

# Into The Cloud

## How To Plug

**THE BEAUTY OF CLOUD COMPUTING** is that little more than a user ID and a credit card will get you started. That's the problem, too. Anything this easy is bound to create problems for IT departments that aren't prepared.

We've experienced this phenomenon many times before, where a technology's ease of adoption translates into unforeseen management challenges. Virtualization resulted in virtual machine sprawl; smartphones ushered in new security risks; instant messaging raised corporate governance concerns.

The purpose of this report is to show IT managers how to maximize the benefits of cloud computing—including ease of use, flexibility, and lower costs—while minimizing the risks. It's a how-to guide to licensing, management tools, bandwidth, security, and architecture.

This report shows we're still in the early stages of cloud computing, which means the tools and techniques are still evolving. After two years of testing, for example, Amazon Web Services' Elastic Compute Cloud service became generally available just a few weeks ago, and enterprise capabilities such as monitoring, management, and load balancing are still on the road map. Likewise, Google's App Engine is in preview mode. Microsoft's Azure cloud services are in preview, too, available only with limited functionality to Windows developers, not other early adopters.

Yet the time to begin planning is now, both as a way of gaining hands-on experience with this new IT delivery model—including the glitches and gotchas—and of getting ahead of people inside your company who might be contemplating tapping cloud services on their own. Here's how to get started.

This article, the fourth in a four-part series, is one element of a multimedia package on business innovation. For links to related stories and additional editorial content, go to [businessinnovation.cmp.com](http://businessinnovation.cmp.com).



**It's easy  
to get started,  
but hard  
to get right**



MANAGEMENT

# The Cloud, Under Control

**T**OOLS FOR MANAGING cloud services range from easy-to-use dashboards that let you create virtual software stacks in minutes to enterprise-class platforms that handle a full range of provisioning and management tasks. The further you get into cloud computing, the more you'll need those higher-end tools.

Amazon.com, Google, and other cloud service vendors provide the basic tools to get started. The administrative console in Google's App Engine, for example, shows traffic levels, bandwidth and CPU utilization, and error rates of Google-hosted apps, and lets you dig into the log file for other detailed data. You also can use it to control administrative rights and manage application upgrades.

However, App Engine is still in "preview" mode, which means those tools will max out as requirements increase. "We're still missing some pieces," admits Google product manager Pete Koomen.

Cloud service providers, startups, and systems management vendors are scrambling to give customers more full-featured tools to manage resources in cloud environments. Amazon says a new management console and cloud-monitoring capabilities for its Elastic Compute Cloud service are "coming soon." Amazon already provides basic functionality, such as the ability to create Amazon Machine Images using a command-line interface. The console will let users configure and manage EC2 resources, while the monitoring capabilities will include real-time metrics on EC2 instances and "availability zones"—those parts of the Amazon infrastructure that customers select for redundancy and maximum availability. Load balancing and auto-scaling also are in Amazon's 2009 plans.

Companies that specialize in cloud management are another option. RightScale's platform—offered as a hosted service—includes a management dashboard, database and Web site management, batch processing, multiserver deployment capability, and the ability to scale automatically. A bare-bones developer's edition is available free, but most IT departments will need one of RightScale's three other editions (Website, Grid,

## Cloud Checklist

- ✓ **EXTEND** your IT architecture to work with cloud services
- ✓ **STANDARDIZE** on one or two cloud service providers
- ✓ **DEPLOY** enterprise-class monitoring and management tools
- ✓ **MOVE** toward federated identity management if warranted
- ✓ **ENCRYPT** data stored in the cloud where appropriate
- ✓ **DEVELOP** a backup plan in case your cloud service fails
- ✓ **ADD** bandwidth to support an increase in network traffic
- ✓ **AVOID** vendor lock-in by opting for open standards

and Premium), which start at \$500 a month plus a one-time fee of \$2,500.

Founded last year, RightScale got its start with Amazon Web Services and is now expanding to manage other public cloud services, including FlexiScale's and Go-Grid's. RightScale also has a version of its platform for the University of California at Santa Barbara's Eucalyptus Public Cloud, an implementation of the open source Eucalyptus software for cloud computing on clustered servers. It's essentially a research and testing project, but the goal is to be able to manage public clouds and Eucalyptus-based private clouds from RightScale's dashboard.

