Cloud Computing:

A Reality Check & Guide to Risk Mitigation

By Dr. Jim Metzler

Sponsored by:

Produced by

December 2009

Webtorials

Table of Contents

Executive Summary

Introduction and Risk Mitigation

This is the second of two closely-linked reports on the topic of cloud computing. The first report[1] identifies the value proposition, the goals and characteristics of a cloud computing solution. That report also provides an overview of existing cloud computing solutions, including some representative pricing and service level agreements. Throughout this document that report will be referred to as *The Cloud Computing Guide.*

This document identifies the current and planned use of public and private cloud computing solutions. It also identifies the risk associated with cloud computing. The purpose of identifying these risks is not to throw up roadblocks relative to the adoption of cloud computing. That follows in part because relative to cloud computing, the biggest risks accrue to those companies that don't implement any form of cloud computing. It also follows because the identification of the risks associated with cloud computing is intended to enable IT organizations to implement appropriate risk mitigation strategies so that they can implement cloud computing.

The Characterization and Interest in Cloud Computing

The three characteristics that IT organizations most associate with cloud computing are:

- Cloud computing includes SaaS
- Cloud computing involves access to services over the Internet
- As part of cloud computing, IT organizations can implement cloud computing concepts inside of their own IT organization

The first two characteristics represent the traditional view of cloud computing, which is typically referred to as public cloud computing. The third characteristic represents a view of cloud computing that has evolved more recently and is often referred to as private cloud computing.

The most popular form of public cloud computing is Software-as-a-Service (SaaS). The current SaaS marketplace consists of a few players with sizable revenues (i.e., Salesforce, WebEx) and several hundred small players. This market is unstable. As such, IT organizations should be wary of a coming shakeout in the SaaS marketplace.

While not as popular as SaaS, Infrastructure-as-a-Service (IaaS) is a large market and the number of IaaS providers is growing. It is likely in that over the next two years that

[1] http://webtorials.com/abstracts/understanding-cloud-computing.htm

there will be an ongoing segmentation of the IaaS market. As part of this segmentation, there will be low-cost, best-effort services from companies such as Amazon and Google. There will also be services that are more robust and more expensive from companies such as IBM[2]. IT organizations that are evaluating services from a particular IaaS provider will have to determine if those services are *good enough* or if a more robust service is required.

The smallest segment of the public cloud computing market is Platform-as-a-Service (PaaS). It is unlikely that there will be a very significant surge in the use of PaaS within the next two years in part because of the concerns that IT organizations have about vendor lock-in as well as the technical immaturity and the complexity of the current solutions.

The Twelve Characteristics of Cloud Computing

There are twelve characteristics that are typically associated with a cloud computing solution; e.g., virtualization, automation, simplification. However, as detailed in the corresponding section of this report, while each characteristic has distinct benefits, each also has challenges. For example, one of the challenges associated with virtual servers is that today VMWare dominates the server virtualization marketplace. However, it is very reasonable to expect that over the next two years that Microsoft will become a major player in that marketplace. It is unlikely in the near term that it will be possible to dynamically move a VM between VMWare and Microsoft environments in a seamless fashion. As such, IT organizations will either have to standardize on one vendor, or accept server islands between which it is difficult to dynamically move VMs.

Impediments to the Broad Deployment of Cloud Computing

Cloud computing faces challenges in a number of areas including: networking; security and confidentiality; management; performance. These challenges are described in detail in the corresponding sections of this report.

Networking

The movement to implement private cloud computing is causing the majority of IT organizations to re-think their data center LAN design. Many LAN vendors are saying that the traditional three-tier data center LAN architecture is no longer appropriate and that IT organizations should adopt a flat layer 2 design. Another component of the emerging design for data center LANs is the convergence of block-level storage and data traffic over a common 10 GbE data center switching fabric.

[2] http://www.networksasia.net/content/big-blue-sees-saas-hype-new-iaas-offering

The performance of the WAN has more impact on the performance of a complex cloud computing solution than it does in a more traditional computing model. IT organizations recognize this and the majority of them are looking at WAN services other than the Internet. This includes an Internet overlay service from a company such as Akamai or a traditional WAN service such as MPLS that has WAN optimization capabilities.

Security and Confidentiality

IT organizations that are interested in using a cloud computing solution must understand where their data will be hosted. In addition, before adopting a public cloud computing solution, IT organizations need to understand what impact, if any, that action will have on their compliance requirements.

Management

Before adopting a cloud computing solution, IT organizations need to do a thorough assessment of the CCSP's management capabilities. This includes identifying:

- The ability of the CCSP to troubleshoot performance or availability issues
- Whether or not the CCSP provides a meaningful SLA.

Performance

There are performance challenges associated with cloud computing relative to the centralization of IT resources, desktop virtualization and the heavy reliance on the Internet. Other performance issues that IT organizations need to address include:

- The ability of a CCSP to identify and eliminate performance issues.
- Whether or not the IT organization has any control over the performance of the service.

Summary: Developing the Cloud Computing Plan

In order to maximize the benefit of moving to a cloud computing model, IT organizations need to develop a plan (The Cloud Computing Plan) which they update on a regular basis. The Cloud Computing Plan should identify the opportunities and risks associated with both public and private cloud computing and must identify a roadmap of what steps the IT organization will take on a quarter-by-quarter basis for the next two or three years.

The Cloud Computing Plan should look systematically across multiple technologies because of the interconnected nature of the technologies. As part of creating this plan,

IT organizations need to understand the cloud computing strategy of their existing and potential suppliers, including the partnerships that they are establishing between and amongst themselves.

Introduction and Goals

This is the second 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. The first report identifies the value proposition, the goals and characteristics of a cloud computing solution. That report also provides an overview of existing cloud computing solutions, including some representative pricing and service level agreements. Throughout this document that report will be referred to as *The Cloud Computing Guide.*

The goal of this document is to provide a reality check on cloud computing from the perspective of an IT organization. Toward that end, this report will:

- Identify the current and planned use of public and private cloud computing solutions.
- Identify the issues that IT organizations need to consider before making either an initial or additional use of public or private cloud computing solutions.

One of the advantages of identifying the current and planned use of any technology or approach to providing IT services is that is enables IT organizations to see how their approach compares to that of their peers. That is important because most IT organizations do not want to be too far outside of the mainstream.

There is, however, an additional advantage of understanding the current and planned use of cloud computing. As shown in *The Cloud Computing Guide,* many of the early adopters of cloud computing will experience significant cost and agility gains. Unfortunately, as will be described in the next section, cloud computing is inherently risky. As such, some early adopters of cloud computing will experience problems, which in a few instances are likely to be severe.

One-way that IT organizations can manage the risks that are associated with cloud computing is to understand what early adopters have done to blaze a trail and reduce the risk for other IT organizations. For example, the fact that a relatively high percentage of IT organizations already use a Software-as-a-Service (SaaS) provider for collaboration and office productivity means that it is somewhat likely that at least the leading providers of those applications are making progress towards eliminating major problem areas. Analogously, the fact that relatively few IT organizations use a SaaS provider for supply chain management or enterprise resource planning means that there

has most likely been less progress towards eliminating major problem areas for those classes of applications.

This report has a two-year planning window; i.e., it will discuss the likely deployment of cloud computing and the issues that need to be resolved to enable that deployment over the next two years. This report is based in part on multiple surveys given to hundreds of IT professionals. Throughout this report those IT professionals will be referred to as The Survey Respondents. This report is also based on twenty interviews with IT professionals who will be referred to in this report as The Interviewees. The Interviewees work for IT organizations, providers of cloud computing services as well as vendors to both cloud computing providers and IT organizations.

Risk Mitigation

Over the last several years there has been a lot of discussion in industry publications about both reducing the cost of IT and about IT organizations becoming more agile. The reality is that in most cases business managers regard IT as expensive and slow moving. That view of IT began in the late 1960s with the advent of mainframe computing. At that time it could take eighteen months or longer once the need for a new application was identified, to develop and implement that application. Since that time there have been multiple waves of innovation in terms of both technology and how it is delivered that were targeted at making IT more responsive to business requirements. This includes the deployment of mini-computers and the movement to outsourcing and/or out-tasking. It also includes the deployment of PCs and the shift to implementing first client-server and then more general n-tier applications. More recently it includes the adoption of new application architectures such as a Service-Oriented Architecture.

As discussed in *The Cloud Computing Guide* the primary goal of cloud computing is a significant improvement in the cost-effect, elastic provisioning of IT services. As that document also discussed, there is evidence that at least in some instances that goal is achievable. Given that goal and that evidence, it is possible to look at cloud computing as the next major wave in giving business managers the low-cost, highly agile IT support it has always coveted.

As the rest of this report will describe, no matter what approach an IT organization take to cloud computing, there are significant risks and challenges. Possible approaches to cloud computing and a sampling of the attendant risks and challenges that will be elaborated on in this report are:

- Utilize Public Cloud Computing Services

    - Results in much less control on the part of the IT organization
    - May require that sensitive data ends up in the hands of third parties
    - Results in performance issues
    - May require that there is less customization and more overall simplification of IT service offerings
    - Results in significant management challenges
    - Creates significant interoperability issues

- Implement Private Cloud Computing

    - Requires considerable investment in emerging technologies
    - Results in significant management challenges
    - May require vendor consolidation
    - Results in somewhat less control on the part of the IT organization
    - Results in performance issues
    - May require that there is less customization and more overall simplification of IT service offerings

However, the biggest risk accrue to those companies that don't implement any form of cloud computing. IT organizations that don't implement any form of cloud computing guarantee that their company will not get the benefits described *The Cloud Computing Guide.* Also, as will be elaborated on in the next section, these IT organizations run the risk that they will get bypassed by business and functional managers that are demanding solutions that have a level of cost and agility that the IT organization cannot provide.

