The 2012 Cloud Networking Report

Part 1: The Emergence of Cloud Computing and Cloud Networking

By Dr. Jim Metzler
Ashton Metzler & Associates
Distinguished Research Fellow and Co-Founder
Webtorials Analyst Division

Sponsored in part by:

Alcatel·Lucent Enterprise

Produced by:

Webtorials

Table of Contents

Executive Summary

The 2012 Cloud Networking Report (The Report) will be published both in its entirety and in a serial fashion. This is the first of the serial publications. As pointed out in this publication, the phrase ***cloud networking*** refers to the LAN, WAN and management functionality that must be in place to enable cloud computing. In order for the report to intelligently describe the networking challenges that are associated with enabling cloud computing, this publication will identify what cloud computing is today and will also describe how cloud computing is likely to evolve in the near term. Subsequent publications will focus on the key components of a cloud network: Data Center LANs, WANs, and Network Management. Given the growing interest in the topic, there will also be a separate section on Software Defined Networks.

The Emergence of Cloud Computing and Cloud Networking

Introduction and Forward to the 2012 Edition

Numerous analyst reports have pointed out the broad interest that IT organizations have in deploying one or more classes of cloud computing. For example, Gartner[1] recently stated that they expected that cloud computing would grow 19% in 2012, and would become a $109 billion industry. Gartner also stated that that they expected that by 2016, that cloud computing would be a $207 billion industry. The high growth rate in the cloud computing market is in sharp contrast to the annual growth rate of the overall IT market, which Garters estimates to be 3%. The broad interest in cloud computing is understandable given that the goal of cloud computing is to enable IT organizations to become dramatically more agile and cost effective and that evidence exists that that goal is achievable.

The primary goal of this report is to describe the network related challenges and solutions that are associated with cloud networking.

The phrase cloud networking refers to the LAN, WAN and management functionality that must be in place to enable cloud computing.

As will be discussed in this report, a traditional network will not be able to successfully support cloud computing.

In order to support cloud computing, a cloud network must be dramatically more agile and cost effective than a traditional network.

In order to describe the networking challenges that are associated with enabling cloud computing, the rest of this section of the report will identify what cloud computing is today and will also describe how cloud computing is likely to evolve in the near term. Subsequent sections focus on the key components of a cloud network: Data Center LANs, WANs, and Network Management. A subsequent section will also focus on an emerging component of a cloud network: software defined networks (SDNs). Given the breadth of fundamental technology changes that are impacting the data center LAN, the data center LAN section is very technical. The sections on WANs, SDNs and Network Management are moderately technical. This year's edition of the cloud networking report leverages last year's edition of the report[2]. However, every section of The 2011 Cloud Networking Report has been significantly updated to reflect the changes that have occurred in the last year.

As noted, the primary goal of this report is to describe the network related challenges and solutions that are associated with cloud networking. A secondary goal of this report is to identify how IT organizations are currently approaching both cloud computing and cloud networking and where possible, indicate how that approach is changing. To accomplish that goal, this report includes the results of surveys that were recently given to the subscribers of Webtorials.com. Throughout this report, the IT professionals who responded to those surveys will be referred to

[1] http://www.networkworld.com/news/2012/071312-gartner-cloud-260882.html
[2] http://www.webtorials.com/content/2011/11/2011-cloud-networking-report.html

as the ***Survey Respondents***. In some cases, the results of the surveys given to the Survey Respondents will be compared to the results of surveys given in 2011. In addition, the SDN section of The Report will include the results of a survey that was conducted in conjunction with Information Week. Throughout the SDN section of this report, the IT professionals who responded to the SDN survey will be respectively to as the ***Information Week Respondents***.

The results of surveys such as the ones described in the preceding paragraph that ask IT organizations about their plans are always helpful because they enable IT organizations to see how their own plans fit with broad industry trends. Such surveys are particularly beneficial in the current environment when so much change is occurring.

The Goal of Cloud Computing

Within the IT industry there still isn't a universally accepted definition of what is meant by cloud computing. The Report takes the position that it is notably less important to define exactly what is meant by the phrase *cloud computing* than it is to identify the goal of cloud computing.

The goal of cloud computing is to enable IT organizations to achieve a dramatic improvement in the cost effective, elastic provisioning of IT services that are good enough.