### AS EASY AS WEB APPS

IT departments experienced at managing Web apps and infrastructure will find that cloud computing has similarities. "If you can manage Web apps, you can manage cloud apps," says Javier Soltero, CEO of Hyperic, which has a version of its Web application monitoring software that runs in Amazon Web Services.

Hyperic IQ consists of a central management server—which typically runs on a company's on-premises server—and agents that reside on Web servers and report back to the management server with availability, performance, and other metrics. With the just-released IQ 4.0, the Hyperic server has been configured as an Amazon Machine Image in EC2. For IT administrators, that means ease of deployment, subscription pricing, and faster performance. Hyperic IQ's capabilities include auto-discovery of software, diagnostics, alerts, analysis and reporting, and other tools.

Beware of an out-of-sight, out-of-mind attitude toward cloud apps. "The notion that, because you're de-

ploying an application in the cloud, it's inherently free from monitoring and management is one of the great lies of cloud computing," Soltero says. "Code is inherently flawed and technology breaks, so you've got to be able to monitor that."

Kaavo also specializes in multicloud management. The startup's platform supports server monitoring, LAMP software configuration in the cloud, load management, software audits, patch management, run-time configuration management, and notifications and alerts. Its Infrastructure and Middleware On Demand software has been

out in a free test version; a general release is due soon. In Kaavo's favor is its management team: Founder and CEO Jamal Mazhar is a Sun-certified J2EE architect, and CTO Shahzad Pervez is a former director of IT and enterprise architect at major companies.

Leading systems management vendors are bringing new controls to the cloud, too. IBM's Tivoli unit plans to inject cloud management into its Service Request Manager, Provisioning Manager, and Monitoring prod-

ucts, says Dennis Quan, IBM Software's director of development for autonomic computing. IBM also wants to boost confidence in cloud security by giving customers greater "control" over the systems that house their data in the cloud, although Quan didn't say how IBM will do that.

Microsoft's answer to cloud management is still in development. It introduced the Windows Azure operating

system and related Azure Services Platform in October but hasn't said when Azure cloud services would be available, although the development tools and basic building blocks for getting started are

available to developers. Also in October, senior VP Bob Muglia demoed a version of Microsoft's System Center enterprise management platform, code-named Atlanta, that will run in Microsoft's cloud.

As all this activity shows, vendors are hurriedly developing enterprise-class controls for emerging cloud services. The challenge for IT administrators is to get the tools in place before cloud service adoption takes on a life of its own. —JOHN FOLEY (jfoley@techweb.com)

### Hands-On Experience

**IN PERSON** Learn how to harness the cloud for your company at TechWeb's Cloud Connect event, Jan. 20-22 in Mountain View, Calif. See the agenda and register to attend:

[cloudconnectevent.com](http://cloudconnectevent.com)

## The Architecture Behind It All

### IT WOULD BE EASY TO IGNORE

the technologies behind cloud services, but it also would be a mistake.

Business technology pros must ensure that cloud services integrate with their enterprise infrastructures. That requires an architecture that spans both.

The components of cloud computing are the same as those in data centers: programming languages, operating systems, databases, Web servers, protocols, APIs. The task is to identify cloud services that are a good fit with your internal systems, applications, and expertise. A comparison of Amazon's Elastic Compute Cloud, Google App Engine, and Windows Azure services shows how that might work.

Amazon's EC2 lets customers pick from a software smorgasbord: Windows Server, OpenSolaris, and seven Linux flavors; the MySQL, SQL Server, and Oracle 11g databases; and the Java, JBoss, and

Ruby on Rails development environments.

Google's forte is simplicity. App Engine lets users tap into Google's homegrown database and other infrastructure software, and APIs provide access to caching, imaging, mail, and other application services. Python is the only programming language supported, though Google intends to add support for others in the future.

