# A Guide for Understanding

# Cloud Computing

***By Dr. Jim Metzler***

Sponsored by:






Produced by

# Table of Contents

## Introduction

The hyperbole-to-reality ratio that currently surrounds cloud computing is the biggest this ratio has been since the initial development of ATM. The main difference between these two ratios is that when ATM was first being discussed there were agreed to definitions as to what was meant by ATM. That is not the case with cloud computing.

This is the first of two closely-linked reports the goal of which is to reduce the hyperbole-to-reality ratio that currently surrounds cloud computing and hence enable IT organizations to rationally plan for its implementation and adoption. In order to achieve that goal, this report will:

- Describe a set of characteristics that are typically associated with a cloud computing solution.
- Define what is currently meant by private, public and hybrid cloud computing solutions and describe how those definitions are likely to change over time.
- Identify the potential cost savings and other benefits that are associated with cloud computing.
- Describe the three primary classes of public cloud computing: Software-as-a-Service, Infrastructure-as-a-Service, Platform-as-a-Service.
- Provide multiple examples of each class of public cloud computing, along with representative pricing and service level agreements.

The second report, which is entitled The Cloud Reality Report, will be available in December, 2009. The goal of *The Cloud Reality Report* is to provide a reality check on cloud computing from the perspective of an IT organization. Both of the two reports have been influenced by the quote that if often attributed to Bill Gates: “We always overestimate the change that will occur in the next two years and underestimate the change that will occur in the next ten. Don't let yourself be lulled into inaction.”

To help IT organizations not overestimate what is possible over the next two years, *The Cloud Reality Report* will focus on the current and planned implementation of cloud computing over the next two years, as well as the steps that IT organizations should take to mitigate the risk of cloud computing. *The Cloud Reality Report* is based in part on multiple surveys given to hundreds of IT professionals as well as approximately thirty

interviews. The interviewees work for IT organizations, providers of cloud computing services as well as vendors to both cloud computing providers and IT organizations.

# Characterizing Cloud Computing

## The Ambiguity over the Phrase *Cloud Computing*

There is no agreed to definition of cloud computing. For example, according to Wikipedia[1], cloud computing requires the use of the Internet to access services provided by a type of vendor that will be referred to in this report as Cloud Computing Service Providers (CCSPs).

The Wikipedia definition of cloud computing does not match up with the services currently offered by some CCSPs, nor does it match up with how some companies will access cloud computing services. For example, in June 2009 Verizon announced a set of cloud computing services[2] that can be accessed via MPLS. In at least some instances, companies will access services from Verizon as well as from other CCSPs not by using the traditional Internet, but by using a WAN service such as an Internet overlay or MPLS. Companies will do that in those instances where there is a need for low, predictable levels of delay and packet loss.

Another problem with the Wikipedia definition of cloud computing is that, as will be subsequently explained in this report, a growing number of IT organizations are implementing some or all of the techniques used by CCSPs in order to implement cloud computing within their own IT environment. This approach is often referred to as implementing private cloud computing[3]. The Wikipedia definition of cloud computing excludes this possibility.

At least one major analyst firm includes in their revenue projection for cloud computing the revenue that results from activities such as cloud-based advertising as well as some of the retail services from companies such as eBay and Expedia. Including this revenue has the impact of greatly overstating the revenue that results from the types of functionality (e.g., applications, storage, networking, computing) that are relevant to

[1] http://en.wikipedia.org/wiki/Cloud_computing

[2] http://www.computerworld.com/s/article/9133846/Verizon_Business_launches_cloud_computing_service

[3] As will be explained in the next section, there is growing disagreement as to what is meant by the phrase *private cloud computing.*

most IT organizations. In addition, when these revenue projections are quoted out of context, it tends to increase the hyperbole that surrounds cloud computing.

Further confusing the definition of cloud computing, the January 2009 edition of The Cloud Computing Journal published an article[4] that had 21 definitions of cloud computing. The position taken in this report is that the IT industry does not need yet one more definition of cloud computing. Instead of trying to define cloud computing, this report will identify the primary goal of cloud computing, will describe the key characteristics of a cloud computing solution and will demonstrate the breadth of cloud computing solutions.

## The Goal and Primary Characteristics of Cloud Computing

In spite of the confusion as to the exact definition of cloud computing, the primary goal of cloud computing is remarkably clear. That goal is to make a dramatic improvement in the cost effective, elastic provisioning of IT services.

The following set of bullets identifies the primary characteristics of existing and emerging cloud computing solutions. There is not, however, a litmus test to determine if a particular service is or is not a cloud computing service. For example, most cloud computing solutions are based on extensive virtualization of the IT infrastructure. As previously noted, Verizon recently announced a cloud computing service. As part of that service, IT organizations can gain access to a virtual server. They can alternatively gain access to a server that has not been virtualized. As such, the Verizon cloud computing service is an example of a cloud computing service that does not necessarily depend on virtualization.

The twelve primary characteristics of a cloud computing solution are:

- **Centralization** of applications, servers and storage resources. Many companies either already have, or are currently in the process of consolidating applications, servers and storage out of branch offices and into centralized data centers. This consolidation reduces cost and enables IT organizations to have better control over the company's data; e.g., who has access to the data; what security procedures are in place.

- Extensive **virtualization** of every component of IT, including servers, desktops, applications, storage, networks and appliances such as WAN optimization

[4] Twenty-One Experts Define Cloud Computing, http://cloudcomputing.sys-con.com/node/612375

controllers, application delivery controllers and firewalls. The reason that virtualization is so often associated with cloud computing is that virtualization tends to reduce cost and increase the elasticity of service provisioning. In some instances, virtualization also reduces downtime and gives the IT organization more control over IT resources.

- **Automation** of as many tasks as possible; e.g., provisioning, troubleshooting, change and configuration management. Automation can enable IT organizations to reduce cost, improve quality and increase the elasticity of service provisioning.

- The **dynamic movement of resources** such as virtual machines and the associated storage. This capability also helps to streamline the provisioning of new applications, improve backup and restoration operations and enable zero-downtime maintenance.

- Heavy reliance on the **Internet**. Virtually all companies currently make extensive use of the Internet to reduce cost, reduce the amount of time and effort it takes to deploy a new solution, and extend access to a wider range of resources and users.

- **Self-service** allows end users to select and modify their use of IT resources without the IT organization being an intermediary. This concept is often linked with usage sensitive chargeback (see below) as well as the concept of providing IT services on-demand.

- **Usage sensitive chargeback** is often referred to as pay-as-you-go. The typical rational for implementing usage sensitive chargeback is that it gives the user greater control over their IT spend because they control how much of the service they consume and at the same time enables the organization providing the service to focus on what they can control - the unit cost of the service.

- **Simplification** of the applications and services provided by IT. In a Network World article[5], Geir Ramleth the CIO of Bechtel stated that he benchmarked his organization against some Internet-based companies. According to that article, "Bechtel operates 230 applications, and it runs 3.5 versions per application. That means it maintains approximately 800 applications at any given time. 'When you

[5] The Google-ization of Bechtel, Carolyn Duffy Marsan, Network World, October 28, 2008

look at Salesforce.com, not only are they running one application, but they are running one version and they are only running it in one location,' Ramleth says."

In a simplified IT environment, the IT organization rarely develops a custom application or customizes a third party application, has just one system for functions such as ERP and SCM, and only supports one version of a given application.

- **Standardization** of the IT infrastructure. Complexity drives up cost and reduces agility and elasticity. As such, complexity is the antithesis of cloud computing. One source of complexity is having multiple suppliers of equipment such as switches and routers, as well as having multiple operating systems (i.e., Linux, Windows, Solaris), or even multiple versions of the same network operating system such as IOS.

  In a standardized IT infrastructure, the IT organization would use one primary WAN vendor and have one primary vendor of switches and routers. The IT organization would also support just one operating system and database management system and assuming that they used Cisco switches and routers, they would support just one version of IOS.

- **Technology convergence**. Cisco recently announced its Unified Computing System[6] (UCS) and HP announced its BladeSystem Matrix[7]. These systems enable the convergence of technologies such as servers, networks, storage (or storage access) and virtualization. In each case the vendor's stated rational for technology convergence is to lower the cost and improve the elasticity of the data center infrastructure.