In order to demonstrate the concept behind the phrase *good enough,* consider just the availability of an IT service. In those cases in which the IT service is business critical, *good enough* could mean five or six 9's of availability. However, in many other cases *good enough* has the same meaning as *best effort* and in these cases *good enough* could mean two or three 9's of availability. The instances in which an approach that provides two or three 9's of availability is acceptable are those instances in which the IT service isn't business critical and that approach is notably less expensive than an alternative approach that offers higher availability.

On a going forward basis, IT organizations will continue to need to provide the highest levels of availability and performance for a small number of key services. However, an ever-increasing number of services will be provided on a best effort basis.

In most instances the SLAs that are associated with public cloud computing services such as Salesforce.com or Amazon's Simple Storage System are weak and as such, it is reasonable to say that these services are delivered on a best effort basis. For example, the SLA[3] that Amazon offers for its Amazon Web Services (AWS) states that, "AWS will use commercially reasonable efforts to make Amazon EC2 available with an Annual Uptime Percentage of at least 99.95% during the Service Year." As part of the Amazon definition of Annual Uptime Percentage, Amazon excludes any outage of 5 minutes or less. The Amazon SLA also states that if their service doesn't meet the Annual Uptime Percentage commitment, the customer will receive 10% off its bill for the most recent month that the customer included in the SLA claim that it filed.

A key attribute of the vast majority of the SLAs that are associated with public cloud computing services is that they don't contain a goal for the end-to-end performance of the service. The reason for the lack of performance guarantees stems from the way that most public cloud computing services are delivered. As shown in **Figure 1**, one approach to providing public cloud computing services is based on the service being delivered to the customer directly from an independent software vendor's (ISV's) data center via the Internet. This is the distribution model currently used for Salesforce.com's CRM application. Another approach is for an ISV to leverage an IaaS provider such as Amazon to host their application on the Internet. Lawson Software's Enterprise Management Systems (ERP application) and Adobe's LiveCycle Enterprise Suite are two examples of applications hosted by Amazon EC2. Both of these approaches rely on the Internet and it is not possible to provide end-to-end quality of service (QoS) over the Internet. As a result, neither of these two approaches lends itself to providing an

[3] http://aws.amazon.com/ec2-sla/

SLA that includes a meaningful commitment to critical network performance metrics such as delay, jitter and packet loss.

Figure 1: Distribution Models for Cloud-Based Solutions

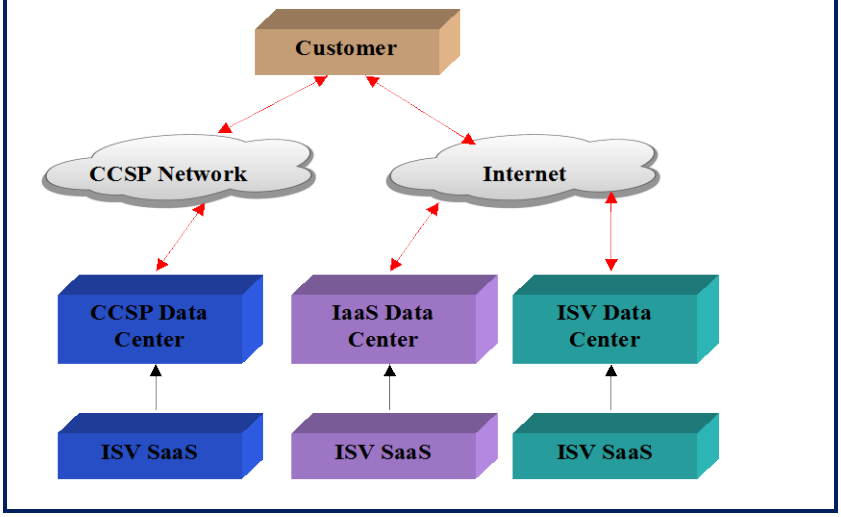

The fact that cloud computing service providers (CCSPs) don't provide an end-to-end performance SLA for applications delivered over the Internet will not change in the foreseeable future. However, as will be described in a subsequent section of this report, there are things that can be done to improve the performance of applications delivered over the Internet.

An approach to providing public cloud computing services that does lend itself to offering more meaningful SLAs is based on a CCSP providing these solutions to customers from the CCSP's data center and over a network that is provided by the CCSP and based on a technology such as MPLS.

