A Revolutionary Approach to Cloud Building

As enterprises consider cloud computing, many are looking to private clouds to address some of the risks that come with public clouds.

With the emergence of large-scale public clouds from Web behemoths such as Amazon and Google, the notion of a ready-for-the-taking, pay-as-you-go resource model has caught the attention of many IT users, corporate and otherwise.

But large enterprises in particular are more likely to be watching the trend rather than jumping feet first into cloud computing. They are concerned about the privacy, security and management implications of putting sensitive corporate data onto a shared, public infrastructure. A more logical starting point, many enterprises are beginning to figure, is to build private clouds, either on- or off-premises, that would provide many of the benefits but address much of the risk associated with public clouds.

Over the next three to five years, industry watchers expect private clouds to become widely used, if not predominant, within the enterprise portfolio of compute models. This portfolio also includes public cloud as well as traditional computing and virtualized resources.

Most companies will move toward their private clouds in a slow, steady course, methodically adapting and evolving their infrastructure and virtualization efforts toward a cloud environment capable of being shared among any number of corporate applications and users.

Dell, however, believes some instances call for a complete clean-sheet approach.

> "In time, the balance will shift between evolutionary and revolutionary"
>
> — Barton George, Dell

Evolutionary vs. Revolutionary Path to the Private Cloud

For many enterprises, building a private cloud is simply the next step on an evolutionary path that began with data center consolidation. When a company has established a strong virtualization underpinning and is working with traditional applications, an evolutionary approach to the private cloud makes perfect sense, says Barton George, Dell's cloud computing evangelist.

With an evolutionary approach, a company gradually builds its private cloud by layering capabilities such as automated provisioning, usage-based billing, resource pooling and a self-service portal on top of virtualized infrastructure as it adapts processes and policies to the new model. With every step along the way, the enterprise becomes more agile and achieves greater efficiencies, thus deriving value even before realizing the full potential of a private cloud, George says.

While this evolutionary approach is appropriate in many cases, for some enterprises in some instances, taking what George calls a "revolutionary" approach to private clouds will be more efficient and much more appropriate.

This revolutionary approach makes use of "new world" applications that are both written and deployed in the cloud. These cloud-native applications are designed from the ground up for greater scalability and use across a multitude of servers. As a result, they run more efficiently, are more responsive and deliver a better end-user experience.

Considering the Revolutionary Approach

George suggests that enterprise IT leaders should think about taking a revolutionary approach to cloud computing when faced with the following:

- Greenfield opportunities involving no legacy infrastructure
- The need to support scalable cloud-native Web 2.0 applications (i.e., applications written in and for the cloud)
- A requirement to build out at such a scale that resiliency and availability are architected into the software rather than handled by physical infrastructure.

"The idea of scale is really important in the revolutionary concept," says Kevin Van Mondfrans, Senior Manager, Dell Cloud Solutions. "What you're doing is taking an application and scaling it out across multiple cloud compute nodes, and utilizing virtual load-balancing following best practices. For example, a cloud built with scale in mind will distribute the application or database across multiple physical servers, not just different virtualized instances on one server."

Also, if Web page load times are not satisfactory, the company could rapidly add more instances of the Web site rather than build bigger Web servers, Van Mondfrans says.

"Rather than having three really large Web servers, the company would have 10 smaller Web servers spread out across the infrastructure. This adds to availability and overall performance, and provides the ability to address unexpected peaks more easily," he explains.

This revolutionary approach requires a new way of thinking about the cloud, but one that Van Mondfrans says enterprise IT executives should undertake sooner rather than later. "This is where the application paradigm is going," he says.

The Public Spark to the Private Cloud Revolution

To guide enterprises on a revolutionary path to the private cloud, Dell draws from its experience helping a select group of the world's largest Internet companies build public cloud infrastructures. Dell has worked directly with Microsoft Azure and other companies that operate at the leading edge of efficiency and view their data centers as their factories. For these organizations, Dell has created customized systems built to operate in a scaled-out environment, where the availability and resiliency are baked into the application, rather than the hardware.