How much risk an IT organization is willing to take is influenced by many factors including their size, industry segment, culture and the general economic environment. The current general economic environment has created for at least some IT organizations the mantra of *doing less with less.* IT organizations that espouse that mantra state that they will not be able to withstand a significant budget cut (e.g., 20%) and still provide the same level of services as they did previously. That concept has driven at least some IT organizations to take a two-pronged approach to providing applications and services. As part of this approach, IT organizations still strive for one hundred percent availability and acceptable performance for a handful of business critical applications. However, in order to reduce cost, these same IT organizations are eliminating some of their infrastructure redundancy and disaster recovery capabilities for the rest of the applications and providing a level of service that they hope is *good enough*. One recent report[3] highlighted the trend to *do less with less* and offer services

[3] http://mmadan.wordpress.com/2009/10/23/cio-executive-board-releases-2010-it-budget-outlook/

that were *good enough* when it reported that 26% of IT organizations proactively lowered end-to-end availability commitments for mission crucial applications in 2009.

Some of the specific risks and challenges associated with cloud computing will be identified **in bold font** in a subsequent chapter (The Drivers and Use of Public Cloud Computing) of this report. Other risks and challenges will be detailed in the remaining chapters of this report.

The purpose of identifying these risks and challenges is not to imply that IT organizations should not adopt cloud computing. As previously stated, the biggest risks accrue to those IT organizations that don't adopt cloud computing. The purpose of identifying the risks and challenges associated with cloud computing is to help IT organizations develop risk mitigation plans.

The Characterization and Interest in Cloud Computing

As described in *The Cloud Computing Guide,* there is considerable confusion within the IT industry as to what is meant by the phrase *cloud computing*. The Survey Respondents were given eight possible characteristics of cloud computing and were asked to identify which of the characteristics were included in their definition of cloud computing. The three characteristics with the highest recognition were:

- Cloud computing includes SaaS
- Cloud computing involves access to services over the Internet
- As part of cloud computing, IT organizations can implement cloud computing concepts inside of their own IT organization

The responses show an interesting balance between the traditional view of cloud computing and a view of cloud computing that has emerged over the last eighteen months. The traditional view is referred to as public cloud computing and it is reflected both by the Wikipedia definition of cloud computing[4] and by the first two characteristics listed above. The emerging view is typically referred to as private cloud computing and is reflected by the third bullet listed above. Some purists believe that the phrase *private cloud computing* is an oxymoron – that cloud computing by definition implies the use of services provided by a third party. The survey responses indicate that that view is out of step with how IT organizations view cloud computing.

The survey respondents were then told that the phrase *public cloud computing* refers to IT organizations acquiring applications or services from a third party and that the phrase *private cloud computing* refers to IT organizations implementing public cloud computing

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

concepts inside of their own IT environment. They were then asked to indicate their organization's approach to both public and private cloud computing.

Just under a quarter of The Survey Respondents indicated that their organization had already started to use public cloud computing services and just over a quarter indicated that their organization had started to implement private cloud computing. In addition, roughly half of the IT organizations that were using public cloud computing services had also started to implement private cloud computing. One of the challenges associated with interpreting these results is that while most IT organizations have a good idea as to what is meant by public cloud computing, they don't always know when one or more of their company's business or functional managers are using an application that was acquired from a SaaS provider. In addition, there is a wide range in terms of how IT organizations interpret what it means to implement private cloud computing. Because of these factors, the results mentioned in this paragraph should be regarded as understating the actual use of cloud computing.

Some of the other insight gained from the responses to this question include that:

- Companies with fewer than 99 employees were more likely to be using public cloud computing services than they are to have implemented private cloud computing.
- Companies with 10,000 or more employees are more likely to have implemented private cloud computing than they are to be using public cloud computing services.
- A third of The Survey Respondents indicated that their organization would either make an initial or additional use of public and/or private cloud computing in the next year.
- A quarter of The Survey Respondents indicated that their organization sees value in cloud computing but that their belief is that the impediments outweigh the benefits.

The most important piece of insight gained from the responses to this question is that less than eight percent of The Survey Respondents indicated that their organization did not see any significant value in cloud computing. That piece of data marks the end of the discussion of 'Why should I care about cloud computing" and trumpets the entrance of the discussion of, "How do I balance the obvious benefits of cloud computing with the obvious risks?" As noted, answering that question is one of the primary goals of this report.

The Drivers and Use of Public Cloud Computing

The Survey Respondents were asked to indicate the two primary factors that are driving, or would likely drive their company to use public cloud computing services. Their responses are shown in **Figure 1**.

Figure 1: The Drivers of Public Cloud Computing

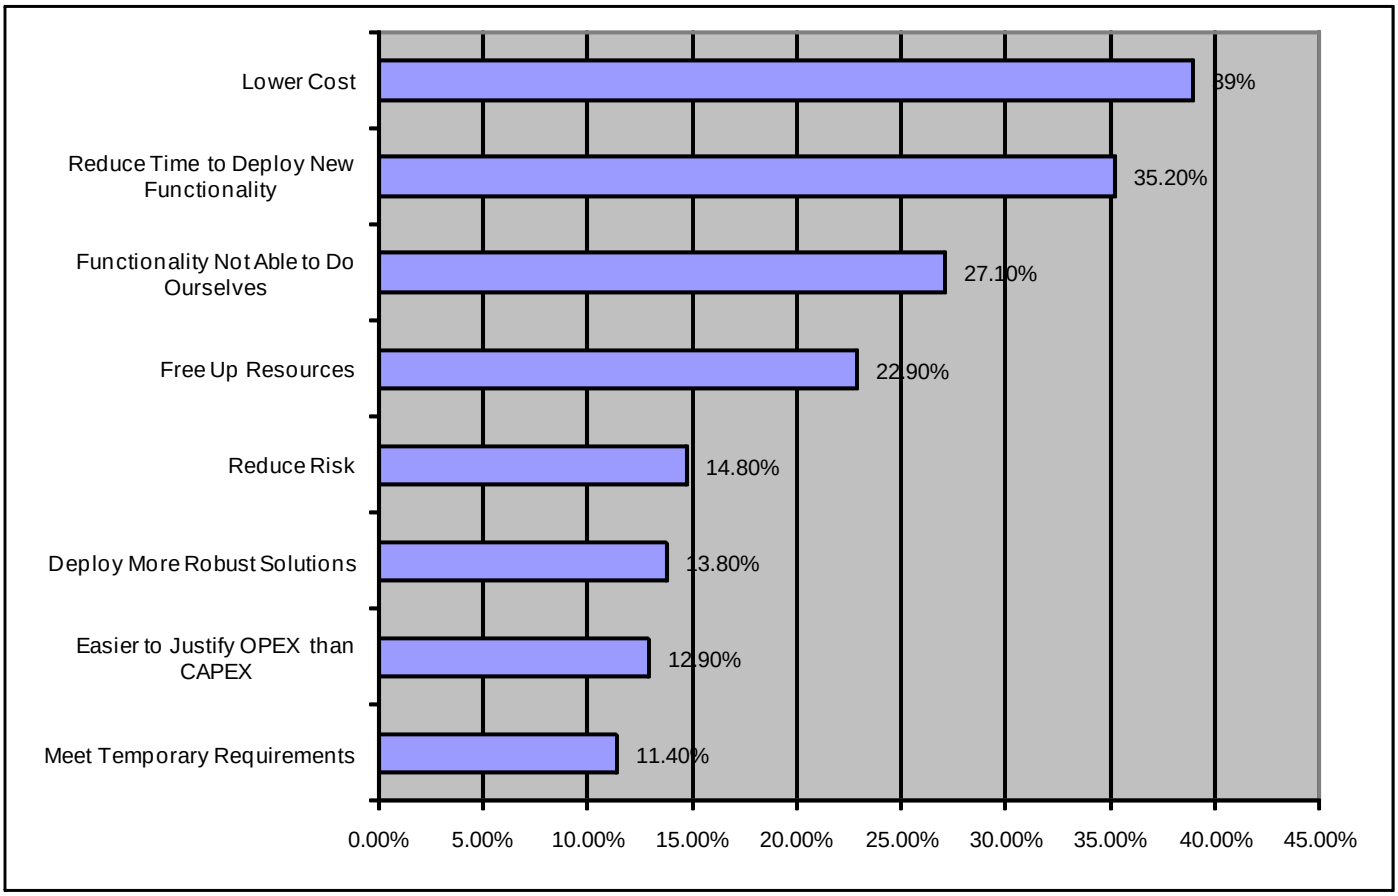

Some of the observations that can be drawn from Figure 1 include:

- The primary factors that are driving the use of public cloud computing solutions are the same factors that drive any form of out-tasking - performing an IT function faster or more cost effectively or enabling access to functionality that would otherwise be unavailable.

- Industry pundits often make the statement that it is notably easier to justify OPEX than CAPEX and use that assertion as a rational for any form of out-tasking. Some pundits go so far as to imply that IT organizations have virtually no capital

budget and a virtually unlimited expense budget. The responses of The Survey Respondents dispel this myth.

- The use of public cloud computing solutions to meet temporary requirements is relatively unimportant to The Survey Respondents. That is likely to change over the next two years.

The Interviewees stated their belief that business managers who are using public cloud services in their personal life (i.e., Google Docs[5], Quicken[6]) will put increasing pressure on the IT organization to be able to provision IT resources as dynamically and inexpensively as can cloud computing service providers (CCSPs). **IT organizations should be wary of the possibility that frustrated business and functional managers will continue to bypass the IT organization and acquire services directly from CCSPs.**

Interest in SaaS

The Survey Respondents were asked to indicate which classes of applications they either currently acquire from a SaaS provider or would likely acquire over the next year. The three classes of applications that were indicated the most often are:

- Collaboration
- Office productivity; e.g., mail, word processing, spreadsheets, etc

.

