virtual

DATA CENTER

Volume 32

Weighing Your Options

ARE YOU in the cloud? A recent TechTarget survey said that more than 15% of respondents are using some type of cloud, while 17% plan to implement the technology in the next 12 months.

Before taking the plunge, find out which cloud computing architecture is best for your organization. And perhaps more important, learn the drawbacks of each option to make sure they aren't deal-breakers.

For example, organizations that opt for private clouds can look forward to maintaining their investment in their existing infrastructures. But private clouds may raise compliance issues because it's difficult to track the physical location of data.

And those choosing a public cloud might be enticed by the cost savings because they don't have to buy, install, operate or maintain servers or other equipment. At the same time, though, organizations working with a cloud provider will have to deal with tradeoffs in reliability and control. Dig deep into the differences in Stephen J. Bigelow's "Public vs. Private: Comparing Cloud Architectures."

In the cloud, it's important to remember that your users don't know and probably don't care where an application is running. They expect the application to perform just as well as it did when it was hosted in your data center. So how can you make sure that your applications are performing adequately in a cloud environment? Read Brien M. Posey's "Measuring and Optimizing the Cloud" for a new focus on cloud performance.

There are no hard-and-fast rules for what you can or should virtualize in your environment. In some cases, you might think twice about virtualizing an application if it would no longer look the same on the users' end. Mike Nelson categorizes the applications that can be hosted in a virtual environment in "How to Set Priorities for Virtualizing Applications."

How are you taking your IT operations to the next level? Send me an email at ccasatelli@techtarget.com, and tell me all about it. ■

CHRISTINE CASATELLI
Editor, *Virtual Data Center*

BROCADE VDX™ CLOUD NETWORKING NOW HAS AN ETHERNET FABRIC

PUT THE CLOUD **WITHIN REACH.**

A cloud-optimized network may be closer than you think. With revolutionary Brocade® VDX™ 6720 Data Center Switches featuring Brocade VCS™ technology, you get the industry's first true Ethernet fabric solution that can help dramatically simplify your architecture and streamline management. The result is unequalled network performance, greater utilization—and a clear path to the cloud.

LEARN MORE AT www.brocade.com/vdx

BROCADE

Public vs. Private: Comparing Cloud Architectures

BEFORE PICKING EITHER OPTION, IT'S IMPORTANT TO EXAMINE THE PRACTICAL LIMITATIONS OF PRIVATE AND PUBLIC CLOUDS.

BY STEPHEN J. BIGELOW

IT'S NO secret that cloud computing has been a misunderstood victim of marketing hype. But regardless of how you define the technologies, it's clear that both private and public cloud computing are slowly being adopted among modern businesses. According to a recent TechTarget survey, roughly 15% of more than 1,000 IT respondents indicated that they have implemented some type of cloud, while almost 17% expect to implement some type of cloud computing over the next 12 months.

The successful move to a cloud is hardly quick or simple, and perhaps the biggest challenge that IT administrators face is choosing between private and public cloud computing architectures. Before picking either option, it's important to examine the practical limitations of private and public cloud computing.

First, let's consider the private cloud, which may be a dedicated, or leased, portion of a provider's hosted facilities. But in most cases, a private cloud is the conversion of an existing traditional data center into a utility-based computing architecture—you own it, secure it and maintain it just as you always have.

The main difference is that a private cloud shifts the emphasis to computing on demand. This generally changes the way that data center hardware is used and shifts the way that departments or stakeholders pay for those computing resources.

There are many reasons for organizations to select private clouds. Cost is always an obvious issue, but it's not necessarily the principal motivator. Organizations that opt for private clouds are usually most concerned with preserving their investment in their existing infrastructures. This was the reason cited by more than 35% of TechTarget survey respondents, and it

makes sense for businesses that have already spent significant capital to create their data centers in the first place.

Another 17% said they were motivated by disaster recovery (DR) and business continuity benefits, while 16% said they wanted the security of continued IT ownership. Eleven percent said they sought to leverage the self-service and automation capabilities that are a hallmark of true cloud computing.

PRIVATE CLOUDS AND COMPLIANCE

Data location may also become a compliance or regulatory issue for private clouds if the organization deploys them across geographic boundaries. It's extremely difficult to track the physical location of data in a true cloud—as opposed to a virtualized data center—so close attention to regulatory adherence is still required.

Even when an organization deploys a private cloud successfully, making the most of that cloud can be a formidable task. Efficient cloud management involves the establishment of standard operating procedures, a high level of automation, sharing of computing resources among departments or stakeholders, along with a regimented standardization through a single service-level agreement (SLA).

This set of criteria allows an administrator to run the entire pool of cloud resources at an exceptionally high utilization rate and take advantage of significant cost benefits. But these criteria are not normal considerations in a typical data center. "If you don't have the discipline of the regimented resources, if you don't have automation to drive this forward, and if you don't have sharing, you can't reach that degree of utilization," said James Staten, vice president and principal analyst at Forrester Research.

NOT AN ALL-OR-NOTHING DECISION

The transition to a private cloud is not an all-or-nothing decision. Experts recommend a gradual transition, starting with a limited cloud infrastructure and running new or low-priority applications on that "mini cloud" within the environment. This cautious approach allows ample time to gain familiarity and comfort with cloud technology and examine the impact of tradeoffs on your particular organization. Success may accelerate subsequent projects, while problems may slow further transition. But a well tested alpha phase of adoption can minimize unnecessary disruption to your business.