- The development of **standards** enable, among other things, the federation of disparate cloud computing infrastructures with one another (see below). Standards could also enable functions such as the dynamic movement of virtual machines (VM) between servers with VM software from different vendors; i.e., VMWare and Microsoft.

- The **federation** of disparate cloud computing infrastructures with one another. This is accomplished by interconnecting their individual management infrastructures and it allows cloud resources to be dynamically shared.

[6] http://newsroom.cisco.com/dlls/2009/prod_031609.html
[7] http://www.hp.com/hpinfo/newsroom/press/2009/090420c.html

## Cloud Computing: A Lot of Things Old, A Few Things New

Most of the concepts that underlie cloud computing are not new. Buying computing cycles from a third party (currently referred to as Infrastructure as a Service) was first done in the 1960s and was called time-sharing. It is hard to argue that virtualized servers are a new concept since IBM was shipping virtualized mainframes thirty-five years ago. As is pointed out in a subsequent section of this report, the dot.com era spawned an early attempt at outsourcing applications from a class of provider referred to at the time as Application Service Providers (ASPs). It is certainly possible to say that from a conceptual perspective that ASPs and the current generation of Software-as-a-Service (SaaS) providers are very similar and that the only real difference between them is how they implement their services.

There are, however, some new concepts that are associated with cloud computing. One of the new concepts associated with cloud computing is the on-demand provisioning of IT resources. For example, in a traditional IT environment it can take weeks or even months to provision a new server. In a cloud computing environment, a new virtual machine can be implemented in seconds or minutes.

Another one of the new concepts associated with cloud computing is the dynamic movement of IT resources. One way that his concept is manifested is through functionality that VMware refers to as VMotion. VMotion enables the live migration of a virtual machine from one physical server to another. This functionality is intended to:

- Optimize IT resources for maximum utilization, flexibility and availability
- Allow IT organizations to perform hardware maintenance without scheduled downtime
- Improve performance by moving virtual machines away from failing or under-performing servers

As noted, it is possible to conclude that the only real difference between an ASP and a SaaS provider is how they implement their services. Those implementation differences, however, can make a huge difference in terms of the likely success or failure of each type of vendor. For example, as will be discussed in a subsequent section of this report, the IT infrastructure that the ASP model of the 1990s was built on did not differ significantly from the typical enterprise IT infrastructure of that era. As such, the typical ASP did not provide a significant cost savings vs. what an enterprise IT organization could provide on its own.

Many of the current generation of SaaS providers, however, have implemented some or all of the characteristics of cloud computing described in the preceding section of this report; i.e., centralization, virtualization, automation, simplification and standardization. If a SaaS provider has successfully implemented these characteristics it potentially gives that provider a significant cost advantage when compared to an IT organization (or a traditional ASP) that has not.

## The Theoretical Economic Value Proposition

One of the underpinnings of the theoretical economic value proposition of CCSPs such as Amazon or Google is that they are large and as such enjoy an economy of scale that most IT organizations don't. Another one of the underpinnings is that some CCSPs have invested to develop functionality that is not commercially available. An example of this which is discussed in the subsequent section is Google's development of Bigtable. A third underpinning is that CCSPs have broadly implemented the majority of the characteristics of cloud computing that were discussed in the preceding section and as a result are notably more efficient than are the vast majority of IT organizations.

There is some anecdotal evidence that indicates that at least in some instances that the theoretical value proposition of CCSPs is indeed accurate. One piece of that anecdotal evidence, the work that Bechtel did to benchmark themselves vs. Internet startups, is described in subsequent sections of this report. However, there is nothing that says that all CCSPs are so notably cost effective that even when they layer their cost of sales and their profit margin on top of the direct costs of providing the services that they offer, that they are notably more cost effective than are the vast majority of IT organizations. In addition, while most IT organizations are not as large as Amazon of Google, many are large enough to enjoy a significant economy of scale. If these large IT organizations also successfully implement some of the key characteristics of cloud computing, based on both the currently available technologies as well as those that will be deployed in the near term, then they can potentially significantly reduce, or even eliminate, the potential CCSP cost advantage.

## Classes of Cloud Computing

**Table 1** identifies and defines the various classes of cloud computing. Table 1 also highlights the fact that not only is there a lack of wide spread agreement as to the definition of cloud computing, there is also a growing lack of agreement in the industry as to the primary classes of cloud computing.

The first column of Table 1 reflects the dominant current view that there are three general classes of cloud computing: private, public and hybrid. The definition of these classes will be expanded upon below. The second column of Table 1 reflects an emerging view of cloud computing that is being advocated by a number of vendors including Cisco[8]. The primary conceptual difference between the current and the emerging views is how a private cloud is defined. In the current view, private cloud refers to IT resources provided by an IT organization. In the emerging view of a private cloud, there is management functionality that allows the enterprise IT organization to control not just the resources that it provides, but also the resources of the CCSPs. This builds on the concept of the federation of disparate cloud computing infrastructures

**Table 1: Classes of Cloud Computing**

| Current View | Emerging View | Definition |
|---|---|---|
| Private Cloud | Internal Cloud | Apply the characteristics of cloud computing solutions to the resources wholly owned by the organization consuming the services. |
| Public Cloud | External Cloud; Public Cloud | Obtain IT services from a CCSP. |
| Hybrid Cloud | Private Cloud | Provide IT services from a combination of external and internal resources but both sets of services are under the control of the IT organization that is acquiring the services. |
| Hybrid Cloud | Hybrid Cloud | Uses a combination of external and internal resources. The external resources are under the control of the CCSP and the internal resources are under the control of the IT organization that is acquiring the services. |

[8] Private Cloud Computing for Enterprises: Meet the Demands of High Utilization and Rapid Change, http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns836/ns976/white_paper_c11-543729.html

with one another that was previously discussed. Because the federation of disparate cloud computing infrastructures is difficult and with the goal of keeping this report as simple and readable as possible, this report will use the current view that there are three general classes of cloud computing: private, public and hybrid.

CCSPs that provide their services over the public Internet or other WAN services are considered to be part of the *Public Cloud*. Public cloud is what many IT organizations think of when they hear the phrase *cloud computing*. CCSPs such as Amazon or Rackspace (see the subsequent section on Infrastructure as a Service) generally have significant expertise in building and managing large-scale, virtualized data centers. For example, Google has developed and implemented a distributed storage system called Bigtable[9]. According to Google, Bigtable is designed to scale to a very large size: petabytes of data across thousands of commodity servers.

In the previously mentioned article[10], Geir Ramleth (the CIO of Bechtel) stated that he benchmarked his organization against some Internet-based companies and one of the results of that activity was that it identified an opportunity for his organization to become more notably efficient. One of his goals is to get the cost that he pays for IT functionality to be in line with the costs of industry leaders such as Amazon and YouTube. Ramleth stated that the price that Amazon charges for storage is roughly one fortieth of his internal cost for storage. Ramleth also estimated that YouTube spends between $10 and $15 per megabit/second for WAN bandwidth, while Bechtel is spending $500 per megabit/second for its Internet-based VPN. These dramatic cost savings should get the attention of every CIO and serve as an indicator that significant savings are potentially possible by either using one or more CCSPs or by the vigorous implementation of the cloud computing characteristics within the enterprise IT environment.

The decision that Ramleth reached was not that he was going to rely on third parties to supply all of his IT requirements. Rather, he decided that Bechtel would adopt cloud computing concepts within their internal IT environment that are similar to what have been deployed by companies such as Amazon and YouTube. This is what is meant today by the phrases *Private Cloud* or *Private Cloud Computing*. Private Clouds have the advantages of not being burdened by many of the potential security exposures, data confidentiality and control issues that are associated with public clouds. As will be demonstrated in *The Cloud Reality Report*, this is a very significant advantage.

[9] http://labs.google.com/papers/bigtable.html
[10] The Google-ization of Bechtel, Carolyn Duffy Marsan, Network World, October 28, 2008

Where an enterprise IT department uses a mixture of public and private cloud services, the result is referred to as a *Hybrid Cloud*. The hybrid cloud approach can offer the scalability of the public cloud coupled with the higher degree of control offered by the private cloud. In particular, a hybrid cloud might prove useful for enterprises that could benefit from offloading some of its storage requirements to a lower cost CCSP such as Amazon or some of its processing requirements to a lower cost CCSP such as Rackspace. Hybrid clouds also readily support disaster recovery solutions and provide an evolutionary path to more complete outsourcing of IT resources to the public cloud.