- Customer Relationship Management (CRM)

The responses listed above come as no surprise for many reasons. One reason is that two of the largest SaaS vendors are Webex and Salesforce.com, which respectively provide collaboration and CRM solutions. Another reason is the incredible interest recently shown in Google Wave[7] and IBM's recent announcement of an online email service[8]. Other reasons will be described below.

There is also no surprise that the two types of applications that IT organizations are least interested in acquiring from a SaaS provider are Enterprise Resource Planning

[5] https://www.google.com/accounts/ServiceLogin?service=writely&passive=true&nui=1&continue=http%3A%2F%2Fdocs.google.com%2F&followup=http%3A%2F%2Fdocs.google.com%2F&ltmpl=homepage&rm=false

[6] http://www.iphonestalk.com/at-3-per-month-iphone-quicken-the-tip-of-the-saas-iceberg/

[7] https://www.google.com/accounts/ServiceLogin?service=wave&passive=true&nui=1&continue=https%3A%2F%2Fwave.google.com%2Fwave%2F&followup=https%3A%2F%2Fwave.google.com%2Fwave%2F&ltmpl=standard

[8] http://www.channelinsider.com/c/a/IBM/IBM-Targets-Google-With-Online-Email-809495/?kc=TCIBESTOF10032009STR3

(ERP) and Supply Chain Management (SCM). There are significant obstacles to acquiring these types of applications from a SaaS vendor, including:

- They require significant integration with back-end systems.
- It would require putting confidential, business-critical data in the hands of a third party – something IT organizations are very wary of doing.

The Interviewees stated that in many cases today the decision to use an application from a SaaS provider is made by either someone who is not in IT or is not in an IT operational role. As a result, in many cases issues such as the cost of integrating the application with other applications, or the cost to use more of the application's functionality at some time in the future are not well analyzed. **IT organizations need to get better control over the decision to use a SaaS-provided application and to do a more thorough analysis of the issues that are associated with using the application.**

The current SaaS marketplace consists of a few players with sizable revenues (i.e., Salesforce, WebEx) and several hundred small players. This market is unstable. As such, **IT organizations should be wary of a coming shakeout in the SaaS marketplace.** Before agreeing to use an application provided by a SaaS provider, IT organizations should evaluate their financial viability. Seeing that most SaaS providers are privately held, this may mean determining if their current financial backers are big enough to be with them for the long haul and if the SaaS provider has sufficient funding, given their burn rate, to be in operation for the foreseeable future.

Interest in IaaS

As previously noted, when The Survey Respondents were asked to indicate which characteristics were included in their definition of cloud computing, the characteristic mentioned the most often was that cloud computing includes SaaS. The recognition that SaaS is a part of cloud computing was significantly stronger than the recognition that either Infrastructure-as-a-Service (IaaS) or Platform-as-a-Service (PaaS) is part of cloud computing. That is in line with the current and planned usage that The Survey Respondents are making of SaaS, IaaS and PaaS solutions. In particular, there is significantly more interest in SaaS, both today and in the near term, than there is in either IaaS or PaaS.

The Survey Respondents were asked to indicate which infrastructure services they are acquiring from an IaaS provider currently and which they expect to acquire within the next year. The following list indicates the IaaS services that The Survey Respondents are currently using. The services are mentioned in descending order of importance; i.e., compute services are the IaaS solution most frequently used by The Survey Respondents.

- Compute; e.g., use of servers and/or virtual machines (VMs) provided by a 3rd party such as Rackspace.
- Disaster recovery from a company such as Iron Mountain.
- Storage from a company such as Amazon

The Survey Respondents stated that within the next year that they expect to double their use of compute and disaster recovery services from an IaaS and triple their use of storage services.

The Survey Respondents were asked to assume that it costs 25% less to use either a server or a VM provided by a 3rd party than it does to use a server or a VM provided by their IT organization. Given that assumption, they were asked to indicate which classes of applications their IT organization would likely host on 3rd party VMs. The top three application classes, in descending order of importance were:

- Office productivity
- Collaboration
- Customer Relationship Management

While the ordering is slightly different, these are the same application classes that The Survey Respondents indicated that they were most comfortable acquiring from a SaaS provider.

The Survey Respondents were asked to assume that it costs 25% less to use storage provided by a 3rd party than it does to use storage provided by their IT organization. Given that assumption, they were asked to indicate for which classes of applications and types of data their IT organization would most likely utilize 3rd party storage. The top responses, in descending order of importance were:

- Backup data
- Training videos
- Office productivity
- Collaboration
- Social Networking

- Customer Relationship Management

One observation that can be drawn is that in response to three different questions, The Survey Respondents indicated that they were relatively comfortable with having office productivity, collaboration and CRM provided as part of a public cloud computing service. The interviewees also stated that they would feel comfortable running a compute-intensive application on infrastructure acquired from an IaaS provider, particularly if the need to run the application was either temporary or sporadic.

The Cloud Computing Guide states that 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. This approach is appealing because it has the potential to reduce the time and the cost it takes to bring new SaaS-provided applications to market. 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[9] 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."

The Oracle assertion raises some obvious questions: How would a user know if a problem was related to virtualization? How could the user prove that the problem was not related to virtualization? **IT organizations need to understand whether or not a potential SaaS provider uses an IaaS provider and if so, what that means for problem resolution.** Many industry pundits make the statement that as part of cloud computing users neither know nor care about the mechanics of how the IT services are provided. The Oracle assertion dispels that myth and indicates that **IT organizations need to understand in detail how a cloud computing service will be provided.**

While the number of IaaS providers is growing, there are far fewer IaaS providers than there are SaaS providers and the overall IaaS market is more stable than is the SaaS market. As such, while some IaaS providers could go out of business or make significant changes to their service offerings, it is unlikely that there will be a major shakeout in the IaaS market over the next two years. What is more likely in that timeframe is that there will be an ongoing segmentation of the IaaS market. One way that the IaaS market will likely segment is that at least some providers will focus on certain applications. For example, whereas simple web hosting has been a key application for most IaaS providers, Skytap focuses on software testing and development and IBM Computing on Demand focuses on high performance computing.

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

As part of the segmentation of the IaaS market there will continue to be low-cost, best-effort services from companies such as Amazon and Google. There will, however, also be services that are more robust and more expensive from companies such as IBM[10]. IT organizations that are evaluating services from a particular IaaS provider will have to determine if those services are *good enough* or if a more robust service is required.

Interest in PaaS

The Survey Respondents were asked to indicate the types of services that they are acquiring from a PaaS provider currently and which ones they expect to acquire within the next year. The following list indicates the PaaS services that The Survey Respondents are currently using. The services are mentioned in descending order of importance.

- Web application platforms such as an API to access Google's mapping, calendar and spreadsheets.
- Social Application Platforms such as using an API to Facebook
- Business application platforms such as Force.com
- Raw compute platforms such as those provided by Amazon

Some of the interviewees stated that their belief that the use of IaaS solutions on the part of ISVs would grow. That statement was based in part on their belief that the current generation of software developers feels comfortable utilizing an IaaS solution such as Amazon's EC2. As part of this approach, a software developer creates an Amazon Machine Image (AMI)[11] 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.

However, these IT interviewees also speculated that a newer generation of software developers would eschew this approach and make more use of PaaS services such as those from Force.com[12]. Part of the justification for that speculation stems from the belief that **consumerism will make an increasing impact on the IT organization – both in terms of what is expected of IT as well as how those expectations are met.** As was previously noted, business managers who are using public cloud services in their personal life will put increasing pressure on the IT organization to be able to provision IT resources as dynamically and inexpensively as can cloud CCSPs.

---

[10] http://www.networksasia.net/content/big-blue-sees-saas-hype-new-iaas-offering
[11] Amazon Elastic Compute Cloud (Amazon EC2), http://aws.amazon.com/ec2/
[12] https://www.salesforce.com/form/signup/freeforce-platform-sem.jsp?d=70130000000EgBk&DCMP=KNC-Google&gcid=3522852779&gclid=CJPZndWToZ0CFRBbagods2oN_A

Analogously, software developers who grew up with an iPhone and instant access to applications will be attracted to a PaaS vendor such as Force.com. One of Force.com's features is the AppExchange, which a couple of interviewees referred to as being analogous to iPhone's application store. The AppExchange is a set of over eight hundred applications and roughly forty percent of the applications can be accessed for free. Customers can test the application prior to purchase to see how they would work in the customer's salesforce.com environment.

The time frame for a significant surge in the use of PaaS is outside of the planning window for this report. Some of the risks associated with using PaaS include:

- Vendor Lock In: It is difficult to move from one PaaS provider to another
- Technical Immaturity: The enabling technologies will likely change significantly in the near term and this will most likely not be transparent to IT organizations who are using PaaS solutions.
- Complexity of the current generation of PaaS solutions[13]

The Twelve Characteristics of Cloud Computing

This section will list the primary characteristics of a cloud computing solution[14] and will describe the degree to which IT organizations currently implement those characteristics, and how likely that is to change over the next two years. This section will also identify the key risks and challenges associated with each of the twelve characteristics of cloud computing.

In order to maximize the benefit of moving to cloud computing, IT organizations need to develop a plan (The Cloud Computing Plan) that they update on a regular basis. A primary component of that plan is how the IT organization intends to evolve over time each of the characteristics of a cloud computing solution listed in this section; i.e., what components of IT will be virtualized, what management tasks will be automated, how much simplification of IT services will occur.