Revolution comes to Data Analytics

An enterprise can apply this revolutionary approach to data analytics as well.

Case in point: the data warehouse environment. Enterprises caught in old-school thinking will typically piece together the data warehouse

Partners in the Cloud Revolution

To help enterprises understand and begin moving down the revolutionary path to cloud computing, Dell has initiated a number of key strategic partnerships through its cloud ISV partner program and its development of specialized private cloud solutions.

For example, Dell has teamed with public cloud platform provider Joyent to create the turnkey Dell Cloud Solution for Web Applications, the first revolutionary cloud solution Dell sells and supports. With this offering, Dell combines Joyent's Web application cloud software with its cloud-optimized servers, reference architecture and custom services, to provide a platform-as-a-service solution.

"We'll be offering quite a range on the services side—services an enterprise would need to assess, design and deploy a private cloud, plus the operational help once they have the cloud deployed," explains Kevin Van Mondfrans, Senior Manager, Dell Cloud Solutions. Dell has also recently announced that it is teaming with Microsoft to build a revolutionary PaaS appliance based on the Windows Azure platform.

In addition, Dell is working with three cloud ISVs to offer easy-to-buy and -deploy cloud solutions accompanied by reference architectures optimized for and validated on its hardware, according to Barton George, cloud computing evangelist at Dell. Initially, Dell will work with its cloud ISV partners to help develop and deploy solutions, but the partners will handle ongoing support.

The ISV partners include Aster Data, for Web 2.0 data analytics, and, in roughly the same realm, Greenplum, for self-service data warehousing within traditional large enterprises.

The third ISV partner is Canonical, which, as the commercial sponsor behind Ubuntu Linux, has teamed with open source private cloud provider Eucalyptus Systems to create the Ubuntu Enterprise Cloud (UEC). In this partnership, Dell will provide infrastructure support services and create a reference architecture for the open source infrastructure-as-a-service private cloud, while Canonical will handle the application support.

"If a customer buys the UEC platform, our reference architecture will say, 'Use these servers, connected this way, and put this software on these boxes,' for example," George explains. "We'll tell them how to deploy the cloud, and Canonical will provide software-level support."

All four companies have embraced the revolutionary cloud approach, Van Mondfrans says.

"They're building on the concept that company's flexibilities are going to build scaled-out environments to get better economics and densities," he says. "They understand that these solutions will run better than they would in traditional environments."

architecture using a couple of large front-end servers and a monolithic SAN in the background. But following the new school of thought, an enterprise IT executive would say, "I can achieve a very large data warehouse with commodity x86 servers. What I will do is continue to add servers as needed with internal storage. I will replicate the data across multiple servers from an availability standpoint, and I can get basically infinite capacity scaling. And, if I scale further along the capacity path by adding more servers, I'll also get better performance because I've got more servers r unning the workloads for me."

But not just any old commodity servers, even high-end ones, will do. A revolutionary cloud requires specialized servers optimized for specific workloads—heavy analytics in the case of the data warehouse, for example. "You have to approach system design from a different perspective, providing less of what you don't need and more of what you do need in this type of environment," George explains.

Systems built for Scale

Inspired by the heavily customized servers created for its large, hyper-scale customers, Dell offers a next-generation server series that is particularly suited to revolutionary cloud environments: the PowerEdge C series, powered by Intel® Xeon® quad- and six-core processors, are extremely dense, have ultra-fast memory subsystems and expansive disk capacity. "The CPU, memory and disk meet the workloads at optimal cloud performance levels," Van Mondfrans says.

"We're using the same enterprise-quality components," George says, "but we've taken out things such as management capabilities that would be redundant when you're operating in a highly intelligent software environment. And we've doubled-down on things that matter particularly to the applications or when operating at scale. Power efficiency, TCO and density become much more important when you're talking about this volume of servers."