## Cloud Computing Services

This section will discuss the three primary categories of cloud computing services: Software as a Service, Infrastructure as a Service (IaaS) and Platform as a Service (PaaS).

### Software as a Service (SaaS)

As is the case with cloud computing in general, there is not a litmus test of characteristics that identifies conclusively if a software service offering qualifies as being a software-as-a-service offering. For example, the dot.com era spawned an early attempt at outsourcing applications from a class of provider referred to at the time as an Application Service Provider (ASP). It is easy to look at SaaS as just the ASP business model with a new name. While there certainly are similarities between the previous ASP and the current SaaS business models, there are also some significant differences.

The previous ASP model was usually based on the enterprise's custom or off-the-shelf applications being hosted at the ASP's site. In this model, the ASP was generally not an independent software vendor (ISV), but simply an operator of application servers dedicated to a particular customer. The applications were still typically enterprise-owned, or were licensed by the enterprise from an ISV. In addition, in most cases customers had the same ability to customize their applications that were hosted by an ASP as they had with the applications that they hosted themselves. Consistent with data center architectures of the time, ASPs usually dedicated a physical server to each application and/or customer instance. With this level of dedicated resources it was very difficult to achieve the economy, scalability, and flexibility that characterize any of the current virtualized data center resource models. Some of the weaknesses of the ASP model included high up-front costs and relatively limited ability to leverage the web to provide remote access to widely dispersed workforces.

In the current SaaS model, an independent software vendor (ISV) offers usage of applications as a subscription service delivered over a network. An example of this is Salesforce.com, which is a popular application that many firms access over the Internet. Salesforce.com is what most people think of when they hear the phrase *software as a service*. Typically, but not always, the service is delivered over the Internet from application servers at the ISV's data center[11]. The ISV encourages the subscriber to

[11] See the subsequent discussion of Spiceworks.

use a standard version of the software as a means of controlling operating costs and providing higher levels of security and reliability to the customer.

SaaS offers a number of advantages. For example, an application can be deployed without requiring expansion of the enterprise data center infrastructure with additional servers, storage, and networking resources. This approach can also reduce the required number of software licenses and the staffing requirements for software maintenance and updates while providing a highly predictable cost structure based on a given number of users. Because in most cases the service is accessible over the Internet, branch office and mobile workers can readily access the application.

Of the three primary classes of cloud computing services, SaaS is currently the most widely deployed and that is highly likely to be the case two years from now. According to the SaaS Showplace[12] there are an over 3,000 SaaS solutions being offered by over 700 companies. The great number of SaaS solutions stems in part from the fact that the barrier to entry is very low. While there are thousands of SaaS solutions, only a few SaaS solutions (i.e., Salesforce, WebEx) have sizeable revenues.

Any web application can be considered as a cloud application in the sense that it resides in the Internet. However, over the next two years a number of SaaS-based applications will be adapted to the platform interfaces of one or more IaaS vendors in order to become a cloud-based SaaS[13]. Using PaaS interfaces, an ISV can adapt its existing SaaS solutions to the cloud's IaaS solutions and can develop new cloud-enabled SaaS applications. By cloud-enabling the SaaS, the ISV no longer needs to dedicate part of its own data center resources to SaaS application delivery because now the application can be delivered from infrastructure provided by cloud computing service providers. Some examples of existing cloud-based SaaS applications include Google gmail and Google docs as well as IBM DB2 database software packaged as an AMI[14] for Amazon's Elastic Compute Cloud (EC2)[15].

A number of key enterprise software vendors have begun to discuss their plans for cloud computing. For example, SAP recently provided more details on their SaaS strategy[16]. As part of their strategy SAP will provide function-specific software

[12] http://www.saas-showplace.com/
[13] See the subsequent discussion of how Oracle is utilizing Amazon's EC2 platform. Also see the discussion of IaaS in general and the discussion of Force.com in particular
[14] AMI is defined in the section of this report entitled *Compute-as-a-Service*
[15] http://aws.amazon.com/ec2/
[16] SAP Unveils SaaS Strategy, http://www.informationweek.com/news/services/saas/showArticle.jhtml?articleID=217800410

applications, available by subscription, that plug into their customers' on-site SAP Business Suite systems, that SAP will host for customers using a multi-tenant architecture.

However, just as many large IT organizations are reticent to use public cloud computing services because they are concerned with loosing control, in the near term many large ISVs will make at most cautious use of IaaS vendors. Oracle, for example, has announced that it is partnering with Amazon to deploy cloud solutions that run on Amazon's EC2 platform. However, Oracle has stated[17] that it "strongly recommends that customers take time to fully understand Amazon Web Services offerings before using any of the associated Oracle products and services." Oracle further stated that, "Amazon EC2 is a fully virtualized environment and uses a virtualization engine that is not supported by Oracle. Users will therefore be directed to Amazon for any virtualization related issues." That raises the obvious question: how would a user know if a problem was related to virtualization?

## Taxonomy of SaaS Applications

This section will create a taxonomy of the most common categories of SaaS applications and will identify at least one representative product for each category[18]. *The Cloud Reality Report* will quantify the interest that IT organizations have in acquiring each of the following classes of applications from a SaaS provider.

### *Customer Relationship Management (CRM)*

As noted, Salesforce.com is the poster child for SaaS claiming to have 51,800 customers as of October 31, 2008. Salesforce.com currently offers CRM solutions tailored for sales, marketing, customer service and IT. They further customize their solutions to certain markets: small business, financial services, high technology, media, public sector and retail.

### *eCommerce*

In this context, the term eCommerce doesn't apply to the business that companies such as L. L. Bean conducts over the Internet. It applies to SaaS applications that enable eCommerce. Agent Technologies provides a system to allow customers to customize, configure and order products online via credit card or purchase orders.

[17] http://www.oracle.com/technology/tech/cloud/faq.html

[18] In many cases a particular application does not fit nicely into just one category. For example, NetSuite is listed in the ERP category. It could also have been listed in the CRM category.

## Enterprise Resourced Planning (ERP)

To be considered an ERP system, a solution must provide the function of at least two systems, such as: Product lifecycle management, Supply chain management (e.g. Purchasing, Manufacturing and Distribution), Warehouse Management, Customer Relationship Management (CRM), Sales Order Processing, Online Sales, Financials, Human Resources, and Decision Support System.

NetSuite has been in business since 1998 and announced its first profitable quarter in February 2009. The NetSuite ERP solution includes financials, order fulfillment, purchasing, inventory, time and billing, payroll, employee self-service and Web presence. As noted, SAP is an example of a traditional provider of ERP software that is evolving its products to be delivered, at least in part, in a SaaS model.

## Collaboration

Zoho provides a number of collaboration/productivity tools including web conferencing, email, word processing, spreadsheets, wikis, chat, online note taker and a presentation tool. Google is attempting to be a major player in the Web based desktop productivity market. Its current offering, Google Apps, includes gmail, calendar, docs, sites Web pages and Postini email security. Micorsoft's offering (Microsoft Office Web suite) is expected to be generally available in mid-2010.

## Supply Chain Management (SCM)

E2open provides SCM solutions for global demand and supply networks. According to E2open, they deliver a working business process including software, deployment, and operations that improves visibility, control, responsiveness, and collaboration across multiple tiers of customers, partners, and suppliers.

## Social networking

Linkedin is geared towards companies and industry professionals looking to make new business contacts or keep in touch with previous co-workers, affiliates, and clients. It is the social networking site recognized by IT professionals as having the most business relevance. However, IT organizations are beginning to make creative use of myriad social networking sites. For example, Peter Whatnell the Chief Information Officer at

Sunoco recently[19] commented on his company's use of YouTube to distribute training videos at no cost.

***Human Resources***

ThinMind provides functionality relative to timesheet tracking, expense reporting, issue tracking and project management.

***Network management***

Apparent Networks' PathView Cloud is a network performance management tool that provides insight into the complete set of end-to-end network paths an application travels. According to Apparent Networks, it enables network managers to monitor performance through the LAN, WAN, Internet and other third-party network segments regardless of who owns or controls them.

- The PathView Cloud solution is available as a free service that allows users to monitor up to five network paths of their choice, with additional paths available at ~$5/path/month.

