Beyond Server

Consolidation

WERNER VOGELS, AMAZON.COM

Server consolidation helps companies improve resource utilization, but virtualization can help in other ways, too.

Virtualization technology was developed in the late 1960s to make more efficient use of hardware. Hardware was expensive, and there was not that much available. Processing was largely outsourced to the few places that did have computers. On a single IBM System/360, one could run in parallel several environments that maintained full isolation and gave each of its customers the illusion of owning the hardware.¹ Virtualization was time sharing implemented at a coarse-grained level, and isolation was the key achievement of the technology. It also provided the ability to manage resources efficiently, as they would be assigned to virtual machines such that deadlines could be met and a certain quality of service could be achieved.

At first glance it appears that not much has changed. Today the main application of virtualization technology in the enterprise is to combat server sprawl through virtualization-based consolidation. Isolation, security, and efficiency remain the main benefits of using virtual machines in this context.

Even though this article is mainly about improving resource utilization, if we consider virtualization only as a tool for

server consolidation, we are underestimating its true potential. Virtualization breaks the 1:1 relationship between applications and the operating system and between the operating system and the hardware. The removal of this constraint not only benefits us in creating N:1 relationships where we run multiple isolated applications on a single shared resource, but also enables 1:N relationships where applications can span multiple physical resources more easily by providing elasticity in their resource usage.

CONSOLIDATION

Classic consolidation is focused on multiplexing physical resources over a number of virtualized environments. The immediate benefits are obvious: reduce the amount of hardware, reduce the data-center footprint, and indirectly reduce power consumption. The latter is an increasingly important driver for consolidation since energy companies are starting to provide significant incentives for cutting back consumption. Consolidation is in essence a cost-reduction activity; by significantly reducing the server footprint (by 30 to 50 percent or even more), the capital investment requirements are directly affected,

Beyond Server Consolidation

which leads to reduced staffing needs and lower operational costs.

One of the main causes for server sprawl in the enterprise has been the requirement by vendors to run their applications in isolation. This requires the IT department to dedicate one or more servers to an application, even if the servers provide more resources than the application requires. Also at the infrastructure level we see that the modern enterprise has many dedicated servers: DNS, DHCP, SMTP, printing, Active Directory/LDAP, etc. Another driver of sprawl is operating system heterogeneity: a mail server that requires Windows Server, a database that is best run on Solaris, a network management package originally acquired for use with AIX, etc.

Add to this the effects of mergers and acquisitions and other integration projects and you will find that an enterprise with a large collection of servers, each dedicated to a single task, is a common pattern. Mergers and acquisitions in particular bring new applications or application versions, additional servers, and, often, new complex integration middleware. It is not uncommon that after a merger the number of servers to support the new infrastructure is larger than the combined server count of the separate companies. Given the complexity of these integration projects, the IT organization relies heavily on coarse-grained server-driven isolation to achieve the integration.

The large number of underutilized servers has become a major problem in many IT departments. Individual companies provide no official numbers about server utilization, but many of the large analyst firms estimate that resource utilization of 15 to 20 percent is common. From personal experience in talking to other CTOs and CIOs around the world, I believe that those numbers are on the high side and the true utilization is often in the 5 to 12 percent range. With more powerful servers enter-

ing the data center every day, the utilization number is decreasing rather than going up.

Single averages seldom tell the whole story, however. Utilization of servers is highly dependent on the type of workloads and is often subject to periodicity. If you inspect utilization over longer periods, you will find that it is more accurately represented by a range that differs depending on the application. In their article on energy-proportional computing, Luiz André Barroso and Urs Hölzle show that in a highly tuned environment such as Google's the utilization tends to fluctuate between 10 and 50 percent when inspected over longer timeframes.² Figure 1 shows the average CPU utilization of more than 5,000 servers at Google during a six-month period. This data reflects our experiences at Amazon; some utilization is driven by customer behavior, but some is triggered by fulfillment process patterns or digital asset conversions.

Cost reduction is an important goal in many IT departments, and server consolidation certainly tops the focus list. Virtualization has become the primary tool in driving server consolidation: 81 percent of CIOs were using virtualization technologies to drive consolidation, according to a recent survey by CIO Research.³ Even though the strategy appears mature, consolidation architects still face significant technical hurdles.

The first challenge in the consolidation process is

how accurately to characterize an application's resource requirements. An engineer's educated guess would, in general, result in an incomplete view of the real constraints. Building a resource-usage profile of an application is essential, and such a profile does not only focus on which resource eventually will bound the application, but also analyzes resource usage over time to determine periodicity and whether dependencies on other system or application components exist. The second part of the profile is how the application behaves when it runs out of capacity: how sensitive is the application to resource shortage, can it adapt, or does the environment need to maintain strict bounds?