Windows Azure and Azure Services Platform are cut from the same cloth as Microsoft's on-premises enterprise line. Azure comprises hosted versions of SQL Server, SharePoint, Dynamics CRM, and .Net Services, and it's developed in Visual Studio and the .Net Framework. Microsoft says Azure will support open protocols (HTTP, REST, SOAP, XML) and non-Microsoft languages (Eclipse, Ruby, PHP, Python).

For IT pros who need to sketch out a cloud architecture, much of the granular information needed is available on service

providers' sites. Amazon has a white paper on cloud architectures that's worth a read for anyone trying to come up to speed.

Your blueprint should take into account the possibility of cloud services from multiple vendors, so think about how you would accomplish interoperability and application integration. Stuart Charlton, senior software architect of cloud computing startup Elastra, recommends REST and the Atom Syndication Format as underlying specifications in a global cloud architecture. Standards for federated identity management also are key, he says.

Dennis Quan, IBM Software's director of development for autonomic computing, says service-oriented architectures already make it possible to connect cloud services in "standards-complaint ways." The next trick will be to transplant services from one cloud to another. The specs to do that, Quan says, are still in their infancy. —JOHN FOLEY

Cloud stories continue on p. 24

## DATA PROTECTION

# Serious About Security

**D**EVELOPERS LOVE the deploy-and-go functionality of cloud computing, businesses like the prospect of reduced infrastructure costs, and users are happy if they get new features faster. People in charge of information security, however, are scratching their heads over how to securely move applications and data to the cloud.

A long-held goal of IT is to consolidate identity management technologies and processes; cloud computing risks setting that back a decade.

Companies could extend directory services authentication outside their environments to cope with apps and even systems in the cloud, though that approach could leave authentication systems vulnerable if the third-party systems are compromised. Or a company could implement a new solution with a separation between the cloud and existing infrastructure management. The downside is having to integrate multiple identity and access management systems. The unappealing alternative is to go back in time and manage the cloud separately.

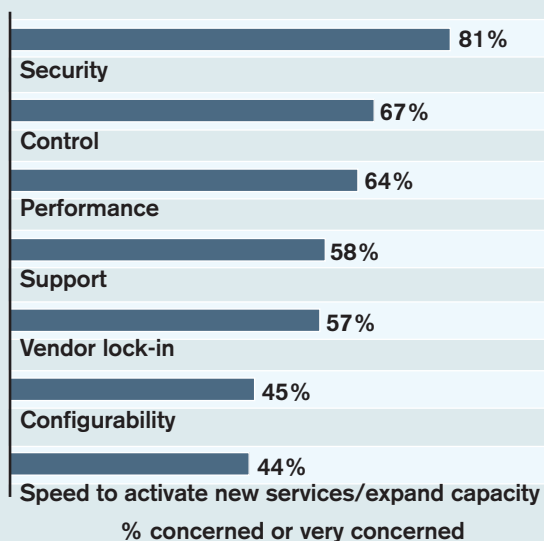
Luckily, some cloud vendors are working on the problem. Google offers the ability to tie Google Apps into existing single-sign-on implementations, increasing security and simplifying management. One company we spoke with that has a large Internet presence deployed an edge authentication server to let cloud systems authenticate via LDAP. Another extended its Web-based authentication protocols to work from external sources and authenticate to its internally hosted systems via Web services.

## DATA LOSS AND BACKUPS

Where's the data stored, who has access, and is it safe? Those are the big questions, since few cloud vendors—with the exception of a number of software-as-a-service vendors—have a long record of handling sensitive data. Unless otherwise advised, expect data to be on shared storage and potentially at risk. Truth is, we take risks with data even inside our own companies. Apply the same benefit-to-risk measure used for in-house data to the cloud, then decide what can go to the cloud and how to protect it. That requires knowing and verifying the vendor's standards and how much they can be adapted.

When using services such as Amazon's Elastic Compute Cloud, companies can apply data encryption within the operating system, application, or database

### Are you concerned with the following issues as they relate to cloud computing?



Data: InformationWeek survey of 172 business technology professionals receiving or considering cloud services

Get the latest cloud research in our Analytics Report, "A Walk In The Clouds": [cloudcomputing.informationweek.com](http://cloudcomputing.informationweek.com)

management system running in the virtual instance. Other services, such as application hosting, require more thought when developing the application to ensure that security measures such as encryption are built in.