- The service is deployed in a hybrid fashion, with a small piece of code (~6MB) installed at the point from which the user wants to perform the monitoring. All other code resides at an Apparent Networks' site, which is accessible via a browser.

Apparent Networks states that users can register for a PathView Cloud account and start testing paths in less than five minutes. They further state that the ability of the service to monitor and measure cloud service provider performance could lead to service level agreements (SLAs) becoming a standard part of cloud services contracts.

Spiceworks is a network management and monitoring, helpdesk, PC inventory and software reporting solution for small and medium businesses. Some of the interesting aspects of Spiceworks include that the solution is:

- Provided for free and that revenues are generated primarily by ad insertion.

[19] Whatnell made the comment in his keynote address at Network World's IT Roadmap conference in Philadelphia, PA on July 14, 2009.

- Supported by crowdsourcing[20] and also leverages its user base to create product and service provider ratings and to exchange best practices and make suggestions about product enhancements.

- Deployed in a hybrid fashion where some functionality resides at the customer's site behind their firewall and some functionality resides at a Spiceworks' site.

***Security***

ScanSafe provides a number of services including Web malware scanning, Web filtering and IM filtering by which all outbound and inbound IM traffic is managed for compliance, and productivity threats. In addition, traditional gateway vendors are increasingly augmenting device strategies with SaaS strategies.

The Blue Coat WebPulse cloud service utilizes Dynamic Link Analysis to identify and assess Web 2.0 threats that rely upon dynamic links to spread malicious content. The WebPulse service receives real-time intelligence on new sites, URLs and dynamic Web 2.0 content and links from a broad, diverse and expanding user community. To identify malware, phishing, and other malicious or inappropriate content within these links, the WebPulse service conducts extensive threat assessments using multiple deep-analysis techniques. Results are immediately available to all users in the community through the WebPulse cloud service without requiring downloads or other update cycles.

According to Blue Coat, more than 54 million users of BlueCoat WebFilter and ProxyClient enterprise products as well as users of K9 Web Protection, Blue Coat's free software for consumers, send more than one billion requests to the WebPulse service weekly.

## Infrastructure as a Service (IaaS)

Infrastructure services are comprised of the basic compute and storage services that are required to run applications or to provide backup. Of the three primary classes of cloud computing services, IaaS has a smaller market than does SaaS, but has a greater market than does PaaS. That relative market positioning is highly likely to still be the case two years from now.

[20] Crowdsourcing is the act of taking a task traditionally performed by an employee or contractor and outsourcing it to an undefined, generally large group of people or community in the form of an open call.

## Compute-as-a-Service

One option is that the cloud-based service could provide an extensive network of virtualized and/or physical compute devices. For example, an independent software vendor (ISV) could use an IaaS solution such as Amazon's EC2 to access a large network of virtual web servers that facilitate the development and testing of distributed applications. Alternatively, an enterprise engineering department could use an IaaS solution to gain access to a large compute cluster that consists of hundreds or even thousands of servers executing a parallel processing high-performance computing (HPC) application.

The availability of EC2 was recently brought into question. As was widely reported[21], in June 2009 a lightning strike caused damage to the power equipment at one of Amazon's data centers. The result was that some customers were out of service for more than four hours. This follows a somewhat similar outage that Amazon's EC2 service experienced in February 2008.

One of EC2's key concepts is the concept of an Amazon Machine Image (AMI)[22]. IT organizations create an AMI that contains their applications, libraries, data and associated configurations. The IT organization uploads the AMI to Amazon and chooses their instance type(s) and operating system. The factors that characterize the different instance types include memory, processing power and storage. For an instance that runs on either Linux or UNIX, the Amazon's standard pricing runs from $0.10 per hour to $0.80 per hour.

One of the services provided by The Rackspace Cloud[23] is cloud servers. According to The Rackspace Cloud:

- Each server is a fully-customizable, pay by the hour, virtualized server instance that is launched, maintained and controlled by the customer.
- Each server can be upgraded, downgraded or removed via an online control panel and an API.

The list price that The Rackspace Cloud charges for servers depends on the amount of physical memory reserved for the server instance as shown in **Table 2**.

[21] Lightning Strike Triggers Amazon EC2 Outage, http://www.datacenterknowledge.com/archives/2009/06/11/lightning-strike-triggers-amazon-ec2-outage/
[22] Amazon Elastic Compute Cloud (Amazon EC2), http://aws.amazon.com/ec2/
[23] http://www.rackspacecloud.com/

In addition to the prices contained in Table 2, there is a cost for bandwidth. That cost is $0.22 per gigabyte sent out from the server and $0.08 per gigabyte sent in to the server. Relative to support, The Rackspace Cloud provides access to forums, online resources and to group chatrooms manned by Rackspace Cloud personnel at no additional cost. For an optional $100 commitment, Rackspace Cloud provides access to a 1 x 1 chat online or to phone support.

**Table 2: Pricing for *The Rackspace Cloud* Server Offering**

| Memory (Mbytes) | Price Per Hour | Monthly Cost |
|---|---|---|
| 256 | $0.015 | $10.95 |
| 512 | $0.03 | $21.90 |
| 1024 | $0.06 | $43.80 |
| 2048 | $0.12 | $87.60 |
| 4096 | $0,24 | $175.20 |
| 8192 | $0.48 | $350.40 |
| 15872 | $0.96 | $700.80 |

Rackspace Cloud provides an SLA for their server offering. As part of that SLA they guarantee that their data center network will be available 100% of the time in any given monthly billing period, excluding scheduled maintenance. They also guarantee the functioning of all cloud server hosts including compute, storage and hypervisor. If a server host fails, they guarantee that restoration or repair will be complete within an hour of problem identification. If they fail to meet the guarantee that they provide for server hosts, they will pay a fee of 5% of the fees being paid by the user for each additional hour of downtime up to a maximum of 100% of the fees.

Amazon got a lot of attention when it first introduced its IaaS services. Since then, Amazon has generated a small amount of revenue from its IaaS services, primarily from small to medium sized ISVs. Partly in response to Amazon, larger, more established service providers who have traditionally provided managed collocation and related services are beginning to include in their service offering not just a physical server, but also VMs. In addition, other players continue to enter this market.

## Storage-as-a-Service (St-aaS)

The current storage-as-a-service industry got its start roughly a decade ago when ASPs first emerged, some of which focused on storage. Most of the early entrants into this market have either gone out of business or were acquired. The current St-aaS market is comprised of a growing number of players who approach the market somewhat differently than their predecessors did.

### *Amazon's Simple Storage Service (Amazon S3)*

According to Amazon[24], "Amazon S3 is storage for the Internet. It is designed to make web-scale computing easier for developers." Amazon S3 provides the ability to write, read, and delete objects containing from 1 byte to 5 gigabytes of data each. Each object is stored in a bucket and retrieved via a unique, developer-assigned key. A bucket can be located in the United States or in Europe. All objects within the bucket will be stored in the bucket's location, but the objects can be accessed from anywhere. Amazon provides authentication mechanisms to ensure that data is kept secure from unauthorized access. Amazon S3 relies on protocols such as REST and SOAP that were originally designed to support a services oriented architecture (SOA).

Amazon's standard storage costs in the US range from a high of $0.15 per Gigabyte/month to a low of $0.12 per Gigabyte/month. There is also a cost of data transfer. Within the US, there is a standard flat rate for all data that is transferred to Amazon. That rate is $0.100 per Gigabyte. Within the US, there is a tapered rate for data transferred out of Amazon. The standard rate for date transferred out of Amazon ranges from a high of $0.170 per Gigabyte to a low of $0.100 per Gigabyte based on the volume of data that is transmitted.

### *Akamai's NetStorage[25]*

This service is notably more sophisticated than the Amazon service. For example, as part of this service, Akamai places copies of the uploaded files at two of Akamai's data centers. The service uses the Akamai Platform to intelligently place the content close to Web users. In addition, Akamai's global traffic management functionality determines the optimal storage location from which to pull fresh content. The service also integrates with Akamai's content and application delivery services for acceleration, scale, security and availability.

