PTS DATA CENTER SOLUTIONS
For More Information:
(866) 787-3271
Sales@PTSdcs.com

Prepping the Network for Cloud Computing

by Drew Robb

It's been nearly 30 years since Sun Microsystems executive John Gage came up with the slogan "The network is the computer." While some scoffed at the idea at the time, it's clear today that highly secure and robust networks are critical requirements of cloud computing.

"Realizing the vision of seamlessly integrated private/public clouds rests on a network foundation," says Charles King of IT analyst firm PUND-IT. "For certain infrastructures, WANs in particular, the result will be ever-growing demand for throughput that will strain the network and create conflicts between application/workload advocates and add to administrators' stress."

Today's cloud-based infrastructures, therefore, need networks that are capable of carrying multiple types of data, each with its own requirements. And those networks need to be more resilient than ever to ensure that users can always connect to remotely hosted services.

"Networks have to be more available, more reliable, and more performance-deterministic," says Tom Nolle, president of CIMI Corp. "You can't cloud source applications over a network if the network introduces risks that won't meet business needs once it is hosted in the cloud."

So, how does one go about evaluating how ready the network is for cloud computing and then correcting what is needed? Here are five steps to follow:

Step One: Establish the Architecture

There are a variety of different cloud architectures (private, public, hybrid, Software as a Service, Platform as a Service, Infrastructure as a Service, Security as a Service, etc.) that can be used. Without a clear understanding of the type of cloud to implement, the rest of the project will fail. The first step, therefore, is to determine which of these architectures is appropriate to meet business needs.

"While most technology transitions should occur in evolutionary steps, a vision for the ultimate technical architecture needs to be put in place early," says Gartner Vice President Thomas Bittman.

When considering whether to go with a public or private cloud, keep in mind that those terms do not refer to where the services are physically located or who is doing the hosting. Rather, it is who can access the cloud.

"A private cloud service is defined by privacy, not location, ownership or management," says Bittman. "It can be on-premises or off-, customer- or provider-managed, customer- or provider-owned."

Public cloud services, on the other hand, are offered by a service provider to all customers. In between those two are community clouds—offered to a limited number of associated customers—and hybrid clouds which contain a mix of private and public cloud functions.

"A single service request could be deployed in either implementation, or moved from one to the other, or can horizontally grow between the two implementations (also called cloud bursting, or over drafting)," says Bittman. "The primary benefit is flexibility of deployment, managed security and elasticity."