Enterprises can operate a production level-quality private cloud using as few as 15 of these specialized servers, but most companies that embrace the revolutionary model will be looking at much larger server farms.

"We expect these private cloud customers to begin to organize groups of servers into modules across which applications will scale," Van Mondfrans says. "We expect customers to deploy compute in multi-server or full-rack modules. Customers no longer want to think about the deployment or management of individual servers. They want to think in bigger buckets to simplify operations. In addition, since the application scaling reduces the reliance on an individual component, customers will think about the availability of the module, not the individual server."

Dell's intention, as evidenced by key strategic partnerships with cloud pioneer Joyent and cloud ISVs Aster Data, Canonical and Greenplum, is to deliver its cloud hardware in simple-to-deploy modules.

> "Customers no longer want to think about the deployment or management of individual servers. They want to think in bigger buckets to simplify operations. In addition, since the application scaling reduces the reliance on an individual component, customers will think about the availability of the module, not the individual server"
>
> — Kevin Van Mondfrans, Dell

Join the Cloud Revolution

"The revolutionary approach to the cloud is built on the concept that you're going to have a large scaled-out environment," Van Mondfrans says.

It's not a matter of "if," but "when," adds George, noting that over time, scaled-out applications will become increasingly prevalent in the enterprise. Embracing a revolutionary approach to the private cloud, as appropriate, today will help enterprises prepare for that future. ■

www.DELL.com

ADDITIONAL READING

Intel Builds a Private Cloud

By Rick Swanborg, *CIO*

Intel's research and development group was the first internal segment to increase efficiency with an infrastructure- and platform-as-a-service approach, improving its server utilization rate from 59 percent in the first half of 2006 to 80 percent today. By reducing infrastructure spending and avoiding building new data center facilities, the company expects to save nearly $200 million. Following this success, Intel decided in early 2009 to build an enterprise private cloud for its office and enterprise groups.

The Situation: From 2005 to 2006, the demand for new business lines dramatically increased the number of design teams in motion-all of which operated on a non-shared and underutilized infrastructure. This, paired with the fact that RaannddD represents the vast majority of Intel's infrastructure, made it the logical place to start a transformation to IaaS and PaaS.

What They Did: Intel implemented a grid computing solution inside its RaannddD environment that features both IaaS and PaaS attributes. The solution is accessible via a self-service portal and command-line interfaces. The portal allows users to provision key infrastructure services-such as storage-in hours instead of months. Perhaps most important, this approach allows Intel to measure what users are consuming.

The project was led by data center architect Brad Ellison and IT cloud engineering lead Das Kamhout. "Overall, this represents a dramatic change in mind-set," says Kamhout. "On-demand self-service allows IT to get out of the way of the business so we can up level IT and be a strategic business partner. And with transparent costs, RaannddD departments across Intel now know what they're costing and what they're consuming."

Why It Was Unique: The combination of the on-demand self-service portal across Intel's entire infrastructure, together with the ability to measure consumption, ensures there is minimal misuse of capacity, thereby allowing IT to efficiently support rapid growth while maintaining high availability of services.

The Takeaway: Having proved that an IaaS approach could work for RaanndDD, Intel IT built the case for replicating those efforts across the organization. The office and enterprise groups' efforts are now under way, with a focus on pervasive virtualization and on-demand self-service for an enterprise private cloud. ■

www.networkworld.com/news
June 14, 2010

ADDITIONAL READING

Cloud Computing: The Future of IT Application Architectures

Assumptions we have traditionally used to design application architectures are increasingly outmoded, says CIO.com's Bernard Golden. Application architectures will change — just as much as IT operations — during the next five years due to the nature of cloud computing applications.

By Bernard Golden

Last week I wrote about the impact cloud computing will have on IT operations. I noted that the increasing scale of data dramatically changes the expectations of how data centers are operated. This week I want to turn to how cloud computing affects IT application architectures, specifically examining the flip side of the coin of data growth: application load. Succinctly put, the assumptions we have traditionally used to design application architectures are increasingly outmoded due to the changing nature of applications. Application architectures are going to change — just as much as IT operations — over the next five years due to the nature of cloud computing applications.