1. The **centralization** of applications, servers and storage resources[15].

The Survey Respondents were asked to indicate the percentage of their company's servers that had already been consolidated out of branch offices and into centralized

[13] http://www.ajaxwith.com/forum/4970-cloud-computing-nine-features-ideal-paas-cloud.html
[14] A more complete definition and value proposition for each of the twelve characteristics can be found in *The Cloud Computing Guide.*
[15] The term "centralization" is not intended to imply that the resources that all of the resources are in one place.

data centers as well as the percentage that they expected would be consolidated within the next year. Their responses indicated that the vast majority of IT organizations have already consolidated at least some servers out of branch offices. Their responses also indicated that only around forty percent of IT organizations have consolidated the majority of servers into centralized data centers and that percentage will increase only slightly over the next year.

Challenges and Risks:

Prior to centralization, most users accessed IT resources over a high speed, low latency LAN. Once IT resources are centralized, most users access them over a relatively low capacity, high latency WAN. As a result, the centralization of IT resources can lead to significant performance issues, particularly when a chatty application such as CIFS (Common Internet File System) is running over the WAN. IT organizations that are centralizing their IT resources should analyze the viability of implementing WAN and application optimization solutions such as those described in The 2009 Application Delivery Handbook[16].

2. Extensive **virtualization** of every component of IT, including servers, desktops, applications, storage, networks and appliances such as WAN optimization controllers, application delivery controllers and firewalls.

Server Virtualization

The data in **Table 1** shows the current and planned deployment of server virtualization as reported by The Survey Respondents. Two observations that can be drawn from **Table 1** are that within the next year:

Table 1: The Percentage of Servers that Already Have or Will be Virtualized

	None	1% to 25%	26% to 50%	51% to 75%	76% to 100%
Have already been virtualized	21.6%	33.0%	18.9%	15.1%	11.3%
Expect to be virtualized within a year	12.4%	25.6%	21.9%	21.9%	18.2%

- The number of IT organizations that have virtualized the majority of their servers will grow by sixty percent.

[16] http://webtorials.com/abstracts/2009-Application-Delivery-Handbook.htm

- The number of IT organizations that have not implemented server virtualization will be cut almost in half.

Challenges and Risks:

The Interviewees stated that to date the applications that have been placed on VMs have typically been simple applications, such as quality assurance, that are not deemed to be business critical. In addition, many IT organizations don't put applications that are security-sensitive on a VM. Both of those situations are likely to continue to be the case over the next two years.

One of the benefits of server virtualization is the ability to move VMs between physical servers – either manually or dynamically. VMWare® refers to that capability as VMotion. It is reasonable to expect that over the next two years that there will be significant progress made relative to the ability to dynamically move VMs between servers within the same data center as well as the ability to move VMs between servers in disparate data centers.

IT organizations, however, need to be aware of the fact that it is unlikely that within the next two years that it will be possible to dynamically move a VM between disparate environments, such as between VMWare and Microsoft in a completely seamless fashion. This should be a concern for IT organizations because in the current environment VMWare dominates the server virtualization marketplace[17]. However, just based on its overall market dominance, it is very reasonable to expect that over the next two years that Microsoft will become a major player in the server virtualization marketplace. As such, IT organizations have two choices. They can either standardize on one of these two vendors or accept that they will not be able to easily move VMs dynamically between a VMWare environment and a Microsoft environment.

IT organizations also need to be aware that the current form of virtualization does not scale to large applications that can't function using only a fraction of a server's resources. For that you need the ability to run a single application over multiple servers effectively and efficiently. Over the next two years, it is unlikely that much progress will be made on this form of virtualization.

Desktop Virtualization

Desktop virtualization centralizes the management of desktop applications including applications hosted at the central site (server-side virtualized applications) and applications that are streamed on-demand to client devices (client-side virtualized

[17] While VMWare dominates, Xen is used by a number of CCSPs including Amazon and 3Tera.

applications). In both of these models, the application is virtualized in the sense that it appears to be installed on the client device when that is not actually the case.

The data in **Table 2** shows the current and planned deployment of desktop virtualization as reported by The Survey Respondents. While the data in **Table 2** indicates that there is less interest in desktop virtualization than there is in server virtualization, the data also indicates that within the next year:

- The number of IT organizations that have virtualized the majority of their desktops will almost double.
- The number of IT organizations that have not implemented desktop virtualization will be cut in half.

Table 2: The Percentage of Desktops that Already Have or Will be Virtualized

	None	1% to 25%	26% to 50%	51% to 75%	76% to 100%
Have already been virtualized	49.5%	34.7%	8.9%	1.0%	5.9%
Expect to be virtualized within a year	22.0%	46.3%	18.3%	7.3%	6.1%

The Survey Respondents indicated that the following three factors, listed in descending order of importance, are the primary factors that are driving them to implement virtualized desktops:

- Reduce cost
- Increase the productivity of the IT organization
- Have more control over the company's data from a security perspective

Challenges and Risks:

The Survey Respondents were asked to indicate the primary factors that limit their organization's implementation of virtualized desktops. By a wide margin, the primary factor was their concern over the performance of the solution. The performance problems that are associated with virtualized desktops are a result of the fact that all of the data that supports the applications and/or the desktops has to transit the WAN from one or more data centers to the user. Protocols optimized for hosted application virtualization (ICA and RDP) are more efficient in their use of the WAN than are protocols such as CIFS. However, network latency can cause the

performance of these protocols to degrade. This can have a very negative impact on tasks such as a user moving his/her mouse around a computer screen.

Compared with hosted applications, streamed applications are far less efficient as they typically use the same inefficient protocols that are native to the application. Furthermore, streamed applications create additional challenges for the IT organization because of the large amount of data that must be transmitted across the WAN when the application is initially delivered to the branch. IT organizations that are implementing virtualized desktops should analyze the viability of implementing WAN and application optimization solutions such as those described in the 2009 Application Delivery Handbook.

3. **Automation** of as many tasks as possible; e.g., provisioning, troubleshooting, change and configuration management.

The Survey Respondents were asked to indicate if their organization either already had, or has plans to automate key management processes. Their responses are shown in **Figure 2**.
One of the conclusions that can be drawn from **Figure 2** is that many IT organizations are intent on automating at least some of their critical management functions.

Figure 2: The Automation of Management Processes

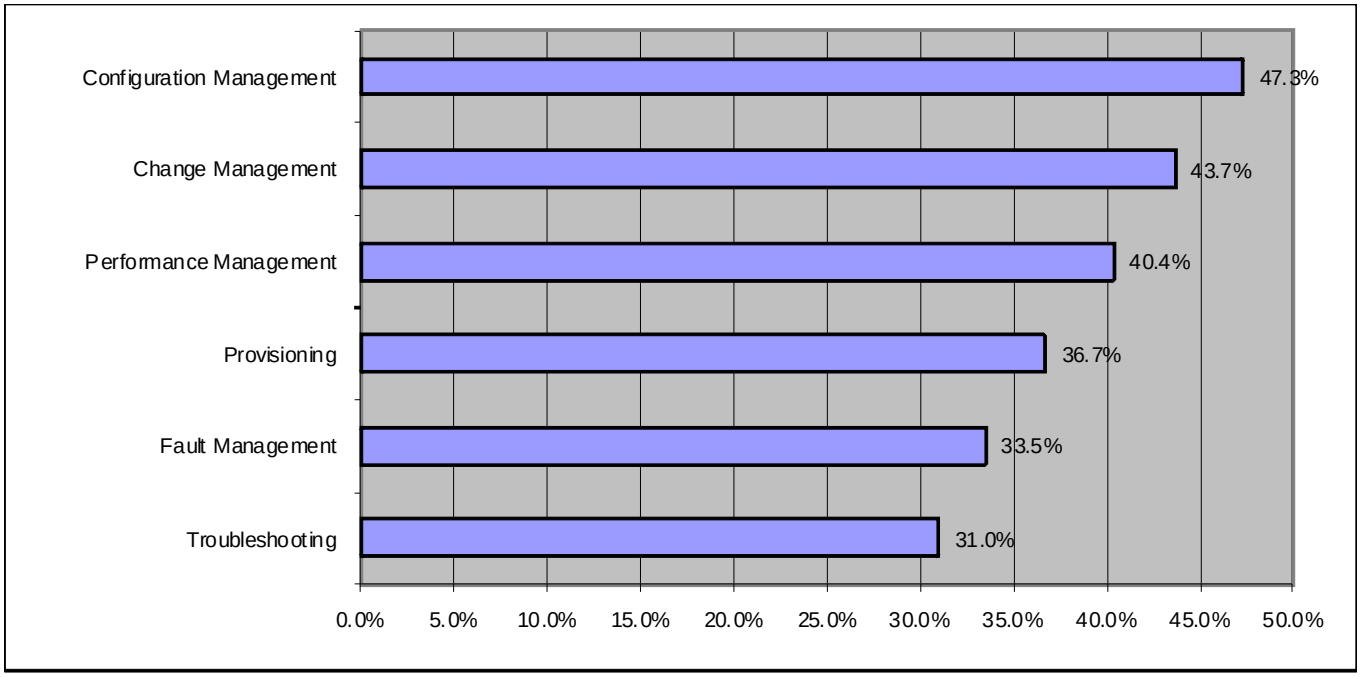

Challenges and Risks:

Automating management is challenging in part due to the sheer complexity of the task. That complexity results from the combination of the number and types of devices that require management; i.e., servers, switches, routers, intrusion detection systems, intrusion protection systems, firewalls, WAN optimization controllers and application delivery controllers. In addition, most environments are multi-vendor which further complicates the task of automation.