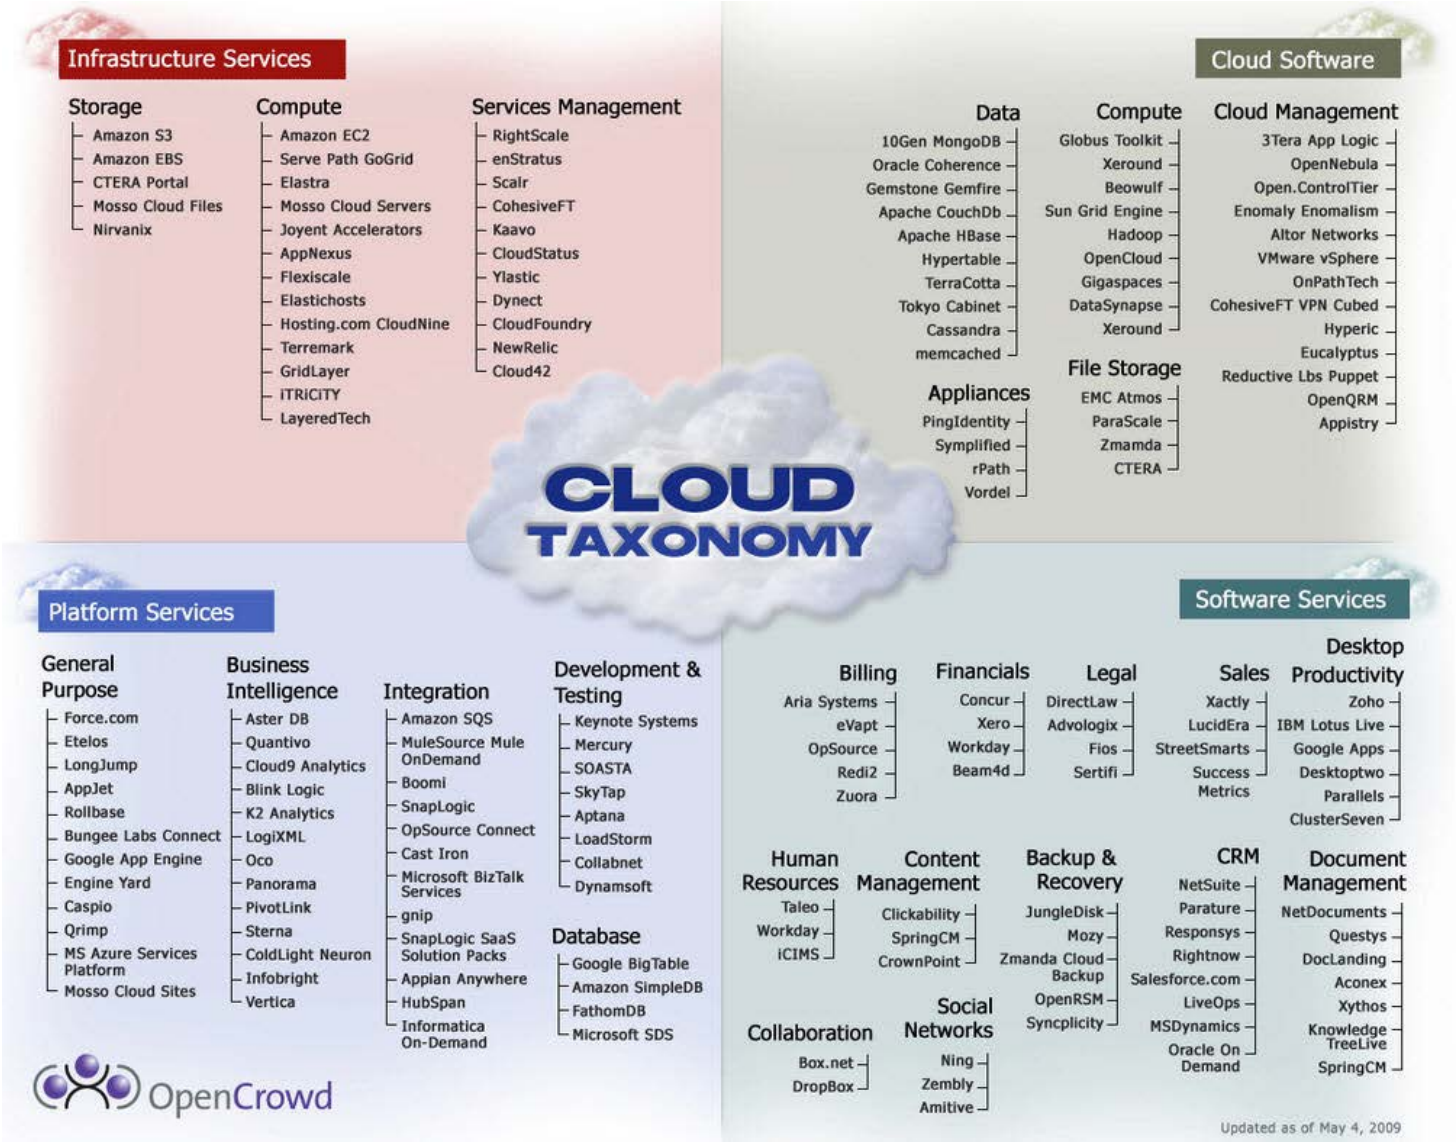

Next, look at what services are going to be provided over the cloud and what networking capabilities will be needed to support those services. Examples include:

• Software as a Service (SaaS) describes when another firm develops and hosts the software and data, and users access the applications and data through a browser or customized interface. Salesforce.com is probably the best-known SaaS provider, but other major software vendors include Oracle (Oracle CRM on Demand) Google, (Google Docs), Microsoft (Office 365) and SAP (OnDemand product line). When using SaaS, the network just has to be adequate to support an internet connection to the service provider.

• With Platform as a Service (PaaS), the vendor provides the development tools, runtime libraries and frameworks, as well as testing and hosting environments. It also allows customers to develop their own applications. Examples include K2 Analytics (an online version of the Oracle/Hyperion Solutions suite), Microsoft Azure and Amazon Relational Database Service (RDS), a hosted version of the MySQL relational database. Like SaaS, PaaS simply requires a network capable of supporting an internet connection.

• Infrastructure as a Service (IaaS) gives customers access to a virtual or physical server or rack for general computing, or a particular part of the infrastructure such as storage or content delivery. IaaS offerings include Terremark's Enterprise Cloud, Rackspace's Cloud Servers, Amazon's EC2 and AT&T's Synaptic Cloud. IaaS may

include actions such as transferring virtual machines, databases or large CAD files between local servers and the IaaS facility, and so the WAN connections must accommodate such traffic.

Note that in describing SaaS, PaaS and IaaS, the examples given were all public cloud providers. However, a central IT department may also provide those same services over a private cloud to its internal customers. When a company creates a private cloud connecting its own data centers together for failover, load balancing or other purposes, the WAN needs enough bandwidth to provide the services without bottlenecks and latency.