Data protection within the private cloud is another important consideration. Organizations that prefer more traditional backup mechanisms can continue to use that approach in the cloud, but the backup tools—such as Symantec's Backup Exec—must be upgraded to a cloud-compatible version.

Other organizations with more exacting recovery needs may opt for some combination of replication/snapshots and clustering as a means of mixing high availability and data

protection for fast restoration. Perhaps the biggest challenge is disaster recovery (DR). The notion of native DR in a cloud is a misconception, and adopters must pay close attention to the DR plan implemented in their private cloud.

Deploying and testing data protection schemes is another central reason why companies should approach private cloud adoption gradually. This approach will leave mission-critical applications with the most stringent data protection requirements until the private cloud is well understood and developed.

A public cloud employs the same basic technologies used to create a private cloud, but public clouds such as Microsoft's Azure serve a fundamentally different purpose. In virtually all cases, a public cloud is accessed as a service through a provider that hosts all of the hardware, services and

ANOTHER LOOK AT SLAs

FOR A PUBLIC cloud to be economical, it's imperative for the provider to share computing resources among different users and automate every possible aspect of the business. The net result is an inexpensive—yet inflexible—source of utility computing.

True public cloud providers—as opposed to hosted service providers that advertise themselves as cloud providers—just cannot afford to cater to the unique needs of individual clients. As a result, service-level agreements (SLAs) are quite rigid and almost always side with the provider, especially when it comes to security and data protection.

The notion of a business committing its important applications and data to the care of a cloud provider that assumes no significant responsibility for the safety, availability or even the physical location of that data is enough to scare off even the hardiest CTOs. But experts say that this is slowly changing and will continue to shift as public cloud providers develop the tools and tactics necessary to address serious business concerns.

Changes are coming to public cloud SLAs, said Laurie McCabe, partner with the SMB Group in Northboro, Mass. But McCabe said she sees the emphasis squarely on compensation for downtime, which can have the most immediate impact on user productivity.

"I'm losing money because I can't enter sales or whatever it is," McCabe said. "That's where cloud providers will focus their efforts," she said. At the moment, there is no generally accepted formula for downtime compensation, but it's an area that public cloud providers are targeting, McCabe said. ■

EDITOR'S LETTER

PUBLIC VS. PRIVATE: COMPARING CLOUD ARCHITECTURES

HOW TO SET PRIORITIES FOR VIRTUALIZING APPLICATIONS

MEASURING AND OPTIMIZING THE CLOUD

applications. Public cloud operation also leverages significant economies of scale, potentially supporting vast server and storage resources spread across numerous geographical areas—even globally.

COST SAVINGS WITH PUBLIC CLOUDS

The most prevalent benefit in a public cloud is cost savings. Organizations do not buy, install, operate or maintain servers or other equipment as they would with private clouds. Organizations may develop their own applications, but they are hosted by a public cloud provider.

Public clouds are also designed as true multi-tenant environments. This allows a huge number of users to share a provider's computing resources, which also makes public cloud services highly cost-efficient for users. Experts agree that the cost of a public cloud is typically far lower than a traditional data center and private cloud because there is no capital expense. There is also greater flexibility for users because organizations pay only for the computing resources that they actually use.

According to the respondents who use a public cloud, more than 31% said they were motivated by a reduction in infrastructure costs, another 26% sought more flexibility than traditional hosted services, almost 21% wanted a more scalable infrastructure without up-front costs, and 11% sought lower IT staff costs.

Security, reliability and regulatory compliance remain the principal concerns with public clouds. The concept of multi-tenancy supports strong economics by allowing high utilization. There is persistent concern about potential vulnerabilities or attacks that might be used to exploit the hypervisor and affect other users' sessions. Although there have been no reports of successful attacks, security remains a serious point of contention for users considering public clouds.

"Security is a shared burden but not shared equally," said Staten, adding that services like Azure provide security to the point of abstraction but no higher. For example, Azure would presumably be responsible for server security by providing logs to ensure that no violations occurred.

Although there have been no reports of successful attacks, security remains a serious point of contention for users considering public clouds.

But organizations are responsible for security on their own applications. That includes design oversights like open ports or weak logins. "This uneven handshake is the mistake most people make when they use cloud

EDITOR'S LETTER

PUBLIC VS. PRIVATE: COMPARING CLOUD ARCHITECTURES

HOW TO SET PRIORITIES FOR VIRTUALIZING APPLICATIONS

MEASURING AND OPTIMIZING THE CLOUD

resources in an unsecure fashion," Staten said.

Reliability and availability are other issues in public clouds. Even when providers implement clusters, redundant sites and other tactics to improve availability, outages can occur periodically. Because downtime can affect the ability to access cloud applications—and affect your organization's ability to function—the move to a public cloud should include a careful evaluation of availability and the consequences of normal and extended downtime.

LOCATION, LOCATION, LOCATION

The physical location of sensitive data may also become a regulatory problem for public clouds. For example, some industries are not allowed to

PREPARING FOR DISASTER

FOR BUSINESSES TO leverage public cloud technology, data must be secure and accessible to users. Today's global challenges make this goal far from certain.

Companies do encounter natural and man-made disasters, fall victim to hacking and go out of business entirely. What happens if your cloud provider falters?