What are the reasons that applications are going to change so much?

All of that big data is going to mean software applications are going to need to change to manage it.

As I noted last week, IDC projections indicate that the average company will experience a seven-fold increase in unstructured data (think click stream capture and video storage, etc., etc.), accompanied by a doubling of

structured data (think database row-and-colum info). I actually think that IDC's projections are understated on the structured data side, because of the constrained assumptions it (very reasonably) brought to its analysis. The remarkable decrease in the cost of IT brought about by cloud computing will — no surprise to economics majors everywhere — lead to much larger amounts of computing being done, which, in its turn, will lead to larger application architectures and topologies.

The Business Use of IT is Changing

In the past, IT was used to automate repeatable business processes — taking something that already exists and computerizing it. The archetype for this kind of transformation is ERP — the automation of ordering, billing, and inventory tracking. That "paving the cow paths" approach to computing is changing. Today, businesses are delivering new services infused and made possible by IT — in other words, creating new offerings that could not exist without IT capabilities. A dramatic example of this is the way music services have developed. Pandora leverages the knowledge of experts to deliver customized song streams to its customers; Pandora tracks the preferences and feedback of every one of its listeners to ensure each receives a personalized offering. Pandora's service could not exist without the support of massive amounts of computing power, which forms the core of the business. Less dramatic, but no less reliant on IT infusion, is the personalized service offered by high-end hotel chains. The personal attention that employees offer guests — going *way* beyond the "prefers non-smoking room" of yore to, say, "likes to see avant-garde theater and new museum exhibitions" — enables highly specific employee interaction with customers. And, guess what, it's all driven by new applications.

The Nature of Applications is Changing

Heretofore, most computing has been driven by human action — someone making a purchase, requesting a Web page, and so on. In the future, a growing percentage of computing will be driven by non-human activities from devices like sensors. As an example, much has been made of the move to smart electric meters — instead of your meter being read by a human walking through your neighborhood, the meter itself will connect back to the electric company data center and upload its billing information. However, one of the other ballyhooed characteristics of these smart meters is their ability to give real-time readouts of load to users. This data about electric usage — the metadata, if you will — will be invaluable to electric companies to help understand how usage changes with immediate pricing feedback. This will result in far more data than just a monthly reading being sent to their data centers. And — wait for it — that data will be transmitted in irregular patterns, leading to highly variable loads, thus affecting the nature of application architectures.

Given how the number, type, and nature of applications are changing, what does this imply for the future of applications and, to address the specific topic of this post, the future architecture of applications? The implications are fourfold:

Application load variability will increase: The driver for the vast changes in resource load variability is application load variability. For hotels, the traditional busy times are early morning (checkout) and late afternoon/early evening (checkin). In the future, personalized attention will mean high application load at other times. In essence, application load will vary throughout the entire day — all 24 hours of it — rather than being focused during business hours. Applications will need to be much more able to dynamically scale.

Application interfaces will change: Instead of being human- (and thereby screen-) focused, data will pour into applications from other applications, sensors, file uploads, and, undoubtedly, things we haven't even thought of yet. So service interfaces and upload interfaces will join terminal interfaces. Applications will need to be able to gracefully — and dynamically — add new data streams as inputs.

Application characteristics will change: The increasing importance of geolocation in apps will necessitate the rapid ability to shift context and data sets. If I'm driving in a taxi, the "nearby" services change quickly as the car moves down the road. Being able to shunt data in and out of working sets quickly (not to mention being able to blend contexts as applications support multiple people sharing a "nearby" context) will become vital. Naturally, this requires high performance.

