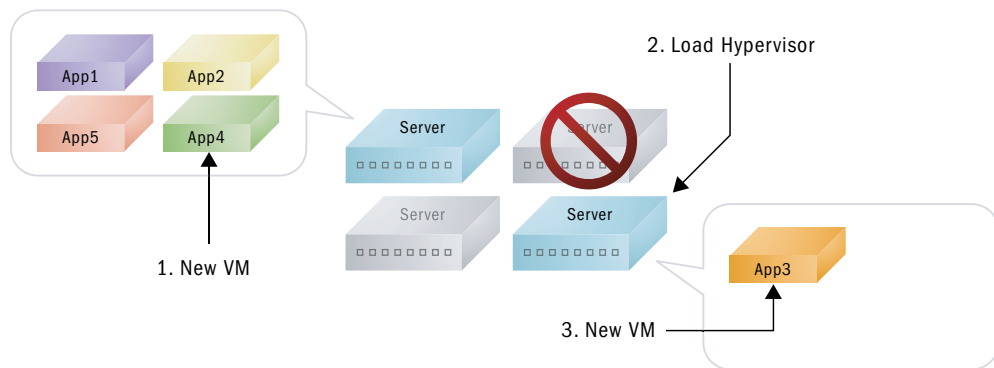


Figure 3b Cassatt automatically responds to the failure by reallocating resources

As the server starts, the Cassatt software causes it to boot using a software image that contains an appropriate operating system, a virtual machine manager, or both. Then the Cassatt software automatically deploys copies of any overloaded applications to the server so that all applications can meet their service level goals. When processing power becomes available on other systems, the Cassatt software will power down the extra server and return it to the global free pool for use by other applications. Continuous, automatic service level monitoring and resource re-allocation lets IT administrators safely eliminate over-provisioning and realize dramatic cost savings. The global free pool of servers also makes it easy to bring up new applications without concern for server provisioning. As applications are added to the data center, IT administrators can easily monitor free server use. If the Cassatt software activates more servers than are returned to the free pool, IT administrators can simply add more servers to the pool to ensure sufficient capacity.

All Cassatt software uses off-the-shelf hardware, operating systems, network switches and applications. Standard components let IT organizations introduce this solution into their infrastructure gradually, without radical change.

A Resource Management Revolution

Virtualization is an important first step to reducing data center costs, but VMs alone only address the hardware aspect of the resource utilization equation. To reduce overall data center costs, a virtualization solution must be coupled with a strategy to automate virtualization control, and management that addresses software and administration issues as well as hardware utilization.

The Cassatt software provides a policy-based, automated virtual machine management solution that dynamically allocates virtual and physical resources throughout the data center according to priorities that are aligned with business requirements. This solution changes a static data center into a dynamic, real-time infrastructure with a single control point for all physical and virtual IT resources. With Cassatt software, administrators can concentrate on delivering the application services that drive business success instead of focusing on constant server management.

The Cassatt Collage platform, in conjunction with the Cassatt Collage Cross Virtualization Manager (XVM), is the catalyst that organizations need to realize the potential of VM to increase server utilization without incurring hidden costs. Combining Cassatt software with virtualization technology will help organizations provide the cost-effective data center services that are vital to business success.

Virtualization Glossary

Hypervisor:	Software that allows multiple operating systems to be hosted on a single computer.
Java Virtual Machine:	A virtual machine that executes Java bytecode. This code is most often generated by Java language compilers, although the JVM has also been targeted by compilers of other languages.
Microsoft Virtual Server:	A virtualization product from Microsoft.
Virtualization:	An abstraction or framework for separating the behavior of a system from its physical implementation.
Virtual Machine:	A software program that emulates a hardware system.
Virtual Machine Monitor:	See Hypervisor.
VMware:	A division of EMC that produces virtualization software including products such as VMware ESX and VMware Server.
Xen:	An open source Virtual Machine monitor that originated at Cambridge University. RedHat and SuSE plan to embed Xen into their popular Linux distributions.

For More Information

- Cassatt: <http://www.cassatt.com/>
- VMware Server: <http://www.vmware.com/products/server/>
- Microsoft Virtual Server: <http://www.microsoft.com/windowsserversystem/virtualsever/>
- Xen: <http://www.cl.cam.ac.uk/Research/SRG/netos/xen/>
- Wikipedia Virtualization Article: <http://en.wikipedia.org/wiki/Virtualization>



Cassatt Corporation t: 408.437.7600
1740 Technology Drive f: 408.437.1364
San Jose, CA 95110 www.cassatt.com