Experts say that any move to the public cloud should include a careful consideration of application and data portability. For example, it may be necessary to move away from one cloud provider and engage another. Otherwise, you might need to restore that application within your own data center—effectively taking the application out of the cloud.

In either case, part of cloud disaster preparation involves some form of transition plan. "The role of IT is changing," said Laurie McCabe, partner with the SMB Group in Northboro, Mass. "IT people have to have a strategy or plan for determining what they are comfortable with, creating a contingency plan should they need to make an adjustment."

Some businesses develop a relationship with a second public cloud provider and have already worked out a process for mapping their applications to the second environment. Other organizations may adopt gateway tools such as CloudSwitch to minimize architectural changes and ease application porting between clouds. And still other organizations may retain and periodically test local servers that can start up virtual machines to run applications in-house.

The good news is that most businesses have time to think ahead and make plans for a transition away from or between clouds. As long as steps are taken before trouble actually strikes, planning for a serious transition is simply another aspect of disaster planning. ■

store data outside certain geographical boundaries, but a true cloud does not observe such boundaries.

A file created by a company in Cleveland may wind up on a cloud provider's storage array in Toronto or London. If an organization cannot provide the physical location of its data, it's impossible to pass a regulatory audit, and this may expose an organization to serious legal issues. Ultimately, it's important to examine the provider's SLA closely and understand the level of security, reliability and regulatory support.

As with the private cloud, a transition to a public cloud should be made slowly and methodically, starting with low-priority applications and services.

As with the private cloud, a transition to a public cloud should be made slowly and methodically, starting with low-priority applications and services. In many cases, these are new applications and services that are not even in data centers yet. For example, an organization may continue using an existing tool in-house but adopt a new cloud version to see how the two versions compare or perform.

The point here is that a true public cloud—as opposed to a hosted service that's simply labeled as a cloud—is a one-trick pony with serious restrictions that may render it inadequate for every application or service. "When you go to the self-service portal where you deploy your application, you're confronted by all of these restrictions," Staten said. "If it's a true public cloud, then you have to adapt to it—it doesn't adapt to you."

Don't assume that a public cloud will protect your data. Although a cloud provider can rebuild lost content to some extent, it may be far below the protection that your application and its data require. For example, it may be necessary to back up data from the cloud to a local storage system and then replicate that data yourself for DR purposes. A risk analysis can help identify the level of protection needed for your data and allow you to plan for that protection.

HYPE VS. REALITY

Deciding between a public and private cloud can be a serious challenge because each one is not appropriate for every application, user group or organization. The hype that surrounds cloud computing has caused confusion, prompting many organizations to steer clear of the technology entirely.

According to TechTarget's recent survey, more than 40% of IT respondents said they have chosen to avoid cloud computing because of security

concerns. Another 30% of respondents said that their business applications do not translate well to a cloud environment, while 27% said that they do not meter or charge back for computing resource use. More than 26% of respondents expressed concern about data storage integrity and compliance, and the remaining respondents cite concerns over virtualization, in-house skills and cloud complexity—among others—as reasons for not using cloud technology.

But experts warn organizations that do venture into cloud technology not to underestimate the challenges. Cloud computing provides enormous potential for efficiency and economy, but achieving those goals within a cloud is difficult. "If you don't put the right kind of workload into a cloud environment, it will not yield cost savings," Staten said. "In fact, it may cost you more."

Generally speaking, the best applications for a cloud deployment are Web-based, customer-facing applications where many users need to access applications from various locations using different levels of connectivity. Examples often include email, collaboration and support portals. These are the low-hanging fruit of cloud computing.

The case for a cloud is not so easy for applications that are heavily tweaked or used only by a small group at one location. These may be better served with a more traditional client/server deployment in an organization's data center.

EFFICIENCY OF CLOUD ECONOMICS

By its very nature, an efficient cloud is a highly automated and regimented environment that can provide only a fixed set of assets and features to users, Staten said. When organizations try to satisfy the unique needs and wants of every stakeholder, the efficiency of cloud economics is quickly lost.

Not only are assets and features limited in the cloud, but applications are also aggressively removed so that computing resources can be reused as needed. This concept may be hard to grasp for many organizations where applications seem to consume computing resources long after their usefulness is over. "This is a concept that is hard for a lot of enterprises to get, which is why we see so few true private clouds," Staten said.

Ultimately, experts see many of the biggest cloud concerns easing in the coming years. For example, cloud vendors are actively working to overcome data location worries. As a result, some form of accreditation may eventually identify cloud providers that adopt verifiable data location and protection practices.

The future may yield a broader range of clouds, each providing different sets of assets and catering to more demanding user requirements, Staten said. At the same time, it's important to remember that neither public nor private clouds will handle every computing need. Traditional data centers will continue well into the future—even as the clouds move in. ■

move fearlessly among the clouds.

it's open, yet secure.
agile, yet efficient.
virtual, yet real.

it builds on your current network,
combining servers, networking,
storage and virtualization.
so you can deploy
applications in minutes,
not weeks.

move among the clouds.
grow your business.
and turn silos into
new ways of working together.

introducing the Cisco data center
business advantage.

cisco.com/go/dcbusiness

together we are
the human network.

cisco

How to Set Priorities for Virtualizing Applications

DON'T FORGET ABOUT USER EXPERIENCE WHEN ESTABLISHING YOUR VIRTUAL APPLICATION PLAN. **BY MIKE NELSON**

DECIDING WHICH applications to run on virtual servers can be difficult. Virtualizing applications requires careful consideration, so it's best to prioritize which apps are the best candidates. Successfully virtualized applications improve resource usage and disaster recovery, reduce hardware and power costs and provide greater flexibility.