Organizations that utilize best effort cloud computing services do so with the implicit understanding that if the level of service they experience is not sufficient; their primary recourse is to change providers. It may seem counter-intuitive that a company would utilize public cloud computing services for which end-to-end performance SLAs are essentially non-existent. However, as described in a subsequent section of this report, two thirds of The Webtorials Respondents indicated that the SLAs that they receive from their network service providers for services such as MPLS are either not worth the paper they are written on, or that the SLAs they receive are not much better than nothing.

SLAs from both traditional network service providers as well as public cloud computing providers are a work in progress.

Characteristics of Cloud Computing Solutions

The following set of bullets identifies the primary characteristics of cloud computing solutions. There is not, however, a litmus test to determine if a particular service is or is not a cloud computing service.

- Centralization of applications, servers, data and storage resources.
- Extensive virtualization of every component of IT, including servers, desktops, applications, storage, switches, routers and appliances such as WAN optimization controllers, application delivery controllers and firewalls.
- Automation and Orchestration of as many tasks as possible; e.g., provisioning, troubleshooting, change and configuration management.
- The dynamic creation and movement of resources such as virtual machines and the associated storage.
- Heavy reliance on the network.
- Self-service to allow end users to select and modify their use of IT resources without the IT organization being an intermediary.
- Usage sensitive chargeback that is often referred to as pay-as-you-go. An alternative is for IT organizations to show the consumption of IT resources by certain individuals or organizations; a.k.a., showback.
- Simplification of the applications and services provided by IT.
- Standardization of the IT infrastructure.
- Technology convergence such as the convergence of LAN and SAN and of switch and server.
- The development of standards that enable, among other things, the federation of disparate cloud computing infrastructures with one another (see below).
- The federation of disparate cloud computing infrastructures with one another.

Classes of Cloud Computing Solutions

There are three classes of cloud computing solutions that will be described in this section of the report. Those classes are private, public and hybrid.

Private Cloud Computing

Many IT organizations have decided to implement some of the characteristics of cloud computing solutions described in the preceding subsection within their internal IT environment. This approach is usually referred to as a *Private Cloud*. One of the primary ways that IT organizations have adopted private cloud computing solutions is by implementing some or all of the previously mentioned characteristics of cloud computing solutions in order to be able to provide Infrastructure-as-a-Service (IaaS) solutions that are similar to the solutions offered by IaaS providers such as Rackspace.

The initial set of IaaS solutions that were brought to market by IaaS providers were the basic compute and storage services that are necessary to run applications. However, the IaaS market is highly dynamic and IaaS providers are deploying myriad new services including:

- Disaster Recovery
- Virtual Private Data Centers
- High Performance Computing

The Survey Respondents were given a set of 7 possible approaches to IaaS and were asked to indicate which approach best described their company's approach to using IaaS solutions, either provided internally by their own IT organization, or provided externally by an IaaS provider. The Survey Respondents were allowed to indicate as many approaches as were appropriate. Their responses are shown in **Table 1**.

Table 1: Approach to IaaS *N=171*

Approach	Percentage of Respondents
We are in the process of developing a strategy	48.0%
We provide IaaS solutions internally for a wide range of applications	19.9%
We provide IaaS solutions internally for a small range of applications	19.9%
We have a well-defined and understood strategy	15.2%
We only use IaaS solutions from a CSP for a small set of applications that are not business critical	14.6%
We use IaaS solutions from a CCSP for a wide range of applications	12.3%
Other	7.0%
We only outsource either a trial of the initial deployment of an application to a CCSP	6.4%
We have a policy against using any IaaS solutions provided by a CCSP	3.5%

One key conclusion that can be drawn from the data in **Table 1** is that:

Roughly half of all IT organizations are currently in the process of developing a strategy for how they will use public and private IaaS solutions.

The Survey Respondents were asked to indicate the two primary factors that limit their company's interest in using internally provided IaaS solution. The five inhibitors to the adoption of private IaaS solutions that were indicated the most times by the Survey Respondents and the percentage of times that they were mentioned were:

- Concerns about the security and confidentiality of data (36.3%)
- Their lack of an internal strategy about IaaS (28.7%)
- Their lack of personnel to design and implement the solutions (25.7%)
- The relative immaturity of the technologies that would have to installed and managed (19.9%)
- The lack of significant enough cost savings (19.3%)

While the conventional wisdom in our industry is that security and confidentiality of data is the major impediment to the adoption of public cloud based IaaS solutions, it is somewhat surprising that:

Concern about the security and confidentiality of data is the primary impediment to the broader adoption of private IaaS solutions.