[24] Amazon Simple Storage Service (Amazon S3), http://aws.amazon.com/s3/
[25] http://www.akamai.com/html/technology/products/netstorage.html

## The Nirvanix Storage Delivery Network™ (SDN)[26]

According to Nirvanis, their SDN is a fully-managed, highly secure cloud storage service. The SDN intelligently stores, delivers and processes storage requests in the best network location, providing the optimal user experience. The SDN is comprised of standards-based access to integrated services, powered by Nirvanix's patent-pending proprietary technology and infrastructure.

## Other Storage CCSPs

The storage as a service market is highly fragmented. For example, vendors such as SunGard offer a disaster recovery service. According to SunGard[27], their Virtual Server Replication services combine virtualization technology, best-in-class recovery capability, and managed production services to deliver the only disaster recovery solution with a contractually guaranteed service level agreement. In addition, companies such as Iron Mountain[28] focus on providing backup services. According to Iron Mountain their service, LiveVault®,is the complete server data backup and recovery solution for remote offices of large enterprises and small & medium-sized businesses.

## Parascale[29]

Parascale does not provide storage services but rather they sell a cloud storage solution which can be used by enterprise IT organization to create their own private storage clouds or by CCSPs to create a St-aaS cloud. Parascale's cloud storage software runs on Linux OS, the Linux XFS file system and IP networking. IT organizations can build a 6-TB storage cloud for roughly $12,000[30] assuming that the IT organization has to acquire new servers. If they can re-purpose servers that have been freed up as a result of server consolidation, then they can build a 6-TB storage cloud for under $6,000[31].

[26] http://www.nirvanix.com/
[27] http://www.availability.sungard.com/sungardsolutions/ITSolutions/managedservices/replication/Pages/replication.aspx
[28] http://www.ironmountain.com/digital/server/
[29] http://www.parascale.com/
[30] The Economics of Private Storage Clouds, John Foley, InformationWeek, February 4, 2009
[31] Ibid.

## Platform as a Service (PaaS)

According to Wikipedia[32], Platform as a Service (PaaS) is the delivery of a computing platform and solution stack as a service. The goal of PaaS is to facilitate the deployment of applications without the cost and complexity of buying and managing the underlying hardware and software layers, providing all of the facilities required to support the complete life cycle of building and delivering web applications and services entirely available from the Internet—with no software downloads or installation for developers, IT managers or end-users.

PaaS offerings include workflow facilities for application design, application development, testing, deployment and hosting as well as application services such as team collaboration, web service integration and marshalling, database integration, security, scalability, storage, persistence, state management, application versioning, application instrumentation and developer community facilitation. These services are provisioned as an integrated solution over the web.

An example of a PaaS offering is Google App Engine[33]. Google App Engine lets you run your web applications on Google's infrastructure. You can serve your app from your own domain name (such as http://www.example.com/) using Google Apps[34]. Or, you can serve your application using a free name on the appspot.com domain. IT organizations can share the application with an unlimited audience, or they can limit access to members of their organization.

The Google App Engine currently offers an API in Python[35] to Google's cloud infrastructure. As previously noted, Amazon's EC2 provides a platform interface called the Amazon Machine Image (AMI). Software vendors, such as Oracle and IBM, are offering AMI versions of their applications that can be downloaded to the EC2. These interfaces serve as a rudimentary form of free PaaS. In the future, providers of fee-based platform services will likely emerge to offer platform level abstraction that supports a wide range of popular infrastructure services. Vendors currently offering platforms for development of traditional SaaS applications can convert these offering to a PaaS. As broader PaaS offerings become available, enterprises can use a PaaS to develop customized business applications that can be run on a variety of IaaS execution services or virtual data center services.

---

[32] http://en.wikipedia.org/wiki/Platform_as_a_service
[33] http://code.google.com/appengine/
[34] http://www.google.com/apps/intl/en/business/index.html
[35] Python is a programming language that is gaining popularity with companies such as Rackspace and Google. See http://www.python.org/

Google's App Engine supports applications written in several programming languages. With App Engine's Java runtime environment, IT organizations can build an application using standard Java technologies, including the JVM, Java servlets, and the Java programming language—or any other language using a JVM-based interpreter or compiler, such as JavaScript or Ruby. App Engine also features a dedicated Python runtime environment, which includes a fast Python interpreter and the Python standard library.

App Engine costs nothing to get started. According to Google, "All applications can use up to 500 MB of storage and enough CPU and bandwidth to support an efficient app serving around 5 million page views a month, absolutely free." Once a user's requirements exceed what Google will provide for free, the resources that the application uses, such as storage and bandwidth, are measured and billed as indicated in **Table 3**.

**Table 3: App Engine Price Structure**

| Resource | Unit of Measure | Unit Cost |
|---|---|---|
| Outgoing Bandwidth | Gigabytes | $0.12 |
| Incoming Bandwidth | Gigabytes | $0.10 |
| CPU Time | CPU Hours | $0.10 |
| Stored Data | Gigabytes per Month | $0.15 |

**Figure 1** depicts the high level three-tier architecture of Salesforce.com. In this architecture, the company's SaaS offerings run over the company's force.com PaaS offerings that draw on the company's Force.com IaaS capabilities. Other SaaS providers can use the general architecture shown in Figure 1. For example, a SaaS provider could offer services based on their use of PaaS and IaaS from other providers.

The Force.com PaaS offering is based on Apex. According to Force.com[36], "Apex Code is a strongly-typed programming language that executes on the Force.com platform. Apex is used to add business logic to applications, to write database triggers, and to program controllers in the user interface layer. It has a tight integration with the database and query language, good web services support, and includes features such as futures and governors for execution in a multi-tenant environment."

**Figure 1: Salesforce.com Architecture**

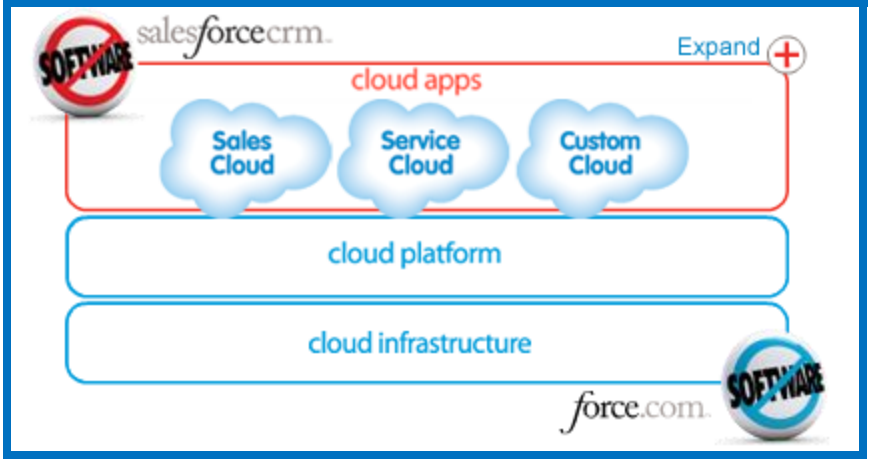


Some of the key features of the Force.com PaaS offering include:

- The ability to customize the database
- A programmable user interface
- Real-time workflows and approvals
- Real-time mobile deployment
- Granular security
- Real-time analytics

Another powerful feature of the Force.com PaaS offering is the AppExchange. The AppExchange is a set of over eight hundred applications. Roughly forty percent of the applications can be accessed for free while there is a fee associated with the remaining applications. Customers can test the application prior to purchase to see how they would work in the customer's salesforce.com environment.

[36] http://wiki.developerforce.com/index.php/Apex

## Summary and a Look Ahead

This is the first of two closely-linked reports the goal of which is to enable IT organizations to rationally plan for cloud computing. With that goal in mind, this report focused on describing and characterizing what cloud computing is. The second report, *The Cloud Reality Report,* will describe how the deployment of cloud computing is likely to play out over the next two years and what steps IT organizations should take in order to mitigate the risk of cloud computing.

It is difficult to precisely identify what is meant by cloud computing. That task is somewhat easier if the topic is public cloud computing. An IT organization can feel somewhat confident in saying that they are using cloud computing if they are using a SaaS or an IaaS provider. There is a question, however, about whether or not the SaaS and IaaS providers have implemented a significantly different technology and business model than have traditional APSs and collocation vendors. If that is not the case, then what is being referred to as public cloud computing is just a renaming of some familiar services.