IT administrators, management and business people alike often ask me the same question: "Which applications do you recommend moving to virtual servers?" There are no hard-and-fast rules for what you can or should virtualize in your environment. The benefits of server consolidation and virtualization have motivated many organizations to run virtualized applications, even mission-critical or resource-intensive apps.

But establishing your priorities for virtualizing applications is not as daunting as you might think. The applications that can be hosted in a virtual environment can be divided into four buckets. The following are my unconventional names for those categories:

- **Heavy hitters.** This category includes tier-one applications—and applications that have consistently high resource utilization. That includes some SAP programs, financial apps and highly transactional databases. This category also includes any mission-critical redundancy-candidate applications.

- **Hills and valleys.** Most of the time, these applications use medium to high levels of resources, but they also have slow periods of barely any resource usage—although these periods do not last long. In this category, month-end financial and sales applications are my favorite candidates for virtualized applications.

■ **Smooth and steady.** Generally, these applications have steady resource usage and experience usage spikes only briefly—if at all. For example, a sudden need for more resources may coincide with events such as backups or antivirus scans.

■ **Hardware lockers.** Very few exist, but some older applications still demand hardware dependencies for license compliance or direct-console access. Most of these application versions are no longer produced or even supported, but for one reason or another, some shops just cannot part with them. Most of these apps cannot or should not function as virtualized applications.

If getting the most immediate bang for your buck is your goal, I suggest virtualizing applications in the smooth-and-steady category. Consider license servers, monitoring applications, patching services and even most file and print applications. Just about any application that has lower-end utilization in any combination of the four core resources—RAM, CPU, disk and network—is a good candidate for virtualization.

Virtualized applications in the top two categories—the heavy-hitters and the hills and valleys—require more resources, testing and legwork. Some IT pros say you should not or cannot virtualize these types of applications, but I don't believe in blanket statements. Even mission-critical applications are possibilities for virtualization. Still, some common virtualization roadblocks include vendor support and licensing issues. It's best to determine support and licensing rules before you invest a lot of time and energy in your plan to virtualize applications.

Just about any application that has lower-end utilization in any combination of the four core resources —RAM, CPU, disk and network —is a good candidate for virtualization

Remember to cover all your bases. Run performance testing and baselines on a physical server, then compare the results to those on a test virtual server. Make sure that you keep the application owners or users informed about the process. And most of all, don't be afraid to push back on application vendors. They sometimes offer the most resistance to virtualizing applications.

Once you prioritize which applications to virtualize, it's a good idea to test their performance. Application performance testing is the best way to learn how virtual applications will fare in your virtual infrastructure and what

the end-user experience will be like.

Virtual application performance testing is exactly that—testing. Some people consider it a science, but I think of it as an art form. You can get creative with testing to generate the best performance, but it won't exactly replicate how your virtual applications will perform in production.

PRE-TESTING FACTORS: SOFTWARE AND USERS

Application performance testing is an essential part of any solid virtual application plan, particularly for tier-one and mission-critical applications. By performing some quick load tests—with some willing users enlisted as guinea pigs—you can experiment with the hypervisor and the apps themselves.

But as you tinker with virtual applications, don't forget about user experience. If a virtualized application runs as well as it did on physical hardware—or better—you shouldn't worry too much about user expectations.

In some cases, though, you might think twice about virtualizing an application if, as a result, it would no longer look the same on the users' end. For example, if an application is visually intensive and is not a true client/server app, but rather a simple Remote Desktop Protocol session-accessed application, a user's daily interactions with the application are most important. Virtualizing applications saves resources and space—which is great for admins—but it won't always benefit users.

Finally, remember that not everyone has the luxury of a full-blown testing lab or the available users for application performance testing. Some software manufacturers produce load-testing software from application vendors themselves to third-party providers.

Application performance testing is an essential part of any solid virtual application plan, particularly for tier-one and mission-critical applications.

PERFORMANCE TESTING CONSIDERATIONS

The following are some important factors to consider as well as questions to ask before you begin virtual application performance testing:

- **Host affinity.** What are the application vendor requirements, if any? Can you run one application on the same host—or even in the same Distributed Resource Scheduler or high-availability cluster—as another application? Host affinity is a requirement in some infrastructures, and in others it just makes good sense.

■ **Resource affinity.** Again, is this a vendor requirement? It may be necessary to dedicate quantities of the core five resources—CPU, RAM, disk, network interface cards and USB—to ensure the best performance and support for virtual applications.

■ **Shares and resource pools.** These resources don't enter into the equation until the end of the process, but you should consider resource quantities up front. Determine the resources your guest machine and its virtual applications will use. Don't necessarily leave the hypervisor's resource allocation at the default amount. You can fine-tune shares throughout the application performance testing process to see what's best for each virtualized application.

■ **Dedicated clusters.** To improve performance and management, try forming a cluster with an isolated set of servers that host common or interacting applications.