Application topologies will become more complex: As scale and variability increase, architecture designs must change. I hinted at this last week, when I mentioned the use of memcached as a data caching mechanism used to increase throughput. Complex applications often incorporate asynchronous processing for compute-intensive tasks; message queues are often used as part of this approach. Therefore, application

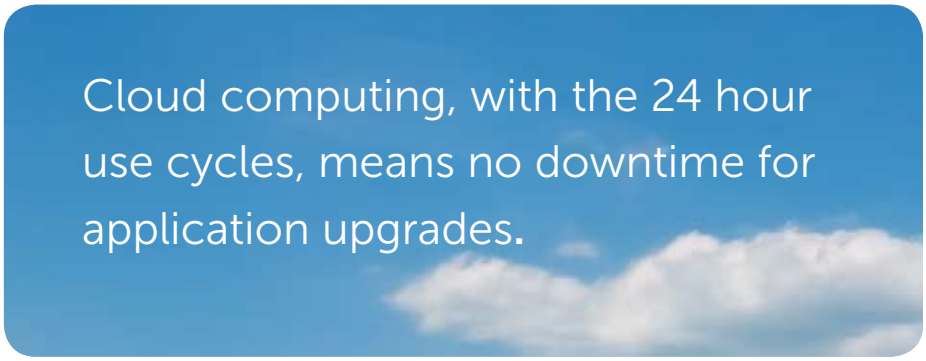

architectures need to change to incorporate new software components and application design.

What are some practical steps you can take to ensure your cloud-targeted application can support these new application requirements? Here are some suggestions:

1. Review software components that you plan to use in the application. Many software components were designed to be used in a static environment with manual configuration and occasional updating. A common design pattern for these components is the use of a "conf" file which is edited by hand to configure the component context. Once the conf file is complete, the component is started (or restarted), reads the configuration information into memory, and goes into operation. In a cloud world, in which context changes constantly as new connections and integration points join and drop, the "edit and restart" model is unsustainable. Look for components that have online interfaces to update context and dynamically add or delete connections. Nothing is worse than rolling out an application and later realizing that some part of it can't really support dynamic topology shifts.

2. Plan for load balancing throughout the application. Many applications support load balancing at the Web server layer, but assume constant numbers (and IP addresses) for application components at other layers. With very large load variability, other layers need to be scalable and need to support load balancing to ensure consistent throughput. Don't design an application with the expectation that only two application components will reside at certain layers. Plan for dynamism and load balancing at all layers.

3. Plan for application scalability. Maybe this is hammering the point home too many times, but double or triple your capacity planning and application architecture assumptions — maybe even factor in a 10X growth possibility. When you plan for much larger scales, you pay attention to bottlenecks and plan to how to relieve them dynamically. If you don't expect scalability, you don't examine your architecture assumptions critically. So review your architecture for scalability bottlenecks.

4. Plan for dynamic application upgrades. Forty years ago, auto manufacturers took two weeks to change over factories to prepare for new model manufacturing. Toyota figured out how to do it in two hours. That meant they had to design for dynamic factory upgrades. Cloud computing, with the 24 hour use cycles, means no downtime for application upgrades. Architecting applications so that the topologies can be changed while users continue to access individual servers requires Toyota-like planning. Likewise, upgrading database schemas (and data sets) to support new application versions necessitates Toyota-like approaches. ■

Bernard Golden is CEO of consulting firm HyperStratus, which specializes in virtualization, cloud computing and related issues. He is also the author of "Virtualization for Dummies," the best-selling book on virtualization to date.

www.cio.com/article
January 19, 2010

ADDITIONAL READING

How to Build a Hybrid Cloud Computing Strategy

Forrester's James Staten explains why a hybrid cloud computing strategy — one that blends elements of public, internal and hosted cloud computing — may be smartest. He also has four tips for IT leaders on building a secure hybrid cloud plan.