Companies should be worried about data loss no matter where their data lies. Amazon knows computers fail, so it advises companies to plan for failure through redundancy and backups. Some cloud vendors provide backup services or ways to export data so companies can create their own backups, while others require customers to use custom or third-party applications.

Keep in mind these critical factors:

» How will backups be performed? Some cloud vendors perform backups, but more likely you'll want to conduct your own. Many customers of Amazon's EC2 also use Amazon's Simple Storage Service or Elastic Block Storage for storage of backup files.

» Can backups be tested? And if the service is down, can you access the backups?

» Where will the backup data reside? It could be on a cloud storage system, hosted by the provider, or

transferred to your own infrastructure. Regardless, you'll still need to know how data's protected when the backups are in storage and transit.

### MANAGEMENT AND MONITORING

Companies' information security teams spend time monitoring vulnerability mailing lists, patching systems, and rewriting code to fix flaws. In the cloud, they're trusting a vendor for at least some of that due diligence. Few vendors provide a way to verify their security practices, though some are becoming more forthcoming. When using cloud systems such as Joyent or Amazon's EC2, companies can apply security at the OS, database, and application layers, but they're still relying on the vendor for network, storage, and virtual infrastructure security.

While cloud customers don't control the actual patching and monitoring for vulnerabilities, they're still accountable for managing their risks. So they need to assess what needs to be protected and how to safeguard those assets, including layering on security measures around cloud infrastructure. Even then, regulations such as the Payment Card Industry (PCI) standards may throw a curve ball, since there's no clarification from the PCI council on how cloud providers are classified. That could mean they're treated slightly differently from auditor to auditor.

Customers of cloud services must demand assurances that they can monitor who has access to their data. Companies that require detailed audit trails should employ data encryption, or they should use cloud providers only for applications that interact with data that's not particularly sensitive.

This is an area that's likely to improve quickly. Google said last month that Google Apps passed a SAS 70 Type II audit of its security processes. Expect to hear more vendors touting their security standards, since security remains a big reason companies balk at moving an application to the cloud.

But internal information security teams shouldn't wait for vendors to up their security game. Cloud computing will become increasingly attractive for everything from desktop apps to server hosting. Applications requiring higher levels of security, such as HIPAA- or PCI-related apps, may be more difficult to certify in the cloud and thus be better served in-house. Community apps and content sites are better candidates. Business technology teams must decide what they're comfortable putting in the cloud. But they also must understand that the cloud ultimately will be part of the infrastructure, and it'll be up to them to figure out how to connect enterprise systems to a cloud infrastructure securely.

—ADAM ELY (aely@nwc.com)

## NETWORKING

# Bandwidth And Beyond

**W**HEN ONE REGIONAL BANK decided to move forward with a Salesforce.com rollout after a successful pilot, it skipped planning for the increased bandwidth it would need. The bank paid for that mistake when employees' Internet access suddenly slowed to a halt.

The explosion of data traveling on the Web could foreshadow a bandwidth crunch for companies that don't invest in bigger pipes. But bandwidth isn't the only potential network problem. The long distance that data travels raises latency worries, and the Internet's uncertain stability coupled with the black box of a service provider's data center make for reliability concerns.

Companies can mitigate some of these worries by upgrading their pipes. One health care company increased bandwidth fivefold to move back-end batch transaction processing into Amazon Web Services. Fortunately,

bandwidth prices continue to fall, but companies still need to plan carefully.

Technologies such as Packeteer's PacketShaper can help assess traffic flows, and most firewall vendors have 30-day free trials of reporting services that can tell companies how much bandwidth they're using. Bandwidth requirements from cloud services providers tend to be unreliable or hard to get, says Mike Healey, CTO of network integrator GreenPages, so companies should estimate bandwidth demand based at least in part on data from pilot tests. To be ready for peak demand, companies should plan for enough bandwidth so that their pipes average no more than 75% utilization, Healey says.