■ **Application stacking.** With the ultrafast and capacity-heavy hosts that exist today, it's possible to stack two or three applications on one virtual server. But with bigger applications, I advise against it. Stacking can be a great thing, but it can also be a problem if virtual applications don't cooperate with one another.

■ **Inter-application cooperation.** Some applications just don't play well with others in the sandbox. We know that all too well from the physical server world. Some virtualized applications won't interact well with others on the same guest, host or cluster. And some just cannot be virtualized at all.

Some virtualized applications won't interact well with others on the same guest, host or cluster. And some just cannot be virtualized at all.

As you begin virtual application performance testing, create a baseline of servers and a load that's as similar as possible to what's on the physical server. It's not always possible to test virtual applications with 100% accuracy, but you should be able to get close.

Also, keep in mind that some problems might emerge only after testing is done and you compare the data. If you think of application performance testing as an art, you'll find some creative ways to boost performance. ■

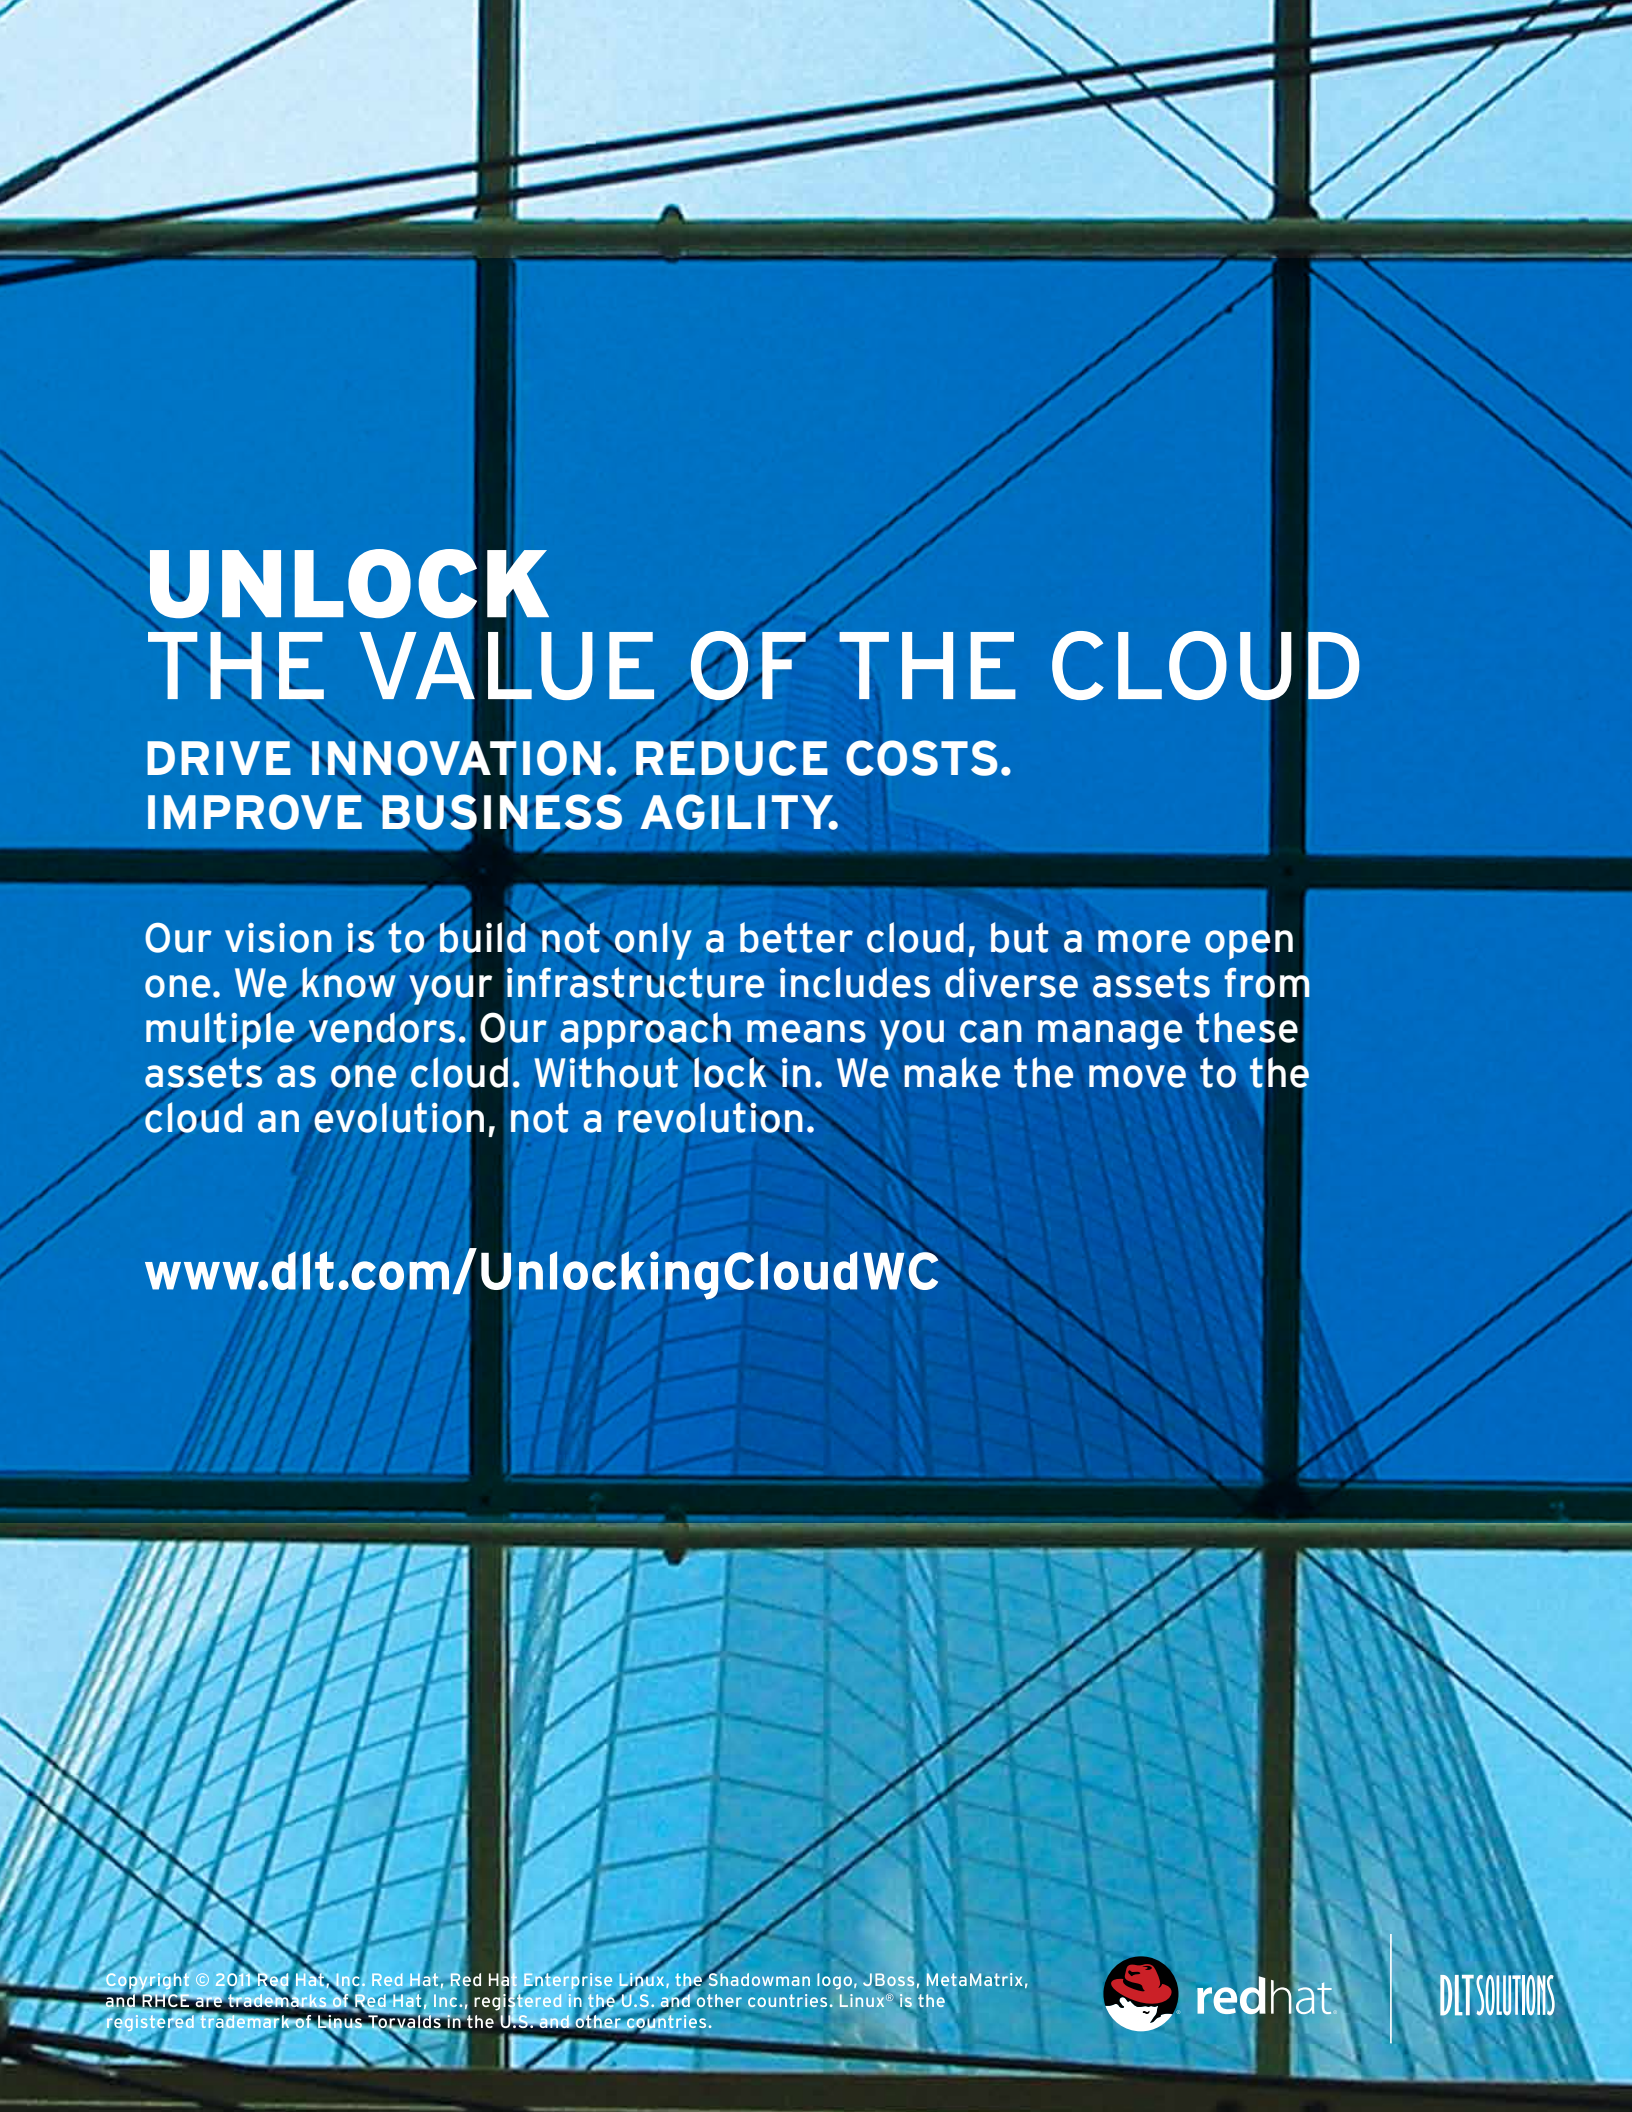
UNLOCK
THE VALUE OF THE CLOUD
DRIVE INNOVATION. REDUCE COSTS.
IMPROVE BUSINESS AGILITY.
Our vision is to build not only a better cloud, but a more open one. We know your infrastructure includes diverse assets from multiple vendors. Our approach means you can manage these assets as one cloud. Without lock in. We make the move to the cloud an evolution, not a revolution.
www.dlt.com/UnlockingCloudWC
Copyright © 2011 Red Hat, Inc. Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, MetaMatrix, and RHCE are trademarks of Red Hat, Inc., registered in the U.S. and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.
redhat.
DLT SOLUTIONS