In spite of the confusion as to the exact definition of cloud computing, the primary goal of cloud computing is remarkably clear. That goal is to make a dramatic improvement in the cost effective elastic provisioning of IT services. There is also a set of twelve characteristics that are typically associated with a cloud computing solution. Those characteristics are:

- Centralization of IT resources
- Virtualization of IT resources
- Automation of IT processes
- Dynamic movement of resources
- Heavy reliance on the Internet
- Self-service
- Usage-sensitive chargeback
- Simplification of IT services
- Standardization of the IT infrastructure
- Technology convergence
- Standards
- Federation of disparate cloud computing infrastructures

Unfortunately there is not a litmus test to determine if a particular service is or is not a cloud computing service based on the service incorporating some specific characteristics or even a specific number of the characteristics listed above.

Each of the twelve characteristics of cloud computing listed above potentially provides value. There are, however, significant challenges associated with each characteristic. *The Cloud Reality Report* will identify those challenges and will identify the degree to which these characteristics are currently implemented as well as how much progress IT organizations will make over the next two years; i.e., how much deployment of desktop virtualization is likely? How much deployment of technology convergence (switching, servers, storage) is likely?

Another part of the confusion that surrounds cloud computing is whether or not there is a compelling economic reason to implement public and/or private cloud computing. The Bechtel case study indicates that in some instances that CCSPs such as Amazon and Google can provide a compelling cost advantage. This cost advantage is based in part on the technologies that they have implemented and in part on the economy of scale that they enjoy due to their sheer size. While most IT organizations are not as large as Amazon or Google, many are large enough to enjoy a significant economy of scale. If these large IT organizations also successfully implement some of the key characteristics of cloud computing, then they can potentially achieve the same cost efficiencies as have CCSPs such as Amazon and Google.

There are three classes of cloud computing services: private, public and hybrid. *The Cloud Reality Report* will detail the current deployment of these classes of services and will discuss how that deployment will likely change over the next two years. The report will identify the concerns that IT organizations have with cloud computing and will also identify how the interest and the concern with cloud computing is influenced by the size of the company.

Within public cloud computing, there are three categories of services: SaaS, IaaS and PaaS. Today there is greater deployment of SaaS than there is of IaaS and there is greater deployment of IaaS than there is of PaaS. The relative size of those three markets will remain the same over the next two years. *The Cloud Reality Report* will expand upon the discussion of these three categories of public cloud computing services. For example, the report will look at various types of applications (e.g., CRM, ERP, SCM, collaboration, human resources, etc.) and identify the types of applications that IT organizations are using in a SaaS model currently and how that usage is likely to change over the next two years. *The Cloud Reality Report* will also expand upon the

use of IaaS – what are IT organizations doing today and how is that likely to change over the next two years.

# Unlocking the WAN Architecture for On-demand Cloud Computing


## Overview

Communications service providers are being called upon by their enterprise customers to help with two seemingly opposing objectives: the expansion of mission-critical WAN services to support global business objectives and cost rationalization. The Alcatel-Lucent Application-Assured (AA) VPN architecture is ideally suited to helping service providers address this requirement. The Alcatel-Lucent next-generation architecture enables service providers to deliver two application-assured business services over the same network infrastructure: VPNs and on-demand cloud computing services, such as Software as a Service (SaaS). This flexible architecture enables service providers to tailor service offerings to each enterprise's specific needs and to grow service offerings as the enterprise's needs change. Leveraging a single architecture to deliver two different services ensures the service provider gets maximum return on its infrastructure investment. Equally important, both services can be delivered using a common management system, which reduces operational overhead and ensures rapid service delivery.

## Two for the price of one

The Alcatel-Lucent AA VPN architecture enables communications service providers to:

- Deliver enhanced application-assured VPN services to enterprises with application-reporting and policy control capabilities. These value-added capabilities enable service providers to differentiate their VPN service offering.
- Deliver application-assured cloud-computing. Cloud computing is of benefit to enterprises who do not wish to take on the high capital investment of software licenses and servers, maintaining 24x7 resources, and so on. They would rather pay a monthly fee for hosted application services to the service/application provider.

Adding on-demand cloud computing into the service mix enables the service provider to strengthen its relationship with enterprise customers and evolve from a connectivity provider to a trusted ICT partner.

## Enterprise challenge equals opportunity

While enterprises are increasingly reliant on their business applications for successful day-to-day operations, they have little visibility of how applications are performing when they are being accessed over a WAN VPN. This limited visibility can be a deterrent to adopting or expanding use of cloud computing offerings.

Service providers have an opportunity to capitalize on this gap by enhancing their existing VPN and offering new on-demand cloud-computing services with the ability to monitor and address application-performance issues. The Alcatel-Lucent AA VPN architecture provides extensive service assurance capabilities, enabling the service provider to expand the reach and value of its business VPN service and to introduce a cloud computing applications portfolio.

## The Solution

The Alcatel-Lucent AA architecture enables service providers to offer application-assured VPNs and cloud computing applications. The architecture is an extension of the Alcatel-Lucent Business VPN Services (BVS) solution, which supports the convergence of IP voice, data and video over Layer 2 and/or Layer 3 business

VPNs. Enhancing the BVS solution with application assurance allows service providers to be highly competitive, leveraging their existing IP/MPLS network and service management infrastructure to offer application-level visibility and policy control with minimal incremental investment.

VPN and on-demand cloud computing services can be delivered over a single, cost-effective Alcatel-Lucent AA architecture.

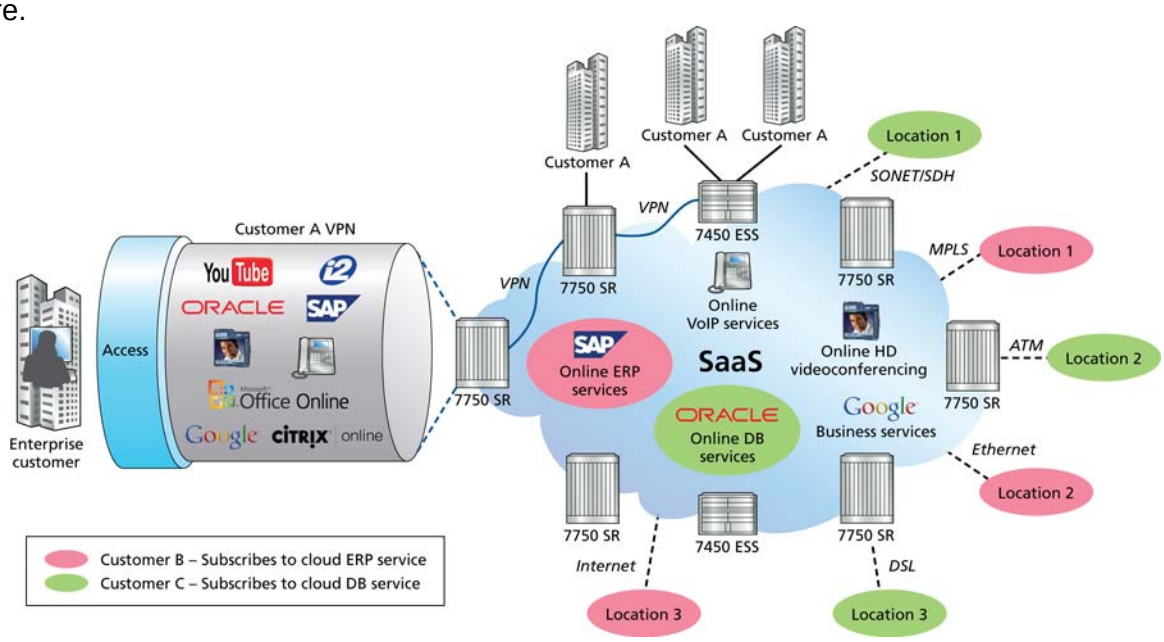


The Alcatel-Lucent AA architecture provides the ability to recognize applications and application flows throughout the network. This in turn enables the service provider to report on the applications and apply application-level QoS controls. The Alcatel-Lucent AA architecture supports multiple small, medium and large enterprise customers cost effectively, providing dramatic cost savings.

## How it works

The Alcatel-Lucent AA architecture relies on the application-assurance feature set of the Alcatel-Lucent Service Router (SR) Operating System (OS), including the purpose-built Alcatel-Lucent Multiservice Integrated Service Adapter (MS-ISA). The Alcatel-Lucent MS-ISA enables stateful application traffic-flow inspection and application assurance for IP/MPLS-based Layer 2 and Layer 3 services. By integrating the application recognition and assurance functions in the Alcatel-Lucent 7750 SR and 7450 Ethernet Service Switch (ESS) situated at the edge of the network, the number of managed network elements and overall power consumption can be dramatically reduced while the services themselves can scale to meet demand.