Automation can cause the IT organization to loose control of their IT resources. For example, as shown in Figure 2, many IT organizations either already have or soon will implement automated fault and performance management. The Interviewees, however, highlighted the fact that few IT organizations feel comfortable using tools to automatically remedy the problems identified by automated fault and performance management tools.

4. The **dynamic movement of IT resources** such as virtual machines (VMs).

As pointed out in *The Cloud Computing Guide,* the dynamic movement of IT resources is a key differentiator between what is currently referred to as cloud computing and previous waves of technology; e.g., time-sharing, virtualized mainframes and application service providers. IT organizations have shown significant interest in having the ability to dynamically move IT resources. For example, according to VMware[18], VMotion technology has been deployed in production by 70% of VMware customers. In some instances, IT organizations have installed a single VM on a physical server just because VMotion makes it easier for them to move the resources to another server.

Challenges and Risks:

Some of the obvious challenges associated with the dynamic movement of VMs is keeping track of where each VM is at each point in time, adjusting DNS so that users are directed to the appropriate VM, and avoiding VM sprawl. Another challenge is ensuring that the migrated VM retains the same security, storage access and QoS configurations and policies as it had previously. If some form of server load balancing is being performed, another challenge is to ensure that the server load balancer can account for the fact that one or more VMs have been moved. If an application firewall was protecting the VM prior to its being moved, an equivalent application firewall must be in place to protect the VM after it has been migrated. As mentioned, given the breadth of challenges associated with the dynamic movement of VMs, it is very likely that over the next two years that there will be interoperability

[18] http://www.vmware.com/products/vmotion/

issues associated with dynamically moving a VM between disparate environments; i.e., VMWare and Microsoft.

5. Heavy reliance on the **Internet**.

The Survey Respondents were asked to indicate the percentage of their company's WAN traffic that supports business critical applications that currently transits the Internet and the percentage that they expect will transit the Internet a year from now. Their responses showed a significant shift in the use of the Internet. For example, only a quarter of The Survey Respondents indicated that today the majority of their business critical applications transit the Internet. However, almost a half of The Survey Respondents indicated that in a year the majority of their business critical applications transit the Internet.

Challenges and Risks:

Virtually all companies currently make 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. There are several factors, however, that can impact the performance of the Internet. This includes the:

- Performance bottlenecks at peering points
- TCP slow start algorithm
- TCP retransmission timeout parameter
- Border Gateway Protocol's ignorance of performance when choosing a route

The poor performance of the Internet can result in unacceptable application performance. As will be discussed in the section of this report entitled cloud networking, many IT organizations will use a WAN service other than the Internet to access cloud computing services.

6. **Self-service** allows end users to access IT resources without the IT organization being an intermediary.

This concept is often linked with usage sensitive chargeback (see below) an approach that is sometimes referred to as pay-as-you-go. Virtually all CCSPs offer some level of self-service, typically by providing a portal interface that enables users to provision, de-provision or reconfigure resources such as servers, firewalls and storage. To date the typical IT organization has not aggressively deployed self-service functionality.

Challenges and Risks:

Some of the risks associated with allowing end users self-service capability with CCSP solutions include:

- Loosing control over the cost of the CCSP solutions
- Sensitive data being put in the hands of third parties
- Integration issues between the CCSP provided services and the services provided by the IT organization

If IT organizations provide self-service functionality to internal resources they potentially loose control over the use of those resources. This use of self-service can drive up cost and make it difficult to plan and manage the IT resources. However, IT organizations that don't provide self-service to internal resources run the risk of being bypassed by internal users who find that it is easier and faster to acquire IT services from a CCSP than it is to acquire services from their IT organization.

7. Usage Sensitive Chargeback - often referred to as pay-as-you-go.

The typical CCSP charges for the use of its resources on some type of usage sensitive basis. In many cases, this usage sensitive pricing is done at a granular level. For example, Rackspace[19] provides VMs and charges for them on an hourly basis. However, the pricing of many other CCSPs is notably less granular. For example, Salesforce.com charges a fix fee per user/per month for the use of its CRM solutions and Google charges a fixed fee per user/per year for usage of Google Apps.

Challenges and Risks:

Today only about 20% of IT organizations make broad use of usage sensitive chargeback, whether that is for IT resources provided directly by the organization or for IT resources acquired from a third party. Traditionally IT organizations have only implemented usage sensitive chargeback in those situations where it would change user behavior in ways that reduce cost and where it fits with the company's culture. The fit with corporate culture is critical because as is often stated, *culture eats strategy for breakfast*[20]. Throughout this report that phrase is used to drive home the concept that if the organizational culture is not aligned with any component of a

[19] www.rackspacecloud.com
[20] http://workingsmarter.typepad.com/my_weblog/2006/01/culture_eats_st.html

cloud computing strategy, that component of the cloud computing strategy will likely be less than successful.

Another impediment to the broader use of usage sensitive chargeback is that in order for an IT organization to effectively implement usage sensitive chargeback at a granular level, the organization has to help each business unit (BU) budget for their use of IT resources. The IT organization also needs the capability to bill the BU in ways that make sense to them and be able to present to them ways that they can save money. This adds a lot of overhead that can only be justified if indeed this form of chargeback significantly lowers cost.

It is highly unlikely that many IT organizations will significantly increase their use of usage sensitive chargeback, particularly at a granular level, over the next two years.

8. Simplification of the services provided by IT.

In a recent article[21], 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 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." Neither Bechtel nor any other IT organization will be able to fundamentally reduce the cost of providing IT services if it continues to offer a complex set of services.

Challenges and Risks:

Overall, seventy percent of The Survey Respondents indicated that their organization currently provides a complex set of service offerings. It was not surprising to see that as a general rule, the larger the company was, the more likely they were to be providing a complex set of services.

Adopting a simpler model of IT service offerings requires significant discussion and negotiation between an organization's CIO and the organization's key business and functional managers. Similar to the situation with usage sensitive chargeback, the success of these discussions depends in large part on the corporate culture. For example, within some organizations there are powerful business and functional managers who have become accustomed to having IT solutions customized for their needs.

As a general rule, corporate culture changes slowly. As such, it is reasonable to expect that there will be at best a very modest reduction in the complexity of the

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

services that IT organizations provide to their users over the next two years. That view is reflected by The Survey Respondents as almost sixty percent of them indicated that they believed that in two years their IT organization would still be providing a complex set of service offerings.

The interest in offering a simpler set of service offerings, however, was not uniform across the survey base. Companies with one thousand or more employees generally expressed their interest in offering a simpler set of service offering while companies that had fewer than one thousand employees often stated their interest in offering a more complex set of service offerings.

9. 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.

Overall, fifty percent of The Survey Respondents indicated that their organization currently has a complex infrastructure. It was somewhat surprising to see that companies with fewer than one thousand employees were at least as likely as larger companies were to indicate that they had a complex infrastructure.

Challenges and Risks:

Unlike simplification, the decision so standardize the IT infrastructure resides largely within the IT organization. One of the risks associated with standardization is that it can eliminate competition between vendors and hence a standardized infrastructure can potentially result in higher costs and/or reduced service. Standardization can also make it difficult for an IT organization to support third party software that requires functionality such as a non-standard server and/or operating system.

Standardization also takes time. For example, even if an IT organization decided to standardize on a given type of server, the organization would most likely not immediately swap out all of their existing servers. Rather, they would replace their existing servers on their normal replacement schedule for servers. This is also an instance where *culture eats strategy for breakfast* as the rank and file of most IT organizations could resist standardization in part because of the allegiance they feel to a given technology or vendor.

Because of these factors, it is reasonable to expect that there will be at best a modest reduction in the complexity of the typical IT infrastructure over the next two years. That view is reflected by The Survey Respondents as forty-two percent of

them indicated that they believed that in two years their IT organization would still be supporting a complex IT infrastructure. Similar to the situation with the simplification of IT services, the interest in implementing a simpler IT infrastructure was not uniform across the survey base. Companies with more than fifty thousand employees were much more likely to indicate that their IT infrastructure would be simplified in two years that were other companies.

10. Technology convergence. Cisco recently announced its Unified Computing System[22] (UCS) and HP announced its BladeSystem Matrix[23]. These systems enable the convergence of technologies such as servers, networks, storage (or storage access) and virtualization.

Challenges and Risks:

Both Cisco and HP have stated that the value proposition of this new class of product is to lower cost and improve the elasticity of the data center infrastructure – key components of a cloud computing strategy. While this class of solution is appealing, the responses of The Survey Respondents indicated that over the next two years, this class of product will only be deployed by a small set of early adopters.

Most IT organizations will avoid this class of product due to a set of general concerns about version one of any new product combined with concerns about managing the converged systems and relying too much on a single vendor. In addition, *culture eats strategy for breakfast* as many IT organizations would at least initially resist this class of product as it blurs some traditional organizational boundaries.

11. The development of **standards** to enable, among other things, the federation of disparate cloud computing infrastructures with one another (see below).

There are a number of cloud computing standards efforts currently underway. One such effort is driven by the OGF Open Cloud Computing Interface Working Group[24]. This group, which was founded in March 2009, is attempting to develop an API specification for remote management of a cloud computing infrastructure. The group's charter states that this will allow for the development of interoperable tools for common tasks including deployment, autonomic scaling and monitoring. The charter further states that the scope of the specification will be all of the high level functionality required for the life-cycle management of virtual machines (or workloads) running on virtualization technologies supporting service elasticity.

[22] http://newsroom.cisco.com/dlls/2009/prod_031609.html
[23] http://www.hp.com/hpinfo/newsroom/press/2009/090420c.html
[24] http://www.occi-wg.org/doku.php

Another important group relative to the evolution of cloud computing is the Cloud Security Alliance[25]. This is not a standards group but rather a group whose mission statement is to "Promote the use of best practices for providing security assurance within Cloud Computing, and provide education on the uses of Cloud Computing to help secure all other forms of computing." One of the documents produced by this group[26] provides security guidance for what the group considers to be critical areas of focus in cloud computing.