Measuring and Optimizing the Cloud

EVALUATE PERFORMANCE DATA IN A CONSISTENT MANNER WITH TOOLS SPECIALLY DESIGNED FOR THE CLOUD. **BY BRIEN M. POSEY**

MANY IT departments are feeling an extraordinary amount of pressure lately to move applications to the cloud. This pressure often comes from those who are anxious to capitalize on the vast savings advertised by cloud service providers.

Although there is no denying that running applications in the cloud can save money—especially in the form of startup costs—some organizations are beginning to discover that the cloud experience isn't always all it's cracked up to be.

As cloud technology begins to mature, security and availability become higher priorities. And although performance can be just as important, it is sometimes overlooked.

The bottom line is that your users neither know nor care where an application is running. They expect the application to perform just as well in the cloud as it did when it was hosted in your data center.

MONITORING PERFORMANCE IN THE CLOUD

So how can you make sure that your applications are performing adequately in a cloud environment? This simple question can actually be very difficult to answer. Odds are that the performance monitoring tools that you used when the application was being hosted in your local data center will not work in a cloud environment.

First of all, you can't depend on your cloud service provider when it comes to performance. Typically, cloud service providers will not even address the issue of performance within their service-level agreements (SLAs). These types of agreements usually cover service availability but not application performance.

There are some cloud providers that will give you performance bench-

marks for your hosted applications. The problem is that this benchmarking data probably does not accurately reflect the end-user experience.

Think about it for a moment: Have you ever heard a cloud service provider advertise that your application is going to perform better than it does now once you move that application to the cloud? Of course not. Cloud service providers know that no matter how good their data centers are, there is nothing that they can do about Internet latency.

Cloud service providers know that no matter how good their data centers are, there is nothing that they can do about Internet latency.

So if a cloud service provider gives you performance benchmarking data, it's likely that the data was collected from within the service provider's own data center. If that's the case, it probably does not take Internet latency into account.

Even if a service provider does try to account for Internet latency, geographic proximity to the data center can make a huge difference in end-user experience. Unless a service provider actually measures your application's performance from your facility, then the performance data that it provides probably won't accurately represent true performance.

So how can you monitor the performance of your applications that are running in the cloud? The truth is that there is no one-size-fits-all solution. Each cloud provider exposes different parts of its network infrastructure. As a result, a solution that works with one cloud provider may not work with another.

THREE TYPES OF CLOUDS

Your options for monitoring cloud performance depend on what type of cloud you are using. The three main types of clouds used for hosting applications include:

- **Infrastructure as a Service (Iaas)**—When an organization subscribes to an IaaS cloud, it is typically given access to a virtual server and a storage pool. IaaS clouds such as Amazon's EC2 provide the greatest opportunities for performance monitoring because you have full access to the server's operating system and to the applications that are running on it.

- **Platform as a Service (PaaS)**—Clouds such as Google App Engine are platforms for running Web-based applications. Although it is sometimes possible to build performance metrics into a Web application, PaaS clouds do not provide access to the underlying operating system.

■ **Software as a Service (SaaS)**—These types of clouds allow organizations to run commercial SaaS applications such as Salesforce.com or Microsoft Exchange in a hosted environment. Although the hosting provider may give you access to performance reports in some cases, there are few options for gathering your own performance metrics.