The Alcatel-Lucent AA architecture is fully managed with the Alcatel-Lucent 5620 Service Aware Manager suite, which includes the Alcatel-Lucent 5670 Reporting and Analysis Manager. This fully integrated network- and service-management solution simplifies operational integration and enables service providers to deliver new application reporting and assurance capabilities to their business services customers quickly.

### *Addressing the application visibility issue*

Service providers can enable web-based service portals to allow enterprises to view application reports for their VPN and cloud computing services on a highly granular basis. Enterprises can use this information to understand their service usage patterns and troubleshoot application-related issues. The solution also allows customers to manage their policies to control applications traversing the VPN, ensuring that resources are

used in accordance with business priorities. Web-based portal provides enterprises critical insight into their applications' performance.

## Reaping application assurance rewards

The Alcatel-Lucent AA architecture allows service providers to introduce a variety of tiered business VPN service plans as well as cloud computing services. Both types of service can be delivered with stringent SLAs for more direct support of enterprises' business objectives and strengthened customer relationships. These enhanced relationships provide increased opportunities for up-selling services to enterprises.

The Alcatel-Lucent AA architecture's cloud-based approach for application assurance delivers many benefits:

- Ability to identify business applications and prioritize them according to their performance behavior to ensure that an enterprise experiences consistent end-to-end application performance
- Better understanding of how an enterprise customer's applications are traversing the WAN, and deeper insight into how to baseline business-critical application traffic to improve performance
- Support for all business VPN and cloud computing applications, regardless of the access type connecting the enterprise site, allowing providers to overcome the limitation of using costly access-specific devices
- Matching of the VPN and cloud computing applications to site-specific application needs, such as volume of use by time and application
- Scalable performance, depending on the service provider's specific infrastructure. For example, a provider with Alcatel-Lucent 7750 SRs and 7450 ESSs can scale up to 70 Gb/s of application processing per node to support hundreds and thousands of enterprise sites.
- Reduced time and investment to operationalize the solution. Integrating application intelligence in the network nodes reduces complexity and eliminates the need for truck rolls to CPE locations. An integrated service management capability also minimizes installation and provisioning times.
- Ability to offer service portals to enterprise customers, allowing them to monitor application statistics, download customized application performance reports, and manage their VPNs and cloud computing applications themselves using application-based policy control.

## Conclusion

The Alcatel-Lucent cloud-based application assurance solution is deployed once by the service provider and can be offered to all enterprise customer sites at minimal cost. Moreover, the application assurance service can be remotely activated in minutes without requiring onsite installation.

These new service capabilities translate directly to increased revenue opportunities, with a faster time-to-market. By entrenching a competitive differentiator in its business service offerings, the service provider is safeguarding future revenue streams from price erosion and reinforcing customer stickiness — something that is extremely valuable in today's highly competitive market.

For the enterprise, these new service capabilities ensure that they are able to expand their mission-critical WAN application services effectively, to support global business objectives and cost rationalization.

**For more information on Alcatel-Lucent's Business VPN Services, please visit**
**http://www.alcatel-lucent.com/businessvpns**

# Effective Network Management for Cloud Computing Services

Apparent NETWORKS™

***Visibility through the cloud ensures reliable service delivery***

Cloud computing is clearly an evolving paradigm, meaning different things and involving different approaches from one organization to the next. Amid all the cloud hype and ambiguity, however, there is a common thread. It's the goal of improving IT effectiveness and lowering costs by leveraging dynamic, distributed environments.

Cloud computing in all its forms – public, private and hybrid – offers compelling benefits. It provides more flexible, efficient and less costly ways to deliver quality IT services and business applications on demand. And like all networking models, achieving optimal performance with cloud-based environments requires effective network management. But the patchwork of resources that typify cloud deployments, including virtualized machines, can present real management challenges.

## Key Cloud Management Challenge – Network Visibility

Leveraging distributed environments is a key ingredient in the cloud computing mix, but it also introduces risk. When a performance issue occurs somewhere in the complex and dynamic array of resources being used for cloud-based service delivery, it can be a nightmare for IT teams.

Those teams are responsible for delivering end-to-end performance, but they lack visibility into the end-to-end network path. Their traditional network management tools don't give them the visibility into the cloud – or the WAN or service providers' networks – that they need to monitor and manage activity properly. Without the ability to see what the problem is and where it is occurring, these teams are flying blind.

## Next-generation Management for Cloud Services

PathView Cloud™ from Apparent Networks solves this visibility problem. PathView Cloud is a network management tool that provides insight into the complete, end-to-end network paths. It enables network managers to see into and through the cloud, and through other third-party network segments regardless of who owns or controls them.

PathView Cloud's patented, light-touch technology troubleshoots network performance problems, monitors service levels, assesses network readiness for new applications, and reports on service delivery. It delivers a comprehensive view into an organization's own network as well as outsourced networks and their service levels. The result is a complete picture of the health of the end-to-end network at any point in time and over time.

When performance problems happen, PathView Cloud automatically performs a hop-by-hop analysis of the impacted path to precisely pinpoint the problem, its cause and probable fixes – whether the problem with an organization's own network – or somewhere out in the cloud.

# PathView Cloud – A Flexible, Affordable Hosted Service

PathView Cloud is an ideal solution for organizations that need effective, next-generation network management but also want to control their IT costs. It makes PathView Cloud's unique and powerful path-based network performance management capabilities available to customers in an easy-to-implement, flexible and low-cost hosted service offering.

PathView Cloud saves organizations money by enabling them to avoid the costs implementing conventional software. These include licensing fees, hardware acquisition costs, and IT staff time for maintenance. Equally appealing is the time companies can save with PathView Cloud. By leveraging the infrastructure Apparent Networks has put in place behind PathView Cloud, customers can dramatically reduce or completely eliminate the implementation phase and go right into production. Customers also get immediate access to all the newest PathView Cloud features without having to wait for delivery and installation of a software upgrade.

PathView Cloud is available now. Customers can get started quickly with a permanently free version of PathView Cloud that provides coverage for five network paths. For customers who require greater coverage, Apparent Networks offers 'pay-as-you-grow' pricing for PathView Cloud: $4.95 per path per month. More information on these offers is available at www.ApparentNetwork.com.

## How PathView Cloud Works

Traditional network management systems can 'see' only the network devices an organization owns – and gets performance information from some form of physical access, such as a probe. PathView take a very different approach. As its name indicates, PathView focuses exclusively on network paths which include the entire string of devices that application traffic uses as moves form source to destination. With this approach, PathView is able to assess network infrastructure, including cloud computing environments, from the application's perspective.


***PathView Cloud monitors and tests network paths from source to destination, through your own and third-party networks.***

PathView automatically measures and monitors key performance characteristics of device on a network path regardless of where on that path they sit. It collects network path data by sending out very small, synthetic packet trains in very fast and carefully timed bursts. Using patented technology, it analyzes signatures (distortions made to the packet trains) to generate network performance intelligence. With this information, PathView can tell IT teams whether a problem is stemming from the network or from an application, and if it is a network problem, its precise location, cause and most likely remedial action.

## PathView Cloud Features and Benefits

With PathView Cloud, IT teams and network management professionals can:

- Troubleshoot network performance problems up to 70% faster with hop-by-hop network path analyses, even through cloud environments and virtual infrastructures.
- Resolve network issues quickly and easily with actionable diagnostic information.
- Proactively monitor network performance, quality of experience (QoE) and service level agreements.
- Track carrier and service provider performance with quantifiable measurements such as bandwidth, jitter, latency and packet loss.
- Easily generate compelling reports that clearly communicate network performance and IT's business value to business and technical audiences.
- Facilitate better capacity planning with historical network usage statistics to provision optimally while avoid costly, unnecessary network upgrades.
- Conduct pre-deployment assessments of network readiness for new applications or infrastructure changes to reduce roll-out risks. Measure impact of change post-deployment.
- Deploy PathView Cloud in minutes with single-ended software requiring no agents, configuration changes or end-points.
- Integrate PathView Cloud easily with third-party network management systems via SNMP and SMTP notification.