Public Cloud Computing

This section of The Report will focus on the two most popular types of public cloud computing solutions: Software-as-a-Service and Infrastructure-as-a-Service.

Software-as-a-Service

According to Gartner[4], the Software as a Service (SaaS) market will have worldwide revenues of $22.1 billion by 2015. One of the key characteristics of the SaaS marketplace is that:

The SaaS marketplace is comprised of a small number of large players such as Salesforce.com, WebEx and Google Docs as well as thousands of smaller players.

One of the reasons why there are so many players in the SaaS market is that the barrier to entry is relatively low.

The Survey Respondents were asked to indicate if their company currently acquires applications from a SaaS provider or if they are likely to within the next twelve months. Their responses are shown in **Figure 2**.

Figure 2: Adoption of SaaS Solutions *N=264*

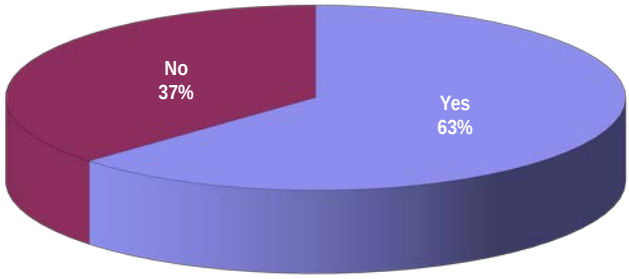

The Survey Respondents were then given a set of 7 types of applications and were asked to indicate the types of applications that their company currently acquires from a SaaS provider and the types of applications that their organization would likely acquire from a SaaS provider over the next twelve months. Their responses are shown in **Table 2**.

Table 2: Interest in SaaS *N=153*

	Currently Acquire	Will Acquire
Collaboration	55%	31%
Customer Relationship Management CRM)	53%	22%
Human Resources	45%	18%
Office Productivity	40%	33%
Project and Portfolio Management	27%	54%
Enterprise Resource Planning (ERP)	24%	16%
Supply Chain Management (SCM)	15%	27%

[4] http://www.slideshare.net/rajeshdgr8/global-saa-s-2012

The Survey Respondents were given a set of ten factors and were asked to indicate the two factors that were the primary drivers of their organization's interest in using SaaS solutions. The responses of the Survey Respondents are shown in **Table 3**. In **Table 3**, the column on the right is labeled *Percentage of Respondents*. That column contains the percentage of the Survey Respondents that indicated that the factor in the left hand column of **Table 3** was one of the two primary drivers of their organization's interest in using SaaS solutions.

Table 3: Factors Driving the Adoption of SaaS Solutions *N=153*	
Factor	**Percentage of Respondents**
Lower cost	39%
Reduce the amount of time it takes to implement an application	35%
Free up resources in the IT organization	29%
Deploy applications that are more robust; e.g., available and scalable	27%
Easier to justify OPEX than CAPEX	26%
Leverage the expertise of the SaaS provider	19%
Reduce risk	11%
Management mandate as our strategic direction	8%
Meet temporary requirements	3%
Other	2%

One conclusion that can be drawn from the data in **Table 3** is that:

The primary factors that are driving the adoption of SaaS are the same factors that drive the adoption of any form of out-tasking.

Given the concerns that IT organizations have relative to the security and confidentiality of their data, it appears to be counter intuitive that 11% of the Survey Respondents indicated that reducing risk was a factor that would cause them to use a public cloud computing solution. In most cases the Survey Respondents' reasoning was that acquiring and implementing a large software application (e.g., ERP, CRM) presents considerable risk to an IT organization and one way to minimize this risk is to acquire the functionality from a SaaS provider.

Infrastructure as a Service (IaaS)

The barrier to enter the IaaS marketplace is notably higher than is the barrier to enter the SaaS marketplace. That is one of the primary reasons why there are fewer vendors in the IaaS market than there are in the SaaS market. Representative IaaS vendors include Amazon, AT&T, CSC, GoGrid, IBM, Joyent, NTT Communications, Orange Business Services, Rackspace, NaviSite (acquired by Time Warner), Savvis (acquired by Century Link), Terremark (acquired by Verizon) and Verizon. As the preceding sentence indicates, the IaaS market is going through a period that is characterized by mergers and acquisitions. The IaaS market is also expected to exhibit significant growth in the next few years. For example, Gartner[5] estimates that the IaaS market will grow from $3.7 billion in 2011 to $10.5 billion in 2014.