By James Staten

Cloud infrastructures are a highly efficient evolution of server virtualization and the scale-out deployment model—but companies should note this evolutionary path isn't a fit for all applications. That being said, cloud computing platforms are more than just shared, multi-tenant infrastructures on the public Internet. Three infrastructure-as-a-service cloud deployment options are available to enterprises today, each with unique characteristics and economics that can help optimize application and service deployment objectives:

1. Public clouds. These deliver the best economies of scale, but their shared infrastructure model can limit configuration, security, and SLA specificity, making them a less-than-ideal fit for services using sensitive data that is subject to compliancy or safe harbor regulations.

2. Internal clouds. These sit within your data center and behind company-built protections, but they typically have modest economies of scale due to funding limitations and tend to be less automated.

3. Hosted clouds. Hosted clouds run at a service provider on resources that are walled off with enterprise-class protections but managed as a pool. These fall between the first two op-

tions, providing more custom protections like an internal cloud but with the greater economies of scale of being a service from a cloud provider.

Enterprises should build a strategy that leverages all three options via virtual private cloud technologies, which will result in a hybrid cloud strategy that optimizes business service deployment efficiencies.

Virtual private cloud is a technique for extending your organizational trust boundaries over a series of resources regardless of their deployment. It builds off the basic concept of a virtual private network (VPN), but is a more robust networking concept that lets you define and control addressing, topology, protocols, and encrypted communications for instances deployed to cloud computing platforms. Virtual private cloud technology defines the network security boundaries for the business service and the locations (types of deployments) where elements of these services can be placed or moved. These solutions can be enabled by two types of offerings: 1) those that focus on the network security layer; and 2) those that abstract the application tier across cloud deployment boundaries.

The evolution of cloud computing and virtual private cloud technologies add to the ever widening portfolio of infrastructure deployment options that help enterprises match the infrastructure to the needs of the application more efficiently and cost effectively than has been possible before. Here are a few tips to ensure the integration between these deployment infrastructures deliver the greatest value:

1. View cloud platforms as a portfolio of deployment options.
Look at creating a portfolio of cloud resources, composed of public, internal, and hosted cloud deployments linked together with traditional deployments via virtual private cloud. Look for ways to optimize—and automate—the deployment and management of your business services where most appropriate based on the application requirements and business/IT policies. You definitely want to learn to walk in cloud environments before going hog wild and spreading services across lots of cloud deployments but, as you gain experience, begin to craft a strategy that streamlines deployment expenses in much the same way that you optimize your Web site deployment. The ultimate goal is to speed IT service delivery while reducing costs.

2. Partner with enterprise architects to get the deployment model right.
Your Web infrastructure likely spans your data center, a content delivery network, and a hosting provider—each chosen for what they do best and connected to deliver optimum performance and customer experience. Approach cloud computing in the same way. Work with your enterprise architects to determine what types and portions of applications and what data sets are the best fit with the specific cloud deployment options discussed here and create joint policies to guide appropriate use as new applications are built or readied for the cloud.

3. Build a security model and policy with your CISO.
Work with your security and risk management professionals to understand what protections must be taken for what types of data so you can determine where applications can safely be deployed. Create and publish this as a policy for all application development professionals to help guide their use of cloud resources. Also, use this guide to set the SLAs for your private cloud, whether internal or hosted.

4. Ask your hosting providers about their cloud plans.
You likely have relationships with one or more hosting providers today. Get an understanding of what cloud services they provide or plan to provide to you in the future. Determine if a hosted cloud is a possibility from these vendors, what degree of configurability is available, and if they can provide virtual private cloud services between your data center and their cloud offerings. ■

James Staten is a Principal Analyst at Forrester Research, advising IT Infrastructure & Operations professionals on the transformation of the server and data center into more efficient, business-focused ecosystems. He is an expert on cloud computing.

www.cio.com/article
May 27, 2009

ADDITIONAL READING

Cloud Computing Shopping List: 4 Key Ingredients

By Rick Swanborg, *CIO*

Facing strong concerns about control and security, the cloud-computing trend has drifted somewhat-- away from the notion that all computing resources can be had from outside, and toward a vision of a data center magically transformed for easy connections to internal and external IT resources.