A common step before this analysis is to break up applications that run in shared environments and put each one in an isolated environment in the hope that it will become easier to predict the individual applications' resource usage in response to request patterns. They will then be separately managed according to their own resource profiles.

The next challenge arrives once there is a clear picture of resource usage and scaling: how optimally to distribute the virtual machines hosting the applications over the physical resources. This is an area with many emerging tools to assist system architects in finding the right mix, but reports from the field indicate that this is still largely a process of trial and error before a reasonable balance is achieved. This is not a trivial task. As we can see from figure 1, resource usage may change significantly over time, which makes the relevant load testing very hard.

The biggest challenge of the whole consolidation process, however, is without a doubt the balancing of server workloads at runtime; 64 percent of the CIOs mention this as problematic in the CIO Research survey. Because of the reduced slack in the system, the applications are more exposed to resource shortages, especially in situations where workloads are highly dynamic.

A good example of a business with changing resource demands is Powerset. Initially building indexes and updating indexes over time have very different resource demands. Powerset has released a data-center resource analysis tool that helps predict which business-specific scenarios make sense for buying, leasing, or using virtualized resources. Given the changing resource demands, in most cases the virtualized servers are more cost effective.⁴

IS 100 PERCENT UTILIZATION THE GOAL?

There are many reasons why we will never see 100 percent utilization: workloads in the enterprise are heterogeneous, and demand may be uncertain and often

occurs with spikes. As such, some CPU cycles or IOPS (I/O operations per second) will always be unused when you measure utilization at larger time scales. Even at the individual operating system level, however, we know that perfect utilization is not possible. For example, an operating system such as Linux may start to behave unpredictably under combined high CPU/I/O loads. We joke that some of these operating systems exhibit an "Einstein Effect"—at high utilization, space and time are no longer guaranteed to behave the same.

As a consequence, the measure of success of consolidation is set more realistically: for pure CPU-bound environments, 70 percent seems to be achievable for highly tuned applications; for environments with mixed workloads, 40 percent is a major success, and 50 percent has become the Holy Grail.

For applications and servers that do become overloaded, migration is potentially a solution. Transparent migration, however, is hard to achieve, and many legacy applications do not respond favorably to this. Two more coarse-grained techniques seem to be effective: application checkpoint and restart has been built into several applications as a disaster recovery tool and is used to move applications to different physical servers; and a number of applications can be run in *clustered* mode (e.g., MS Cluster Service enabled), where a second node can be brought up and at the application level, state and work can be migrated from the first to the second node.

An extreme example of the use of virtual machine migration is application parking. In this case several applications that are hardly using any resources are each running in their own VMs but are sharing one physical server when they are in rest state. As soon as an application starts using more resources, it is migrated to a server that has sufficient resources available that fit the application's profile.

BEYOND COST SAVING

Until now we have discussed traditional consolidation, as exercised by many IT departments, where the main focus is thoroughly analyzing enterprise-wide resource usage and using virtualization to multiplex those resources as efficiently as possible. Business priorities determine at any given time how efficiency is measured. In the classic environments we see a grow-and-shrink trend; an application is brought in on its own server or added as part of a merger integration. This is followed by a phase of resource usage and risk analysis, which determines where the applications can be collocated in a virtualized manner, after which the server pool shrinks again.

Beyond Server Consolidation

In all of this, however, virtualization is used as a traditional IT cost-saving tool. The real power of virtualization as a strategic enabling technology comes when you consider its role in application deployment and management. With the right virtualization management tools you can get to an environment in which you can significantly speed up the time to market of new applications and have them scale efficiently to customer demand.

A good example of this is the role of virtualization in Amazon's infrastructure. Amazon is the world's largest service-oriented software organization, where not only the technology is service oriented, but also people are organized in teams that mirror the software organization. This gives Amazon great agility in customer-focused business and technology development. In running close to 1,000 services, Amazon ended up with many engineers performing similar tasks, most of them related to resource management: managing application deployments, configuring servers, handling storage failures, configuring load balancers, etc. Conservative estimates indicated that engineers were spending up to 70 percent of their time on general tasks not directly related to the business functionality of their service.

We decided to bring these common activities into an infrastructure-services platform where they could be managed more effectively while maintaining Amazon's focus on reliability and performance. Storage, compute, and messaging were virtualized as infrastructure services. A number of these services have since been made available outside of Amazon: S3 (Amazon Simple Storage Service), EC2 (Elastic Compute Cloud), SQS (Simple Queue Service), and SimpleDB.⁵

Two key requirements in the design of these infrastructure services markedly changed the way resources are managed: the services are fully self-service, allowing engineers to start using them with minimal friction; and resources can be managed dynamically, giving engineers the power to acquire and release resources immediately.