Financial services giant Wells Fargo has dozens of data centers with tens of thousands of servers hosting thousands of applications. While it does use Salesforce.com, its primary strategy is to develop its own private clouds.

“At Wells Fargo, we have a goal to virtualize everything possible,” says Karl Steger, VP/technology manager, Server Infrastructure Capacity & Tools Management. “With the current state of technologies in cloud and grid computing, as well as virtualization support from VMware, Citrix, etc.—and the new super-huge servers from Dell, IBM and others—we see almost no reason for dedicated physical servers.”

In selecting the architecture, keep in mind that all the details won’t need to be worked out at this point. Not only will the company’s needs change over time, but cloud technology itself is still evolving.

“Ultimately, the underlying hardware architecture will become more of a fabric, with combined compute, network and storage capabilities,” says Bittman. “The service may start as a strictly on-premises private cloud, but it may evolve over time to enable over drafting/cloud bursting to external, public cloud services.”

Step Two: Audit the Existing Network Infrastructure

The current infrastructure may be adequate to meet the SLAs required for cloud computing, it may need minor tweaks, or it may need a complete overhaul.

Even if the network doesn’t currently have issues, that doesn’t mean that it will meet the higher SLAs required by cloud computing. Once those SLAs have been established, tests and simulations can be run to determine what changes or upgrades need to be made. In doing so, be sure to look at the entire business process across all layers and locations.

“Most of the existing capacity planning and management processes are component-oriented (network, storage, compute), and run in isolation,” says Forrester Research VP and Principal Analyst Jean-Pierre Garbani. “Virtual and cloud technologies require a holistic approach that is business-service-oriented, not component limited.”

Step Three: Review Network Security

One of the major concerns with implementing cloud computing is how to maintain security and compliance when data is traveling over public networks and services are provided by outside vendors. Even when a private cloud is implemented internally, the new architecture still must prove to be secure.

“Equally important as getting the technical engineering right is getting the organizational alignment right,” says Erik Sebesta, chief architect and technology officer for Cloud Technology Partners Inc. “We are seeing

things like people putting in a private cloud, but the risk team won't allow them to deploy any mission critical applications to it."

So all of the network's access points need to be reviewed and critical information should be encrypted before exiting the firewall. If using a public or hybrid cloud, or hosting part of a private cloud externally, the security provisions at the service provider also need to be reviewed to ensure they meet requirements. On-premise infrastructure doesn't automatically mean greater security.

Step Four: Plan the Upgrade

The network upgrade should run parallel to or a step ahead of service migration to the cloud. It may just involve adding bandwidth or require a complete move to a data center Ethernet fabric. Planning a network upgrade to support cloud computing isn't much different than any other network upgrade.

"It is very similar to the steps that would be taken to support any newly remote or external computing resource, such as colocation hosting or a new data center," says Jim Frey, managing research director, Enterprise Management Associates. "Network capacity planning should be undertaken to prepare for the potential change in workloads and traffic levels, and ongoing monitoring undertaken to ensure that actual volumes and performance remain within expected limits."

Moving services offsite will add latency, so settings may need to be adjusted to ensure that processes don't time out unnecessarily. When properly managed, however, the latency will be invisible to end users. Adding bandwidth may be part of the solution, but it may also involve using WAN optimization and data or application caching to boost speeds and reduce latency. Data compression and storage deduplication can reduce the amount of data that needs to be transmitted.

"Bottom line, if you are going to do anything with the cloud, you need to keep networking services in mind as well as how to optimize, which means more than just throwing bandwidth at it," adds Frey.

Step Five: Execute the Initial Planned Upgrades

With the above preparations in place, everything should be set up for a successful transition to cloud-based computing. As with any major evolution, take a gradual approach; have monitoring tools in place to identify any bandwidth or latency issues; and be prepared to immediately roll back any unsuccessful changes until the root cause can be identified and remedied. Services with lower security and privacy concerns, and minimal business impact should come before business-critical applications.

"Clearly, the requirements to fully develop a private cloud service are pervasive and significant," says Gartner's Bittman. "However, they do not need to be done all at once, and not everything can be planned in advance."

With increased experience, additional services can be added to the cloud as appropriate. Since one of the major advantages of cloud computing is flexibility, these new cloud services won't necessarily follow the same architecture as the initial ones. By taking the first steps toward building a cloud infrastructure, even if it involves only public cloud services, a company can start reaping the benefits of higher service levels and greater agility.

“Enterprise IT should avoid being overly cautious and missing opportunities for experimentation and innovation with public cloud computing,” says Bittman. “There are likely cases in which public cloud-computing services will be ‘good enough’ for specific workload needs.”

Drew Robb is a freelance writer based in Los Angeles, CA.

PTS
DATA CENTER SOLUTIONS
For More Information:
(866) 787-3271
Sales@PTSdcs.com