Sales of cloud-related technology are growing at 26 percent per year—six times the rate of IT spending overall, though they made up only about 5 percent of total IT revenue this year, according to IDC's Cloud Services Overview report. Defining what constitutes cloud-related spending is difficult, the report acknowledges, though it estimates global spending of $17.5

billion on cloud technologies in 2009 will grow to $44.2 billion by 2013.

Hybrid or internal clouds will be the rule, however; even in 2013, only about 10 percent of that spending will go specifically to public clouds, IDC predicts.

Hybrid cloud infrastructure isn't radically different from existing data-center best practices, except that all the pieces are supposed to fit neatly together using Internet-age interoperability standards rather than home-grown kludge, according to Chris Wolf, analyst at The Burton Group.

As you prepare spending plans that line up with a move to the cloud, consider these four items as key for your list.

1. Application Integration

Surprise: Software integration isn't the first thing most companies think about when building a cloud, but it's the most important one, according to Bernard Golden, CEO at cloud consulting firm HyperStratus, and CIO.com blogger.

Integration means more than just batch-processing chunks of data being traded between applications once or twice per day the way that was done on mainframes, according to Tom Fisher, vice president of cloud computing at SuccessFactors.com, a business-application SaaS provider in San Mateo, Calif.

Being able to provision and manage user identities from a single location across a range of applications is critical, especially for companies that have never been in the software-providing business before and don't view their IT as a primary product, he says.

"What you're looking for is take your schema and map it to PeopleSoft or another application so you can get more functional integration," Fisher says. "You're passing messages back and forth to each other with proper error-handling agreement so you can be more responsive. It's still not real time integration, but in most cases you don't really need that."

Sales of cloud-related technology are growing at 26 percent per year—six times the rate of IT spending overall.

2. Security

The second critical factor in building a useful cloud is the ability to federate—securely connect without completely merging—two networks, Golden says.

That requires layers of security, including multifactor authentication, identity brokers, access management and, in some cases, an external service provider who can provide that high a level of administrative control, according to Nico Popp, VP of product development at Verisign, which is considering adding a cloud-based cloud security service.

What it really requires is technology that doesn't yet exist, according to Wolf: an Information Authority that can act as a central repository for security data and control of applications, data and platforms within the cloud. Today, it's possible to assemble that function out of some of the pieces Popp mentions, but there is no single technology able to span all the platforms necessary to provide real control of even an internally housed cloud environment, Wolf says.

3. Virtual I/O

Having to squeeze data for a dozen VMs through a couple of NICs will keep you from scaling your VM cluster to cloud proportions, according to Bill Welty, Manager, IT Enterprise Architecture and Unix Operations at a large digital mapping firm.

"When you're in the dev/test stage, having eight or 10 [Gigabit Ethernet] cables per box is an incredible labeling issue; beyond that, forget it," Welty says. "Moving to virtual I/O is a concept shift--you can't touch most of the connections anymore--but you're moving stuff across a high-bandwidth backplane and you can reconfigure the SAN connections or the LANs without having to change cables."

Virtual I/O servers, such as the Xsiigo I/O Director servers Welty uses, can run 20Gbit/sec through a single cord and as many as 64 cords to a single server, connecting to a backplane with a total of 1,560Gbit/sec of bandwidth.

Concentrating so much bandwidth in one device saves space, power and cabling, Welty says, keeps network performance high and ultimately saves money on network gear.

"It becomes cost effective pretty quickly," Welty says of the Xsigo servers, which start around $28,000 through resellers such as Dell. "You end up getting three, four times the bandwidth at a quarter the price."

4. Storage

As noted, but storage continues to be the weak point—the hole into which one pours money—of both the virtualization and cloud-building world.

"Storage is going to continue to be one of the big costs of virtualization," Golden says. "Even if you turn 90 percent of your servers into images, you still have to store them somewhere." ■

www.networkworld.com/news
June 07, 2010