For more information on the PathView Cloud network management tool and how it can accelerate your cloud computing initiatives and ensure their success, please visit www.ApparentNetworks.com.

## About Apparent Networks

Apparent Networks is the only IT performance management provider that delivers the end-to-end service insight required for today's cloud applications. By experiencing network performance without affecting it, the company's patented path solutions assess network readiness, monitor service levels, and diagnose problems otherwise hidden from sight. Leading companies rely on Apparent Networks to assure application delivery and expand their service portfolios with confidence. **For more information, visit www.apparentnetworks.com**

# Accelerate Secure Access to SaaS and the Cloud as Part of a Complete Application Strategy

Blue Coat

Your business is driven by the demands of centralization, mobilization and globalization. Software-as-a–service (SaaS) and cloud-based IT infrastructure are entirely new models that enable you to adapt and scale to meet these demands; but not without unique challenges. Beyond the obvious concerns of ceding control of corporate data or mission critical functions to an external organization, there are monitoring, optimization and security issues you need to be able to resolve, including:

- Performance and bandwidth issues that arise from network-centric delivery, often over Internet links where you have less control than your corporate WAN
- Ability to differentiate, prioritize and control what is business-critical from what is not
- A strong authentication and tracking of users accessing remote applications to protect against unauthorized access
- Exposure to Web-based malware, not necessarily from the SaaS and utility providers, but from the dynamically moving link threats of social networking and other approved sites
- The need to optimize and secure mobile and home workers, in addition to headquarters and remote sites

## An Application Delivery Network Helps Secure and Optimize

To deliver applications precisely when and where they are needed, your enterprise requires the critical layer of intelligence of an Application Delivery Network (ADN) to provide greater application mobility and security. An ADN helps you to:

- Accelerate SaaS performance and reduce bandwidth requirements along with files, email, video and other business applications
- Filter unauthorized sites and content at any user location
- Authenticate and monitor employee access to external SaaS applications and IT utilities
- Protect users from dynamic link-borne malware with real-time cloud security service
- Secure and optimize remote and home workers with client-based software, along with HQ, data centers and branch offices.

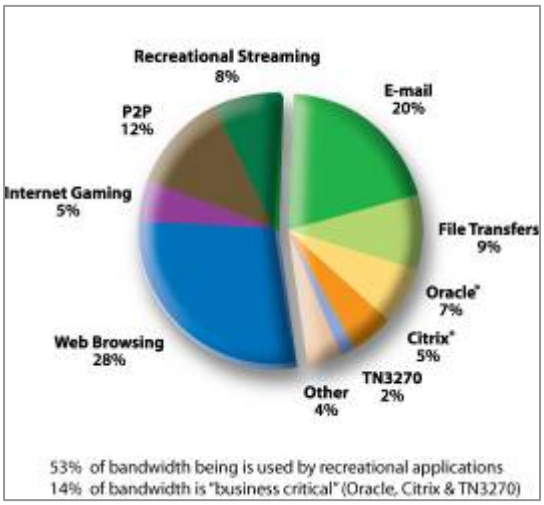


Your business depends on getting the right information to the right people on time, by assuring fast and responsive application delivery while protecting users and data from malicious threats. Are you supporting the delivery of information as reliably and safely as your business demands? Now is the time to assess your ability to monitor application performance, optimize your WAN, and secure your Web access. Your network should be providing three essential technologies, visibility, acceleration and security, to fully optimize and secure the flow of information to any user in your enterprise.

## Visibility – See and Monitor Applications & web sites


Before you can assure the timely delivery of applications you need the ability to discover and classify them. There are hundreds of applications running on a typical enterprise network and thousands of Web sites accessed by employees. They all require valuable bandwidth, and often the least important traffic is the most aggressive at dominating available bandwidth. Recreational applications such as iTunes, YouTube and peer-to-peer (P2P) are notorious for draining performance from important business applications. Having visibility into applications means the ability to distinguish and classify each one.

For business applications like Oracle and SAP you need the ability to sub-classify them to prioritize the right operations and assure reliable end-user experiences. For Web traffic, the ability to categorize, monitor access, prioritize bandwidth and block malware is key to protecting users and information.

## Monitor End-user Experiences & Web Access

Delivering acceptable user experiences requires the right technology to measure and alert on response times for SaaS, voice and other business applications. Multiple performance factors must be monitored for each session. When user experiences are jeopardized, you need the intelligence to know where the problem is, the real-time control to proactively fix it, and the ability to analyze the cause to help fine-tune ongoing delivery.

An integrated solution is more scalable and cost-effective than a disparate group of tools that operate independently. To truly manage the user experiences of each application you need to identify and resolve problems quickly. And that requires closely integrated service-level metrics that automatically monitor performance and guide resolution.

## Accelerate SaaS and Other Mission Critical Applications

Once you have visibility, you can speed up the important applications with acceleration. As with visibility, acceleration must be application aware and intelligent with the ability to apply policies to classified traffic, allowing you to enforce corporate priorities and protect information.

Without control over a SaaS vendor site, optimizing these and other cloud-based applications is a unique challenge. With Blue Coat's innovative asymmetric Web and SSL acceleration technologies, we can optimize and secure external applications over Web & SSL, delivering a better user experience, reducing bandwidth and giving IT control of access.

## Anticipate and Block Threats

New forms of malware present constantly changing threats to your network every day. Once you classify and accelerate your important applications you need to protect users and information from Web-based malware without compromising global agility.


An intelligent Secure Web Gateway solution integrates multiple layers of protection against malware, flexible policy controls and a collaborative cloud-based service to prevent intrusions and block your exposure to elusive Web threats. Enforcing corporate policies regarding recreational and Web applications are essential to maintaining productivity and reducing contention for bandwidth. And having the ability to prevent data loss by validating trusted access limits the compromise of critical business information, whether accidental or intentional.

## With Blue Coat, Control is Yours.

You face a converging set of business drivers — centralization, mobilization and globalization. Blue Coat helps you tackle them head-on as the technology leader in Application Delivery Networking. By integrating Application Performance Monitoring, WAN Optimization and Secure Web Gateway, we offer unmatched visibility into the security and performance of business applications and Web traffic across your network.

Blue Coat provides the ability to identify and classify over 600 applications, monitor and troubleshoot performance, and resolve problems before they impact users.

And Blue Coat accelerates business-critical applications, including bulk, internal, external and real-time applications with secure sessions, including SSL, to assure even remote users will experience headquarters-like performance.

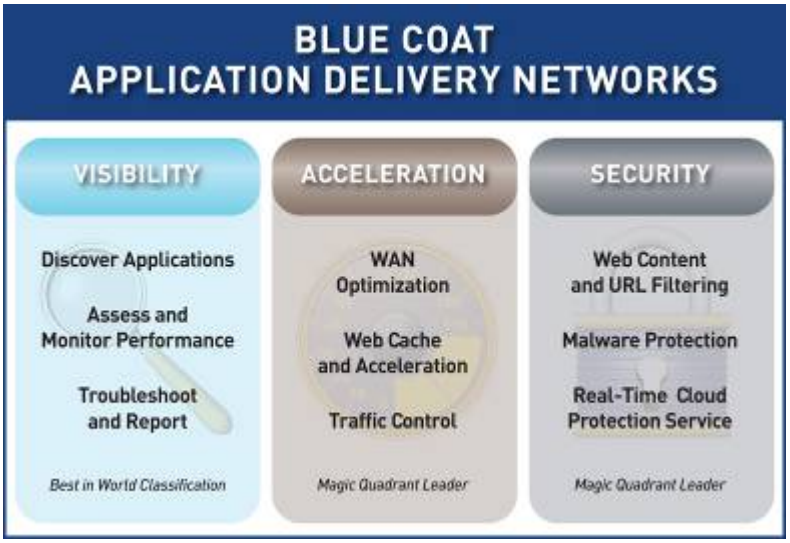


Finally, Blue Coat secures users and data with an industry-leading cloud defense and inline threat analysis for comprehensive malware scanning and filtering, centralized policy management, and in-depth reporting and logging.

With Blue Coat, you are ready to implement an Application Delivery Network with the intelligence you need to centralize, mobilize and globalize your entire IT infrastructure. Find out more at www.bluecoat.com.