UNIFORM APPLICATION DEPLOYMENT

Amazon EC2, the service most similar to traditional virtualization, uses a model where engineers programmatically can start and stop instances that they have previously built.⁶ These instances are virtual machine images that are the output of the application build process, and they are stored in the Amazon S3 storage service. The EC2 management environment places the virtual machine on

a physical server based on resource requirements. This provides engineers with the ability to grow and shrink the resources their services use based on customer demand and other scaling attributes.

This brings us to one of the main strategic advantages of virtualization: it creates a uniform application deployment environment where engineers are shielded from the particulars of the underlying hardware. It is not uncommon to see a single virtual machine running on a physical server, where the goal is not to maximize efficient resource sharing, but to speed up deployment of applications and to scale up and down at a moment's notice.

Feedback from Amazon EC2 customers revealed that they were traditionally confronted with significant overhead in acquiring resources from their IT organizations. Server acquisition times often run into several months, and once a resource has been allocated to an application, teams are unwilling to release it given the long lead times in reacquiring the resource when needed again.

This conservative approach requires long resource planning cycles: teams need to predict their resource usage long ahead of deployment and execution, which triggers overscaling to deal with unexpected higher demands on the application. This model is a stumbling block for enterprises that want to react to demand faster and more efficiently. There is increasing uncertainty in many markets as product and service life cycles are compressed and increased competition makes the success of products more difficult to predict. To adapt to these new realities, enterprises need to shift to different models for their resource management, where acquiring and releasing resources based on demand is becoming an essential strategic tool. In this context the pay-as-you-go model of the Amazon infrastructure services is very attractive.

Having the virtual machine as the standardized unit of deployment is crucial in adapting to shifting resource demands, where it is important not only to acquire resources but also to release them when they are no longer needed. Many of Amazon's EC2 enterprise customers claim that their resource acquisition cycles have changed from months to minutes.

One area characterized by very long cycles in acquiring resources is IT in government. Funding and allocation decisions often require teams to purchase servers at the beginning of a project, many months before the software is completed and before a good usage pattern has been developed. This leads to ultra-conservative planning with

low utilization of the ultimate configuration and results in significant barriers to prototyping and experimentation. One DoD IT architect reported that the department's software prototype normally would cost \$30,000 in server resources, but by building it in virtual machines for Amazon EC2, in the end it consumed only \$5 in resources.⁷

The new agility also caters to other advantages of using virtualized infrastructures. While traditional consolidation based on virtualization only increases the density of resource usage, there still may be barriers to changing the mix of applications and services running at any given time. Incorporating virtual machines in the change management process and adding autonomic management features significantly improves the agility of the enterprise. Using economic models to automate resource allocation to optimize business value remains a Holy Grail.

VIRTUALIZATION OF THE DATA CENTER

Virtualization plays a crucial role in enabling the IT organization to grow beyond its data centers and exploit utility computing infrastructures. Utility computing is the packaging of resources such as computation and storage as metered services similar to public utilities (electricity, water, natural gas, and telephone networks). This has the advantage of low or no initial cost to acquire hardware; instead, computational resources are essentially rented.

In an organization where virtualization is already pervasive to support consolidation and/or application deployment scenarios, the tight dependency between application/operating system and the physical hardware has already been removed. Running the virtual machines on hardware that is not directly controlled by the organization is a logical next step.

Utility computing services are different from traditional application outsourcing where the infrastructure owner runs the application on behalf of the client and has application-specific knowledge. These services are also different from grid environments as they do not impose a particular programming model to be used for application development. Instead a utility computing service allows its customers to launch virtual machine instances on their hardware in a manner similar to running these VMs in their private data centers. Amazon EC2 is one of the prominent services that offer access to compute resources in a utility style using virtual machines; EC2 customers can package virtual machines as they run them in their data centers to run in Amazon EC2 as well.

Using utility computing services benefits the cost-saving targets that often underlie consolidation efforts; capital expenditures are greatly reduced by going to a model

where you pay for the resources only for the period of time that you actually use them. Frequently, enterprises start using these utility computing services to address their needs for overflow and peak capacity; this way they can deal with uncertainty in demand without big investments in hardware that will be idle most of the time. This on-demand acquiring and releasing of resources is additive; once enterprises have become comfortable using a computing utility service for handling peaks, they quickly start using it for other tasks, especially those that do not require around-the-clock resource allocation such as document indexing, daily price calculations, digital asset conversion, etc.