The Survey Respondents were asked to indicate the IaaS services that their organization currently acquires from a CCSP and the services that their organization will likely acquire from a CCSP during the next year. Their responses are shown in **Table 4**.

Table 4: Current and Planned Adoption of IaaS Services *N = 142*

	Currently Acquire	Will Likely Acquire
Storage	26.8%	16.9%
Computing	26.8%	9.2%
Virtual Private Data Center	17.6%	14.1%
Disaster Recovery	16.2%	21.8%
High Performance Computing	10.6%	9.9%

Because storage and computing were the initial set of IaaS services that were brought to market, it was not at all surprising to see that over a quarter of the Survey Respondents indicated that they currently used those services. In addition, given that high performance computing (HPC) is somewhat of a niche application, it was not surprising that there was relatively little interest in acquiring HPC from an IaaS supplier. However it was somewhat of a surprise to see that:

There is strong interest on the part of IT organizations in acquiring both virtual private data center and disaster recovery services from IaaS providers.

5 http://www.qas.com/company/data-quality-news/iaas_market_to_record_strong_growth__7178.htm

Drivers and Inhibitors

This section will discuss the factors that are driving and the factors that are inhibiting the deployment of IaaS solutions.

- **Drivers**
  The Survey Respondents were given a set of eleven factors and were asked to indicate the two factors that were the primary drivers of their organization's interest in using Cloud-based IaaS solutions. The responses of the Survey Respondents are shown in **Table 5**. In **Table 5**, the column on the right is labeled *Percentage of Respondents*. That column contains the percentage of the Survey Respondents that indicated that the factor in the left hand column of **Table 5** was one of the two primary drivers of their organization's interest in using Cloud-based IaaS solutions.

Table 5: Factors Driving the Adoption of IaaS Solutions *N = 171*

Factor	Percentage of Respondents
Lower cost	30.4%
The ability to dynamically add capacity	30.4%
Reduce time to deploy new functionality	26.3%
Obtain functionality we are not able to provide ourselves	22.2%
Deploy more highly available soluti`ons	19.3%
Free up resources	17.0%
Easier to justify OPEX than CAPEX	15.8%
Prefer to only pay for services that we use	14.0%
Satisfy temporary requirements	11.7%
Other	4.7%
Our strategy is to use IaaS providers wherever possible	4.1%
Leverage the security expertise of the provider	4.1%

The conventional wisdom in the IT industry is that lower cost is the primary factor driving the adoption of Cloud-based IaaS solutions and that factors such as the ability to dynamically add new capacity, while important, are nowhere near as important. As the data in **Table 5** highlights, the reality is that the ability to dynamically add new capacity is as important a driver of the adoption of Cloud-based IaaS solutions as is lowering cost. In addition, another very important driver of the adoption of Cloud-based IaaS solutions is the ability to reduce the time it takes to deploy new functionality. It is reasonable to look at the ability to dynamically add capacity and the ability to reduce the time it takes to deploy new functionality as two components of a single factor – agility. Looked at this way,

By a wide margin, agility is the most important factor driving the adoption of Cloud-based IaaS solutions.

- **Inhibitors**
  The Survey Respondents were asked to indicate the two primary factors that limit their company's interest in using a Cloud-based IaaS solution. Those factors and the percentage of times that they were indicated by the Survey Respondents are shown in **Table 6**.

Table 6: Inhibitors to the adoption of Cloud-based IaaS Solutions *N = 171*

Factor	Percentage of Respondents
We are concerned about the security and confidentiality of our data	57.9%
We don't see significant enough cost savings	24.0%
The lack of time and resources to sufficiently analyze the offerings and the providers	19.9%
Uncertainty about the provider living up to their promises	19.9%
We have concerns about the availability of the solutions	16.4%
Our lack of confidence in a shared infrastructure	15.2%
The lack of a meaningful SLA	14.6%
We don't believe that the gains in the agility of these solutions justifies the cost and/or the risk	11.7%
Our policy is to either limit or totally avoid using IaaS providers	8.8%
The provider is not capable of adding capacity in a dynamic enough fashion	4.7%

One conclusion that can be drawn from the data in **Table 6** is:

Concern about the security and confidentiality of data is by a wide margin the number one factor inhibiting the adoption of Cloud-based IaaS solutions

A component of the concerns that IT organization have about security and confidentiality stems from the overall increase in the sophistication of hackers, For example, until relatively recently the majority of security attacks were caused by individual hackers, such as Kevin Mitnick, who served five years in prison in the late 1990s for computer- and communications-related hacking crimes. The goal of this class of hacker is usually to gain notoriety for themselves and they often relied on low-technology techniques such as dumpster diving.