Redundancy is just as critical as extra bandwidth. Not planning for failover is "the biggest mistake we see clients make," says Healey. Multiple telecom companies provide last-mile

### Plug Into The Cloud

**DO IT YOURSELF** Learn how to devise a cloud computing strategy, and find daily news and analysis, at our cloud destination: [plugintothecloud.com](http://plugintothecloud.com)

Internet access in most metropolitan areas.

Also, even if a company upgrades bandwidth, it could encounter performance lags if a cloud service provider's closest data center is 3,000 miles away. "People talk about connectivity and throughput ... but latency is also a big deal, even within the cloud, because you've got distributed environments and customers talking to customers," says Glenn Dasmalchi, technical chief of staff to Cisco's CTO.

High-performance, low-latency demands come from apps such as those for calculating market risk or melding components into a composite app. They're part of the reason Amazon has built its content delivery network with data centers at points around the world, to act similarly to caching services from Akamai or LimeLight. Ask vendors what they're doing to reduce latency.

Companies needing more efficient bandwidth for cloud computing also can use load balancers. One software startup pushed most of its infrastructure—storage, processing, the developer environment—into the cloud and invested in 10 50-Mbps Verizon Fios lines, with one

Radware load balancer to aggregate bandwidth into the equivalent of one 500-Mbps line. With coming WAN optimization standards and data-intensive communication between the cloud and on-premises environments, Dasmalchi expects WAN optimization also will have a role to play in accelerating traffic among cloud computing providers, ISPs, and cloud computing users.

Cloud computing brings some new network headaches, but it also may cure others. For applications moved to the cloud, network administrators should have less work tweaking internal network architecture, since they're only providing a connection to the cloud providers' data center.

While potential cloud customers get their own networks ready for the task, they should ask cloud providers about their networks: who they use for backhaul, whether connections are redundant, where data centers are located. "Ideally, you would want to see their network design," Healey says. While the burden might be on the cloud vendor to build an adequate network, it's on the buyers to do their homework to make sure it's solid. —J. NICHOLAS HOOVER (nhoover@techweb.com)

## CONTRACTS

# Plug, Negotiate, And Play

**B**UYING CLOUD SERVICES is very different from buying packaged software when it comes to the legalese. At its most basic, just about anyone can sign up for services by filling out a few Web forms. Most companies, however, are going to want some more official license agreement tailored to their needs, and that's where things can get more complicated.

Providers typically follow one of two approaches with their cloud service and SaaS licenses—by the person, or by usage. Microsoft Exchange Online costs \$10 per user, per month, for example. Others charge per transaction or per gigabyte of data exchanged; Amazon S3 storage costs between 12 and 15 cents per gigabyte of storage and between 10 and 17 cents per gigabyte of data transfer in the United States. Some providers use a hybrid of the two approaches.

While one of the big appeals of the cloud is its scalability, companies should know what the limits are to that. Service providers often cap service levels based on what they think a customer can pay, says Ed Sullivan, CEO of Aria Systems, which provides billing services to cloud providers. For small businesses, that can limit customers to \$10,000 worth of a service per month.

With SaaS, providers increasingly are selling bundled versions that range from basic to high end. Mi-

crosoft sells a cheaper "deskless worker" version of Exchange Online that gives access to the basic version of Outlook Web Access, rather than the Outlook client. There's a free version of Google Apps, and a premium version with some business guarantees.

Like any technology, the more SaaS and cloud computing services a company buys, the more leverage it'll likely have with the contract and price. Microsoft, for example, gives discounts if buyers wrap their services into an Enterprise Agreement. Amazon charges less as customers use more of S3 and EC2. As cloud computing becomes more popular, and companies start making bigger deals, vendors are having to get flexible.

## SHARED LIABILITY NOT INCLUDED

SaaS and cloud computing carry a degree of uncertainty in terms of security, uptime, performance, and stability. With packaged software, in-house IT pros handle problems. But in the cloud, companies must rely on the service provider to minimize risks, and that needs to be spelled out in the contract, says Robert Scott, an attorney who represents both vendors structuring their licenses and business customers negotiating cloud contracts.