Challenges and Risks:

Standards development is a lengthy process [27] that typically results in a standard that represents the least common denominator in terms of the functionality on which the participants on the standards committee can agree. Over the next two years, few standards with any significant impact relative to cloud computing will be widely deployed.

It is, however, quite likely that some products will be brought to market to fill the standards vacuum. For example, it is reasonable to believe that a product that helps to bridge the gap between the VMWare environment and the Microsoft environment may well come to market sometime in the next two years. If successful, this product could become a de facto standard for interconnecting those environments.

12. The **federation** of disparate cloud computing infrastructures with one another.

Challenges and Risks:

Federation is accomplished by interconnecting the individual management infrastructures and it allows cloud resources to be shared. One way this can be accomplished is through the use of standards. However, as noted above, over the next two years few standards with any significant impact relative to cloud computing will be widely deployed.

Another way this can be accomplished is when the products that support cloud computing support multiple vendors. That does not tend to be the situation currently. For example, VMware recently announced its next generation of cloud computing operating system vsphere4[28]. Unfortunately, vsphere4 only manages VMware machines.

---

[25] http://www.cloudsecurityalliance.org/

[26] http://www.scribd.com/doc/15055401/Security-Guidance-for-Critical-Areas-of-Focus-in-Cloud-Computing

[27] The Mandate for SIP to be Interoperable, http://www.networkworld.com/newsletters/frame/2008/072108wan2.html

[28] http://searchservervirtualization.techtarget.com/news/article/0,289142,sid94_gci1354214,00.html

A third way this can be accomplished is, as discussed, through the use of a product whose enabling technology acts as a de facto standard. A fourth way is if the disparate cloud computing infrastructures have a common technology and vendor base. Over the next two years, these last two approaches are the most likely ways that federation of disparate cloud computing infrastructures will occur.

Other Impediments to the Broad Deployment of Cloud Computing

Cloud Networking

LAN

The deployment of cloud computing introduces new challenges relative to the data center LAN. Some of these challenges are driven by:

- The demands and constraints associated with the dynamic movement of VMs between physical servers. Typically port configurations are physically tied to a switch and require manual intervention to re-bind the port configuration to a new switch when a VM is moved to a new physical server. The management capabilities of at least the first generation of virtual switches (vSwitches) that resided inside of a virtualized server were rudimentary. In addition, the traditional multi-tier LAN architecture (i.e., access, distribution, core) can create bottlenecks that negatively impact the ability to dynamically move compute resources between physical servers. As a result, many LAN vendors are recommending that the data center LAN of tomorrow should have fewer tiers than the traditional data center LAN. Some vendors are suggesting that IT organizations adopt a flat, Layer 2 data center LAN.

- The need for both Storage-as-a-Service providers who provide cloud storage APIs, as well as enterprise IT organizations who want to expose their data to a range or applications, to have fast, reliable access to a range of storage media types through a range of access protocols.

- The deployment of a next generation of data warehousing and large-scale analytics processing software from companies such as Greenplum. The data center LAN that interconnects the data storage and data processing layers needs to be non-blocking and have sub-microsecond latency. It also needs to be able to scale dramatically without complex network redesign and implementation.

Another trend impacting data center LANs is the convergence of block-level storage and data traffic over a common 10 GbE data center switching fabric. The unified fabric offers significant cost savings in multiple areas including converged network adapters on

servers and reductions in rack space, power and cooling capacity, cabling, and network management overhead. In order to emulate the lossless behavior of a Fibre Channel (FC) SAN, Ethernet needs enhanced flow control mechanisms that eliminate buffer overflows for high priority traffic flows, such as storage access flows.

Lossless Ethernet will make iSCSI SANs based on 10 GbE a more attractive alternative to Fibre Channel SANs. Lossless Ethernet will also play a key role in supporting Fibre Channel over Ethernet (FCoE) technology which will allow the installed base of Fibre Channel storage devices and SANs to be accessed by Ethernet-attached servers over the unified data center switching fabric.

To quantify the degree to which IT organization recognize these challenges and opportunities and are responding to them, The Survey Respondents were asked "Has your IT organization already redesigned, or within the next year will it redesign, its data center LAN in order to support cloud computing in general, and virtualized servers in particular?" Their responses are shown in **Table 3**.

Table 3: Redesign of the Data Center LAN

	Already Have	Will Within the Next Year	No Plans
Cloud Computing in General	28.6%	42.9%	28.6%
Virtualized Servers in Particular	50.5%	30.7%	18.8%

The data in **Table 3** indicates that the vast majority of IT organizations are aware of the new challenges and opportunities relative to the data center LAN and either have already started to respond to those challenges, or intend to respond within a year.

The Survey Respondents that either already have, or soon will redesign their data center LAN, indicated that the following, in descending order of importance, are their most important design goals:

- Guaranteed performance
- Scalability
- Effective and extensible management
- Self-healing resilience
- Low latency

WAN

Figure 3 depicts the typical n-tier application architecture that is commonly used by most IT organizations. Considering the three key impairments (delay, jitter, packet loss) that impact application performance, these impairments are far more likely to occur within the WAN than they are to occur in any other component of an n-tier application. As such, the performance of the WAN is the primary determinant of the overall performance of an n-tier application.

There is nothing about cloud computing that fundamentally changes the information flow in an n-tier application. For example, in a simple public or private cloud computing solution, the end user's traffic transits a LAN and a WAN before reaching the relevant data center. At the data center, the application delivery controller (ADC) may or may not be virtualized. The Web, application and database servers would most likely be running on VMs and would potentially utilize virtualization software from different providers. In a growing number of instances, the desktops will be virtualized.

In a more complex cloud computing solution, the end user accesses a broad set of services that are provided internally at one or more of their company's data centers as well as from one or more cloud computing service providers. Whereas in a traditional n-tier application the WAN comes into play just once (connecting the enterprise LAN to the ADC), in a complex cloud computing solution, the WAN comes into play multiple times – supporting traffic between and amongst the users, the enterprise's data

Figure 3: The Typical n-Tier Application

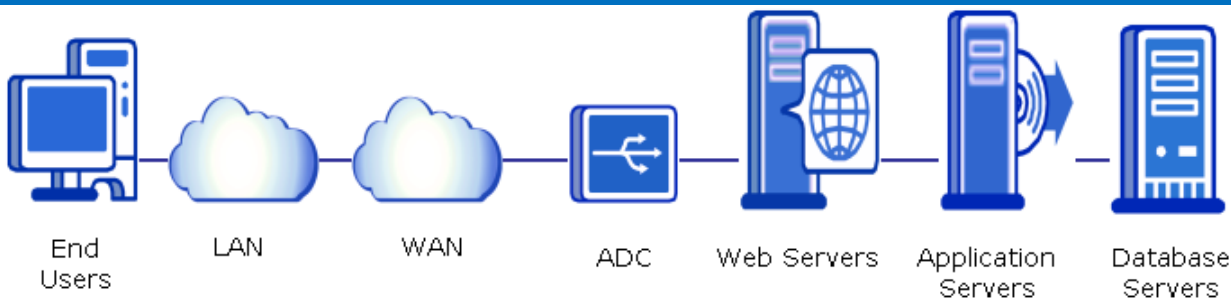

center(s) and the data centers of the cloud computing service providers. Because of this, the performance of the WAN has more impact on the performance of a complex cloud computing solution than it does in a more traditional computing model.

In addition to representing the information flow in an n-tier application, **Figure 3** can also be viewed as depicting end users accessing a public cloud computing service from a CCSP such as Amazon using a Network Service Provider (NSP) such as Verizon. Looked at this way, Figure 3 highlights the three technology and management domains that comprise a public cloud computing services. Those domains being the enterprise, the NSP and the CCSP. The distinction, however, between an NSP and a CCSP can

become fuzzy. For example, in addition to being an NSP, Verizon is also a CCSP[29]. As such, it is possible that end users would use a WAN service provided by Verizon acting as an NSP to access an IaaS solution from Verizon acting as a CCSP.

Due to a variety of well-known issues (e.g., packet loss at peering points, BGP's inability to choose the path with the lowest delay, the TCP slow start algorithm) the Internet often exhibits performance problems. As such, the Internet is not always the most appropriate WAN service to use to access cloud computing solutions. The Survey Respondents were asked to indicate which WAN service they would most likely use when accessing public and private cloud computing services over the next year. Their responses are shown in **Table 4**.

Table 4: WAN Services to Access Cloud Computing Services

	The Internet	An Internet overlay from a company such as Akamai	A traditional WAN service such as MPLS	WAN Optimization combined with a traditional WAN service; e.g. MPLS
Public Cloud Computing Services	57.1%	9.0%	19.5%	14.3%
Private Cloud Computing Services	28.8%	4.0%	30.4%	36.8%

The data in **Table 4** indicates that IT organizations understand the limitations of the Internet relative to accessing cloud computing services. Some of the other conclusions that can be drawn from **Table 4** include:

- When accessing public cloud services, IT organizations are only somewhat more likely to use the Internet than they are to use other WAN services.
- When accessing private cloud services, IT organizations are notably more likely to use a WAN service other than the Internet than they are to use the Internet.

Security and Confidentiality

The Survey Respondents were asked to indicate which of a dozen potential issues concerned them the most relative to using public cloud computing services. By a very wide margin, the top two concerns were the security of their data and the confidentiality

[29] http://telephonyonline.com/business_services/news/verizon-cloud-computing-0603/

of their data. However, as pointed out by one of The Interviewees, the concern that The Survey Respondents expressed over the security and confidentiality of their data is somewhat in conflict with the fact that many organizations use Salesforce.com and hence already place confidential sales data in the hands of a CCSP. Another interviewee expressed his belief that IT professionals who do their personal shopping online, and hence put their credit card data in the hands of a third party, will quickly adapt and come to trust the security procedures of at least some CCSPs.