However, over the last few years a new class of hacker has emerged and this new class of hacker has the ability in the current environment to rent a botnet or to develop their own R&D lab. This new class includes crime families and hactivists such as Anonymous. In addition, some national governments now look to arm themselves with Cyber Warfare units and achieve their political aims by virtual rather than by physical means.

The sophistication of the current generation of hackers was highlighted in the Blue Coat Systems 2012 Web Security Report[6], which focused on a number of topics including malnets and social networking. A malware network, or malnet, gathers users, most frequently when they are visiting trusted sites and routes them to malware. According to the Blue Coat Report, "In 2011, malnets emerged as the next evolution in the threat landscape. These infrastructures last beyond any one attack, allowing cybercriminals to quickly adapt to new vulnerabilities and repeatedly launch malware attacks. By exploiting popular places on the Internet, such as search engines, social networking and email, malnets have become very adept at infecting many users with little added investment."

The report noted the increasing importance of social networking and stated that, "Since 2009, social networking has increasingly eclipsed web-based email as a method of communications." The report added that, "Now, social networking is moving into a new phase in which an individual site is a self-contained web environment for many users – effectively an Internet within an Internet." For example, according to the Blue Coat report 95% content types that are found on the Internet are also found within social networking sites. The five most requested subcategories of content that were requested from social networking sites, and the percentage of times that they were requested are shown in **Table 7**.

Table 7: Most Requested Content from Social Media Sites

Subcategory of Content	Percentage of Times it was Requested
Games	37.9%
Society/Daily Living	23.8%
Personal Pages/Blogs	6.4%
Pornography	4.9%
Entertainment	4.2%

Part of the challenge that is associated with social network sites being so complex is that IT organizations cannot just look at a social media site as one category and either allow or deny access to it. Because these sites contain a variety of classes of content, IT organizations need the granular visibility and control to respond differently to requests at the same social media site for different types of content.

Another component of the concern that IT organizations have about security and confidentiality of their data stems from the fact that in most cases IT organization perceive that there is a higher security risk if their data is being stored on a device that is shared with other users which is tpically the case when an IT organization is using an IaaS solution. The security risk that is associated with all forms of cloud computing was discussed in IBM's X-Force 2011 Trend and Risk Report[7] that was published in March 2012. According to the IBM report, in 2011, there were many high profile cloud breaches affecting well-known organizations and large populations of their customers. IBM recommended that IT security staff should carefully consider which workloads are

[6] http://www.bluecoat.com/sites/default/files/documents/files/BC_2012_Security_Report-v1i-optimized.pdf
[7] X-Force 2011 Trend and Risk Report

sent to third-party cloud providers and what should be kept in-house due to the sensitivity of data. The IBM X-Force report also noted that the most effective means for managing security in the cloud may be through Service Level Agreements (SLAs) and that IT organizations should pay careful consideration to ownership, access management, governance and termination when crafting SLAs.

The Role of Virtualized Network Services

As prevously noted, one of the primary goals of The Report is to identify what functionality is needed in the network to support cloud computing. With that goal in mind, the Survey Respondents were given a number of questions that related to the role that virtualized network services play in their evaluation and selection of Cloud-based IaaS services.

One of the questions contained a set of network services and the Survey Respondents were asked to indicate if they thought the network service should be part of a Cloud-based IaaS service and if they did, whether they preferred to manage the network service themselves or have the CSP manage it. The vast majority of the Survey Respondents (87+%) thought that each one of the network servcies listed in **Table 8** should be part of a Cloud-based IaaS service. Columns two and three of **Table 8** respectively contain the percentage of the Survey Respondents who prefer to manage the service themselves as well as the percentage of the Survey Respondents who prefer to have a CSP manage the service.