Standard terms often say little about many impor-

tant topics related to risk, Scott says. For example, if there's a security failure in a service that compromises financial data, a company might be required to notify customers under state or federal law, and potentially face legal action. "Who pays for that?" Scott asks.

Much of this negotiating is similar to that of an outsourcing agreement, including scrutinizing licenses with the end of the relationship in mind. Companies should make sure the terms and conditions lay out how to get data back if they decide to leave the service or can't pay for it, or if the provider suddenly shuts down. Customers need to know how they'll get data from the service provider, and how to use that data once they have access to it.

### SLAs: COMPLEXITY AND LIMITS

The service-level agreement is another piece to the puzzle. Most cloud providers give some refund if the service is down for a certain percentage of time each month, as measured by the service provider responding pings, or data requests, from the customer. Negotiating a stronger SLA will cost a pretty penny, says Warren Ross, Capgemini's global director of IT product marketing, because services get more expensive with specialized SLAs attached.

Most companies, like Salesforce.com, exclude planned downtime from SLAs, so if the vendor tells you about a planned outage, there's no refund. There tends to be a good bit of planned downtime, and it

takes away from the value of SLAs, says Divakar Jandhyala, CEO of SaaS billing and metering startup eVapt.

That said, SLAs are getting stronger and in some cases more complex, much as they did as Web hosting became mainstream. Microsoft's Exchange Hosted Filtering service includes five SLAs: for uptime, one each for anti-spam and antivirus effectiveness, one for latency, and another for performance.

Microsoft has tiered SLAs for Exchange Online and SharePoint Online. The starting point is a 99.9% availability SLA; if that's missed, customers get a 25% monthly credit. At less than 99%, customers get a 50% credit. And if there's a major outage or a virus outbreak, customers get a 100% refund for the month. But companies negotiate SLAs hoping they won't have to use them. A retailer won't be happy with a 5% refund of monthly fees if its Web site goes down on Cyber Monday.

Other elements of a cloud computing license should include a written understanding that the service provider will meet compliance demands and protect intellectual property.

Licensing cloud services is at once simpler and more complicated than using packaged software. The services are easy to buy, and many come with standard SLAs that offer a reasonable level of protection. But to ensure all areas of risk and liability are covered, keep those negotiating skills—and a lawyer's phone number—handy. —J. NICHOLAS HOOVER

## Beware The Risk Of Lock-In

**T** PROFESSIONALS ARE ALL too familiar with the consequences of addicting their organizations to proprietary programming languages, means of storing information, and other technologies.

Where open standards exist, the likelihood of costly and painful migrations down the road are somewhat mitigated. But where no or few standards exist—as is currently the case in cloud computing—the odds of getting locked in increase, as does the potential cost of switching should such a move become necessary.

Data is one of the biggest concerns. On-premises systems afford more control over how and where applications keep data. With cloud-based systems, particularly turnkey solutions, schemas

are solution-specific. Just because your data can be downloaded out of one cloud doesn't mean it will easily transfer into a competitor's platform, cloud or not.

Source code can be another problem, particularly with platforms in the cloud. Between actual code and any forms that may have been developed in the cloud, can any of it be reused elsewhere, or will a rewrite be required? When Sun's Project Caroline makes its debut, one of the expected features is a scalable cloud for running Java code. Although this speaks nothing of where the data is kept, one advantage of Java is its portability not just to on-premises solutions, but also to something in between, like a Java application server running in Amazon's Elastic Compute Cloud.

Another potential lock-in point occurs when virtualization technologies are in play. To the extent that your "systems" are supported by virtualization, it's important to realize that not all virtualization technologies are created equal.

Many providers advocate the use of virtual machines to bridge the gap between on-premises computing and cloud computing. For example, virtualize your servers locally, then move them into the cloud. But does the target cloud support your virtualization technology of choice? It's an area where the dearth of standards has given rise to specialists like rPath that help level the playing field between dissimilar platforms. —DAVID BERLIND

(dberlind@techweb.com)