IT organizations that are interested in using a cloud computing solution must understand where their data will be hosted. If their data is hosted domestically, IT organizations should verify that their data will be hosted close to their users in order to minimize the impact of delay. If their data can or will be stored internationally, IT organizations should realize that some governments have a habit of snooping on international traffic, while some other governments (primarily in the European Union) have very strict data privacy rules. In all cases storing data internationally will tend to increase network delay and can lead to poor overall performance. This is another example of why the blithe statement made by many industry pundits that as part of cloud computing users neither know nor care about the mechanics of how the IT services are provided, is wrong. A very significant challenge, however, is validating where the IT organization's data actually is stored once it is given to a CCSP.

As a minimum, before adopting a public cloud computing solution, IT organizations need to understand what impact, if any, that action will have on their compliance requirements. In particular, the fact that an IT organization is using a CCSP provided solution does not mitigate the obligation of the IT organization to show their compliance with myriad regulations. Some of the regulations that were of concern to The Survey Respondents included the PCI Data Security Standard (PCI DSS), HIPAA, Office of Foreign Asset Control[30] Compliance and the Family Educational Rights and Privacy Act (FERPA)[31].

Before adopting a public cloud computing solution, IT organizations need to identify what impact the adoption of these solutions will have on its existing security issues. For example, for the second straight year, Deloitte reports[32] that excessive access rights[33] is the top audit finding. If a company is going to use a business critical application from a SaaS provider it needs to understand both its role and the role of the SaaS provider relative to preventing excessive access rights to that application. In addition, it is possible for IT organizations to use a SaaS provider such as Cloud Compliance[34] to

[30] http://terrorism.about.com/od/beingprepared/tp/OFACCompliance.htm
[31] http://www.ed.gov/policy/gen/guid/fpco/ferpa/index.html
[32] http://www.cloud-compliance.com/blog/bid/27055/Excessive-Access-Rights
[33] Excessive access rights refers to people having access to critical applications even though there is no business justification for why they need access
[34] http://www.cloud-compliance.com/

prevent excessive access rights relative to the applications provided directly by the IT organization.

Before adopting a cloud computing solution, IT organizations need to do a thorough assessment of the CCSP's security and confidentiality policies and capabilities. This includes determining:

- Whether or not the CCSP undergoes regular third party risk assessment audits and will make the results of those audits available to both existing and potential customers.
- Whether or not the CCSP can meet the same security audits as can the IT organization that is using the CCSP's services.
- The encryption capabilities that the CCSP provides.
- The degree to which the CCSP follows well established guidelines such as the Federal Information Security Management Act (FISMA)[35] or National Institute of Science and Technology (NIST) guidelines.
- Whether or not the CCSP has SAS 70 Type II security certification.
- Whether or not it is possible to dictate in which countries the data will be stored.
- Whether or not it is possible to validate where the data is actually stored?
- The tools and processes that the CCSP has implemented to avoid unauthorized access to confidential data.
- Whether or not the CCSP informs the IT organization when someone accesses their data.
- What impact using the CCSP will have on existing security issues such as excessive access rights.
- The degree to which the CCSP has the right and/or intention to make some use of the data provided to it by the IT organization; e.g., analyzing it to target potential customers or to identify market trends.
- The CCSP's policies and procedures relative to data recovery.

[35] http://iase.disa.mil/fisma/index.html

- The CCSP's procedures to avoid issues such as virus attacks and man in the middle intercepts.
- The degree to which the CCSP's staff is trained in security matters.
- Whether or not background checks are performed on each of the CCSP's employees.
- The ability of the CCSP to protect against Cross-site scripting (XSS)[36].
- How the CCSP interprets the third party doctrine[37].

In addition, as a general rule when an IT organization enters into a relationship with any vendor the IT organization should assume that business and/or technology change will occur during the course of the relationship and that at some point in time, the relationship will end. Given that reality, a key component of any agreement that an IT organization establishes with a CCSP has to address issues such as how does the IT organization get their data back if there is a dispute or if the contract has expired or if the CCSP goes out of business? Will the IT organization get their data back in the same format as they submitted it to the CCSP or will the CCSP have reformatted it to fit into their environment?

Management

Roughly forty percent of The Survey Respondents indicated that they had significant or very significant management concerns relative to the use of public cloud computing services. Almost twenty-five percent of The Survey Respondents indicated that they had significant or very significant management concerns relative to the use of private cloud computing services.

The section of this report entitled *The Twelve Characteristics of Cloud Computing* identified some of the management challenges associated with virtualization, automation, the dynamic movement of IT resources, self-service, usage sensitive chargeback and the federation of disparate cloud computing infrastructures.

Some additional concerns relative to managing server virtualization were described in a recent report[38]. These challenges include:

- IT organizations must be able to discover the virtual machine environment in a manner similar to how it discovers the traditional physical environment.

[36] http://en.wikipedia.org/wiki/Cross-site_scripting
[37] http://www.networkworld.com/newsletters/frame/2009/100509wan2.html
[38] http://webtorials.com/abstracts/managing-server-virtualization.htm

- In many cases, once a server is virtualized, IT organizations loose visibility into the inter-VM traffic flows.

- IT organizations need the ability to perform traditional management functions (e.g., troubleshooting, application profiling, baselining and capacity planning) for VMs.

Performing a traditional management function such as troubleshooting in a virtualized environment can be extremely difficult. For example, assume that an IT organization is running an n-Tier application (as depicted in Figure 3) in a virtualized environment and that the web server, application server and database servers are running on VMs in separate physical servers. Further assume that the application server experiences sporadic performance problems and that this causes the n-Tier application to perform badly. In order to troubleshoot this problem and identify that it is the application server that is causing the application to perform badly, the IT organization needs detailed information on each of the three VMs and the communications between them. However, just determining that it is the application server that is causing the application to perform badly is not enough. The IT organization needs to understand why the application server is experiencing sporadic performance problems. The answer to that question might be that other VMs on the same physical server as the application server are sporadically consuming resources needed by the application server. Solving this type of problem will require a level of management data that most IT organizations do not currently collect.

As previously noted above, transferring VMs between servers is complex. Regulations can make moving VMs even more complex. For example, if an IT organization is performing specific logging and auditing as a way of ensuring PCI compliance, then that functionality would have to migrate along with the VM that hosts the application.

One of the fundamental issues relative to managing a public cloud computing service is that as previously noted the service has at least three separate management domains: the enterprise, the WAN service provider and the various cloud computing service providers. Effective management requires that detailed, consistent management data be gathered from each of the management domains. In many cases, IT organizations have the ability to gather detailed management data from their own domain. This ability, however, is sometimes either reduced or eliminated by characteristics of a cloud computing solution such as virtualization or the dynamic movement of resources. IT organizations typically do not have the ability to gather management data from either an NSP or a CCSP. That situation is slowly changing. As previously mentioned, the OGF Open Cloud Computing Interface Working Group is attempting to develop an API specification for remote management of a cloud computing infrastructure. In addition,

Amazon provides a service (CloudWatch[39]) that allows the user to gather metrics such as CPU utilization, disk read/write operations, and throughput from Amazon's APIs.

Before adopting a cloud computing solution, IT organizations need to do a thorough assessment of the CCSP's management capabilities. This includes identifying:

- The amount of management data that the CCSP will make available to the IT organization.
- The ability of the CCSP to troubleshoot performance or availability issues.
- The CCSP's management methodologies for key tasks such as troubleshooting.
- Whether or not the CCSP provides tools such as dashboards to allow the IT organization to understand how well the service they are acquiring is performing.
- Whether or not the CCSP provides detailed information that enables the IT organization to report on their compliance with myriad regulations.
- The degree to which the CCSP has a sufficient number of qualified internal staff to perform troubleshooting vs. relying on crowdsourcing[40].
- The primary management tools that the CCSP utilizes.
- The ongoing training and certification of the CCSP's management personnel.
- The CCSP's backup and disaster recovery capabilities.
- The approach that the CCSP takes to patch management.
- The specific mechanisms that the IT organization can use to retrieve its data back.
- Whether or not the CCSP will allow the IT organization to test the mechanisms to retrieve its data on a regular basis.
- The escalation process to be followed when there are issues to be resolved.
- How the service provided by the CCSP can be integrated with other services provided by either another CCSP and/or by the IT organization.

[39] http://aws.amazon.com/cloudwatch/
[40] http://en.wikipedia.org/wiki/Crowdsourcing

- How the management processes performed by the CCSP can be integrated into the end-to-end management processes performed by the IT organization.

Performance

Over forty percent of The Survey Respondents indicated that they had significant or very significant concerns relative to the performance of public cloud computing services. Over a quarter of The Survey Respondents indicated that they had significant or very significant concerns relative to the performance of private cloud computing services. These concerns caused one of The Survey Respondents to comment that, "High performance requirements will not go in the cloud."

The section of this report entitled *The Twelve Characteristics of Cloud Computing* identified some of the performance challenges associated with the centralization of IT resources, desktop virtualization and the heavy reliance on the Internet. Other performance issues that need to be addressed include:

- The ability of a CCSP to identify and eliminate performance issues.
- The procedures by which the IT organization and the CCSPs will work together to identify and resolve performance problems.
- The actual performance of the service and how that varies by time of day, day of week and week of the quarter.
- The growing *turf wars* between the sub-groups within the IT organization and the CCSPs.
- Whether or not the IT organization has any control over the performance of a public cloud service.
- The performance impact of using multiple CCSPs.
- Whether or not the CCSP provides a meaningful SLA. Does that SLA have a goal for availability? Performance? Is there a significant penalty if these goals are not met? Is there a significant penalty if there is a data breach?
- The degree to which it is possible to customize an SLA.
- The potential that the IT organization is subject to fines and penalties from regulatory bodies if performance is substandard.
- The ability of the CCSP to support peak usage.