If the application that you need to monitor was developed in-house, then one of the most effective techniques for monitoring the application's performance is to have its developers create an agent that can be embedded into the application. The general idea is for the agent to note the time just before phoning home. When the agent's transmission is received, the time stamp on the transmission can be compared against the time when the transmission was received, which will allow you to measure latency.

TOOLS YOU CAN USE

If you happen to have an application that is coded in Java, you don't have to worry about the hassles of developing a custom agent. There are at least two software publishers that offer Java-based code for monitoring cloud performance. One is AppDynamics, which is designed to monitor

TAKING IT TO THE NEXT LEVEL

MONITORING THE performance of applications in the cloud is important to ensure that end users' needs are being met. But there are other reasons why you should measure cloud performance.

One of the main objectives that draws organizations toward hosting applications in the cloud is scalability. Running an application in the cloud frees an organization from dealing with the constraints of physical hardware. Some cloud providers are able to dynamically allocate additional hardware resources to applications to help them perform better during peak usage periods.

But as nice as dynamic scalability is, cloud providers tend to charge customers based on the hardware resources that they use. This means if you allocate excessive resources to your hosted application, you will end up paying more than you need to. Alternately, if you don't allocate enough resources to the application, it will perform poorly.

By monitoring an application's performance and comparing the performance metrics to the application's usage statistics, you can determine at what point additional resources need to be dynamically allocated. Finding the optimal hardware usage levels will help to save your organization money while ensuring that applications perform as expected. ■

distributed Java-based applications in the cloud. The software tool provides a dashboard for tracking application performance and helps to isolate the cause of any problems that may be detected. AppDynamics also detects memory leaks in real time and locates their sources.

Another tool that exists for monitoring applications in the cloud is New Relic, which works with Ruby, Java, PHP and .NET applications. Two features make New Relic unique. First, it is designed to be simple to deploy. The other is that the basic monitoring solution is free. If you also need diagnostic and optimization capabilities, New Relic offers those tools for a fee.

AppDynamics and New Relic are suitable for use with IaaS and PaaS clouds, but they do not work with SaaS. But another application called FactFinder from BlueStripe Software is said to work with any application, regardless of whether it is a packaged application like Microsoft Exchange or a custom app that was developed in-house.

FactFinder works by using protocol-level visibility into each node that is involved in running a distributed application. By monitoring packets on a hop-by-hop basis, FactFinder is able to create a map of the infrastructure that is running the distributed application and locate bottlenecks when performance problems occur.

Regardless of which method you use to collect performance data for your cloud-based application, it is critically important that you measure the performance data in a uniform manner.

Regardless of which method you use to collect performance data for your cloud-based application, it is critically important that you measure the performance data in a uniform manner. Unless you are consistent with your choice of metrics and the tools that you use, the historical performance data you collect will be meaningless. ■

thawte®

Building Trust Around The Globe

When you want to establish trusted relationships with anyone, anywhere on the internet, turn to Thawte. Securing Web sites around the globe with:

- strong SSL encryption
- expansive browser support
- multi-lingual customer support
- recognized trust seal in 18 languages

Offering outstanding value, Thawte is for those who know technology. Secure your site today with a Thawte SSL Certificate.

www.thawte.com

ABOUT THE AUTHORS

Stephen J. Bigelow, a senior technology editor in the Data Center and Virtualization Media Group at TechTarget Inc., has more than 15 years of technical writing experience in the PC/technology industry. He holds a bachelor of science in electrical engineering, along with CompTIA A+, Network+, Security+ and Server+ certifications, and has written hundreds of articles and more than 15 feature books on computer troubleshooting, including *Bigelow's PC Hardware Desk Reference* and *Bigelow's PC Hardware Annoyances*. Contact him at sbigelow@techtarget.com.

Mike Nelson has been in IT for more than 20 years, with exposure to a diverse field of technologies and solutions. He has devoted more than half a decade to virtualization and server-based computing. Currently, Nelson is a senior analyst at a Fortune 100 company in the midwest United States.

Brien M. Posey has received Microsoft's Most Valuable Professional award seven times for his work with Windows Server, IIS, file systems/storage and Exchange Server. He has served as CIO for a nationwide chain of hospitals and healthcare facilities and was once a network administrator for Fort Knox.

TechTarget
Where Serious Technology Buyers Decide

Cathleen Gagne
Editorial Director
cgagne@techtarget.com

Christine Casatelli
Editor
ccasatelli@techtarget.com

Linda Koury
Director of Online Design
lkoury@techtarget.com

Marc Laplante
Publisher
mlaplante@techtarget.com

Peter Larkin
Associate Publisher
plarkin@techtarget.com

TechTarget
275 Grove Street
Newton, MA 02466
www.techtarget.com

See ad page 3

- Best Practices for Virtual Infrastructure Management
- Data Center Evolution Without Revolution
- Simplifying the Virtual Data Center

See ad page 11

- Move fearlessly among the clouds. Evolve your data center. Transform your business. Discover better ways to work. Together.

See ad page 16

- Cloud Infrastructure Matters
- Red Hat Cloud Foundations: Cloud 101
- Unlocking the Value of Cloud On-Demand Webcast

See ad page 21

- thawte products
- free thawte guides
- thawte live chat