A good example of using utility computing for excess capacity tasks is the *New York Times's* project to convert 11 million historical articles from TIFF to PDF. Finding sufficient capacity on the corporate server would have been difficult, given the deadlines for the project, and buying additional hardware for such a one-off task would not be very efficient. The *Times* created a virtual machine image containing a special conversion application, moved 4 TB of images into Amazon S3, and fired up 100 instances of the virtual machine in Amazon EC2. Within 24 hours all articles were converted into 1.5 TB of PDF at the cost of a fraction of a single server.⁸

One of the benefits of this model is that measuring TCO (total cost of ownership) becomes easier; instead of amortizing the costs of server, network, power, and cooling over a number of applications running on a server, the absolute infrastructure costs are metered utility cost.

Looking at the wide variety in companies that use virtualization to run their applications in Amazon EC2, one can see that utility computing has many applications beyond enterprise capacity management. Usage ranges from classical parallel computing by financial and pharmaceutical companies to startups running Web services, from large software companies using it for product and release testing to image rendering by movie studios. All this is enabled by virtual machine technology for packaging and instantiating the applications, managing security, and on-demand access to required resources.

TESTING, TESTING, ...

Software testing is another area that is always on the short end of receiving resources and has much to gain from virtualization. The demands of testing on the infrastructure change during the development cycle. Early in the cycle one may use a continuous integrating technique with nightly rebuilds of the environments, changing to load and scale testing later in the cycle. Test

Beyond Server Consolidation

engineers often need to keep many different servers running, each with a different version of an operating system for managing release testing.

Traditionally, QA departments manage their own resources, and in many cases they are highly constrained in the resources available to them. Even in this constrained environment, there would be periods during a day, week, or year where the hardware would go unused. Virtualization has changed the QA process dramatically by acquiring resources on demand when they are needed for particular tests and releasing them when the tests are finished. This has tremendously improved the utilization the testers get out of their environments. No longer is there a need to have many different operating systems running or to have complex multiboot environments around; starting and stopping different operating system images becomes an on-demand activity. Going virtual has in many cases increased the number of resources available for QA at any given time, as the pool of physical resources can be shared with the production environment. This makes load testing at scale more realistic.

SCALE, RELIABILITY, AND SECURITY

While this article focuses on the role of virtualization in utilization management, there are other areas where virtual machines can play an important role. One of those is security, where many innovative uses are possible but where even the simplest brings many benefits. Moving an application from a shared environment into its own dedicated virtual machine allows for straightforward operator and user access control. It can reduce the number of open ports and as such the potential for exposure to vulnerabilities. Many IT groups use this technique to meet compliance requirements for applications that do not have adequate access control and auditing.

Similarly, the use of VMs for uniform application deployment can be the basis for disaster management. Often a simple checkpoint-restart facility is sufficient to do fast failover between machines. If applications are built for incremental scalability, the adaptive management facilities such as those in utility computing infrastructures will allow organizations to quickly grow and shrink capacity based on demand.

SUMMARY

Virtualization's main application in the enterprise is still server consolidation. As effective as that is, we are likely

to see a very different picture a number of years from now, where virtualization will be the key enabling technology for a series of strategic changes in IT.

Adaptive resource management using utility computing will be essential to success in an economy with increasing uncertainty. Adapting quickly to new customer demands, new business relationships, and cancelled contracts will be a key business enabler in the modern enterprise, regardless of whether the enterprise executes a software-as-a-service strategy or uses the resource in a more traditional manner.

Virtualization will change the way we do testing, with QA departments getting access to a greater variety of resources than they ever had before—at a much lower cost to the business. Similarly, companies that were not proficient in handling reliability, fault tolerance, and business continuity will find in virtualization a new tool that will allow them to make significant progress toward these goals without rewriting all of their software. □

REFERENCES

1. Gum, P. H. 1983. System/370 extended architecture: Facilities for virtual machines. *IBM Journal of Research and Development* 27(6): 530-544.
2. Barroso, L. A., Hölzle, U. 2007. The case for energy-proportional computing. *IEEE Computer* 40(12).
3. CIO Research. 2008. Virtualization in the enterprise survey: Your virtualized state in 2008. *CIO Magazine* (January).
4. Powerset Datacenter Dashboard; http://www.powerset.com/flash/datacenter_model.
5. Amazon Web Services; <http://aws.amazon.com>.
6. Amazon Elastic Compute Cloud (EC2); <http://aws.amazon.com/ec2>.
7. Brewin, B. 2006. Amazon.mil? DISA is intrigued by Web services model for creating systems. *Federal Computer Week* (October 30).
8. Gottfrid, D. 2007. Self-service, prorated super computing fun! *New York Times* (November 1).

LOVE IT, HATE IT? LET US KNOW

feedback@acmqueue.com or www.acmqueue.com/forums

WERNER VOGELS is vice president and chief technology officer at Amazon.com, where he is responsible for driving the company's technology vision.

© 2008 ACM 1542-7730/08/0100 \$5.00