- Whether or not the cost of using a WAN service other than the Internet severely erodes the cost advantage of using the CCSP.

- The impact of mobile users accessing a cloud service over a wireless connection.

- The ability to meet state and federal compliance regulations for data availability

In a complex cloud computing environment, the IT organization is using services from multiple CCSPs as part of a single service. One of the key performance issues in this type of situation is how well each CCSP has designed their application to strip off the requisite data and hence minimize the amount of data that has to be transmitted.

As described in the section entitled "Cloud Networking" many IT organizations believe that an optimized WAN will resolve at least some of the performance issues associated with cloud computing. There are many ways to deploy an optimized WAN. This includes implementing:

- An Internet overlay from a company such as Akamai. This provides functionality such as route optimization, TCP and HTTP optimization as well as content offload for applications delivered over the Internet.

- Optimization functionality such as admission control, application based session control and fine grained QoS control on a per application basis delivered as part of an MPLS service.

- The use of an Application Delivery Controller (ADC) to improve the performance of servers in the consolidated data centers.

- The use of a WAN optimization controller (WOC) on each end of the WAN. The WOC can be a traditional physical appliance, a virtual appliance possibly running in a VM, or client software running on a user's PC or laptop or smartphone.

Since WAN optimization controllers (WOCs) are proprietary, in order for a WOC to be beneficial, the WOC on each end of the WAN circuit has to be from the same vendor. This should not be much of a problem for private clouds, but it can an issue be for public clouds.

Summary: Developing The Cloud Computing Plan

The vast majority of The Survey Respondents indicated that their organization sees value in cloud computing. That piece of data marks the end of the discussion of 'Why should I care about cloud computing" and trumpets the entrance of the discussion of "How do I balance the obvious benefits of cloud computing with the obvious risks?"

Relative to the current use of cloud computing, it is no surprise that small organizations are more likely to be using public cloud computing services than they are to be using private cloud computing services. It is also not a surprise that large organizations are more likely to have implemented private cloud computing than they are to have implemented public cloud computing services. That will not change significantly over the next two years. What will change is that two years from now there will be notably more use of both private and public cloud computing than there is today.

Many IT organizations that have already implemented cloud computing have not done so in a highly systematic fashion. In some cases, they used a trial and error approach to choosing a SaaS provider, while in other cases they evaluated one aspect of cloud computing (e.g., server virtualization) in isolation. As previously noted, in order to maximize the benefit of moving to a cloud computing model, IT organizations need to develop a plan (The Cloud Computing Plan) which they update on a regular basis. The Cloud Computing Plan should identify the opportunities and risks associated with both public and private cloud computing. (This report will assist with the identification of risks.) The Cloud Computing Plan must identify a roadmap of what steps the IT organization will take on a quarter-by-quarter basis for the next two to three years and ensure that the steps are in line with the corporate culture. This includes identifying:

- What functionality (e.g., applications, storage) needs to remain under the tight control of the IT organization and what functionality is appropriate to hand over to a CCSP.
- What levels of service are *good enough* for each class of application and for the myriad storage requirements.
- How the IT organization will evolve over time the twelve characteristics of a cloud computing solutions; e.g., virtualization, automation, simplification.
- How the IT organizations will evolve its LAN architecture to support cloud computing.
- How the IT organizations will evolve its use of WAN services to support cloud computing.
- How the IT organization will minimize the security and confidentiality risks associated with public cloud computing services.

- What management functionality must be present in the management domain controlled by the IT organization as well as provided by the relevant network service providers and CCSP(s).
- How the IT organization will overcome potential performance bottlenecks.

The Cloud Computing Plan should look systematically across multiple technologies because of the interconnected nature of the technologies; e.g., when you virtualize the data center servers this has an impact on both the data center LAN and on network management. As part of creating this plan, IT organizations need to understand the cloud computing strategy of their existing and potential suppliers, including the partnerships that they are establishing between and amongst themselves.

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.

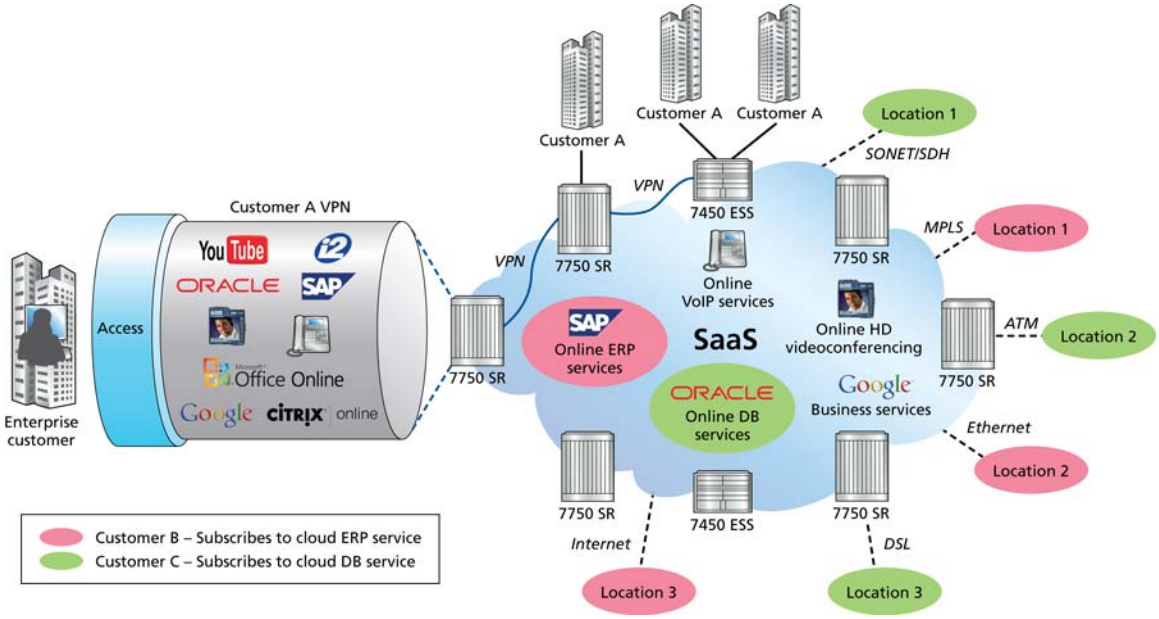

The Alcatel-Lucent AA architecture provides the ability to recognize applications and application flows throughout the network. This in turn enables the service provider to report on the applications and apply application-level QoS controls. The Alcatel-Lucent AA architecture supports multiple small, 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

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.

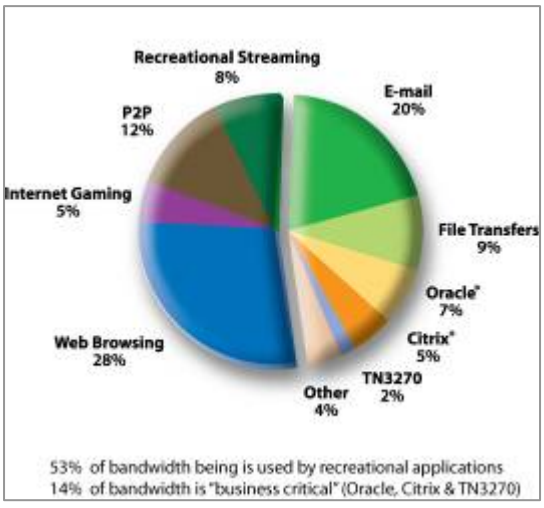

Your business depends on getting the right information to the right people on time, by assuring fast and responsive application delivery while protecting users and data from malicious threats. Are you supporting the delivery of information as reliably and safely as your business demands? Now is the time 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.

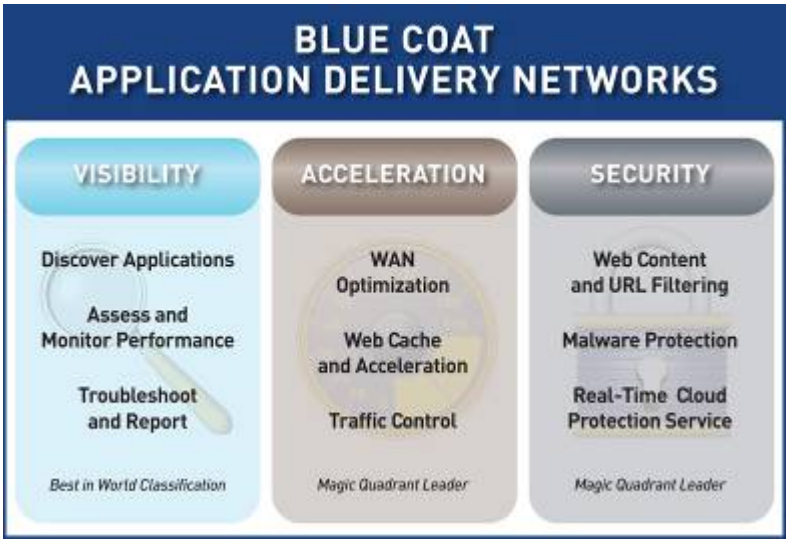

Finally, Blue Coat secures users and data with an industry-leading cloud defense and inline threat analysis for comprehensive malware scanning and filtering, centralized policy management, and in-depth reporting and logging.

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.