Table 8: The Applicability and Management of Network Services ***N = 171***

Network Service	Manage Ourselves	CSP Manage
Load Balancer	61.9%	38.1%
SSL Load Balancer	62.2%	37.8%
Firewall	81.4%	18.6%
WEB application firewall	68.5%	31.5%
IDS/IPS	64.1%	35.9%
VPN	70.2%	29.8%
WAN optimization	50.8%	49.2%

One obvious conclusion that can be drawn from the data in **Table 8** is:

There is a strong desire on the part of IT organizations to manage the security related network services that are part of an IaaS service.

Because IT organizations expect that Cloud-based IaaS services are supported by a wide range of network services, this raises the question, “When evaluating IaaS services, how carefully do IT organizations evaluate the associated network services?” To answer that question, the Survey Respondents were asked, “When your organization evaluates cloud services such as computing, storage and virtual private data centers, how carefully does your organization evaluate the enabling network services such as Load Balancer, SSL Load Balancer, Firewall?” Their answers are contained in **Table 9**.

Table 9: Importance of Network Services	N = 171
How Carefully	**Percentage of Respondents**
We don't evaluate them at all	8.6%
We look at them as a check-off item, but don't evaluate	10.0%
We pay some attention to them, but they are not a major component of the evaluation process	21.4%
They are a major component of the overall evaluation process	33.6%
They are a critical component of the overall evaluation process	26.4%

One obvious conclusion that can be drawn from the data in **Table 9** is:

The evaluation of the supporting network services is a key component of the overall process of evaluating IaaS solutions.

Given the critical role that network services play in the evaluation of Cloud-based IaaS services, the Survey Respondents were asked to indicate the two most important criteria they look for when evaluating network services such as a Load Balancer, an SSL Load Balancer, or a Firewall, that enable cloud services. The criteria and the percentage of times that they were indicated by a survey respondent are shown in **Table 10**.

Table 10: Criteria to Evaluate Networking Services	N = 171
Criteria	**Percentage of Respondents**
A robust feature set similar to traditional networking equipment	25.9%
The ability to grow/shrink the capacity of the service on demand	23.8%
The ability to rapidly provision the network service; e.g., 5 minutes or less	21.1%
The ability to only pay for what we use	17.8%
A brand name vendor	6.3%
The ability to charge back to business units based on usage	5.1%

The conventional wisdom is that when IT organizations evaluate network services, that a name brand vendor is an important criterion. The data in **Table 10** refutes that belief as the data in the table highlights the fact that a robust feature set is the single most important criterion that IT organizations examine with evaluating networks services. However, another way to evaluate the data in **Table 10** is based on the previous definition of agility[8]. Looked at this way, the data in **Table 10** clearly indicates that the agility of network services is the most important criterion that IT organizations examine with evaluating networks services.

In order to understand the organizational dynamic that underlies the decision to use an IaaS solution from a CSP, the Survey Respondents were asked about the roles of the organizations

[8] In this context, agility is the ability to dynamically add capacity and the ability to reduce the time it takes to deploy new functionality.

that are involved in making that decision, Their responses, shown in **Table 11**, indicate how the decision is made.

Table 11: The Decision Making Process
N=160

Role	Percentage of Respondents
Largely by the IT organization with some input from the business or functional unit	40.0%
The IT unit and the business or functional unit participate equally	26.3%
Largely by the business or functionaly unit with some input from the IT organization	15.6%
Entirely by the IT organization	11.3%
Entirely by the business or functional unit	6.9%

One obvious conclusion that can be drawn from the data in **Table 11** is:

> ***Roughly 20% of the times that a company is evaluting public IaaS solutions, the company's IT organization is either not involved at all or plays a minor role.***

Hybrid Cloud Computing

Like so much of the terminology of cloud computing, there is not a uniformly agreed to definition of the phrase ***hybrid cloud computing***. According to Wikipedia[9], "Hybrid cloud is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together, offering the benefits of multiple deployment models. Briefly it can also be defined as a multiple cloud systems which are connected in a way that allows programs and data to be moved easily from one deployment system to another."

Based on this definition, one form of a hybrid cloud is an n-tier application in which the web tier is implemented within one or more public clouds while the application and database tiers are implemented within a private cloud. Another form of hybrid cloud that receives a lot of attention is cloud balancing. The phrase ***cloud balancing*** refers to routing service requests across multiple data centers based on myriad criteria. As shown in **Figure 3**, cloud balancing involves one or more corporate data centers and one or more public cloud data centers.

Cloud balancing can be thought of as the logical extension of global server load balancing (GSLB).

Figure 3: Cloud Balancing

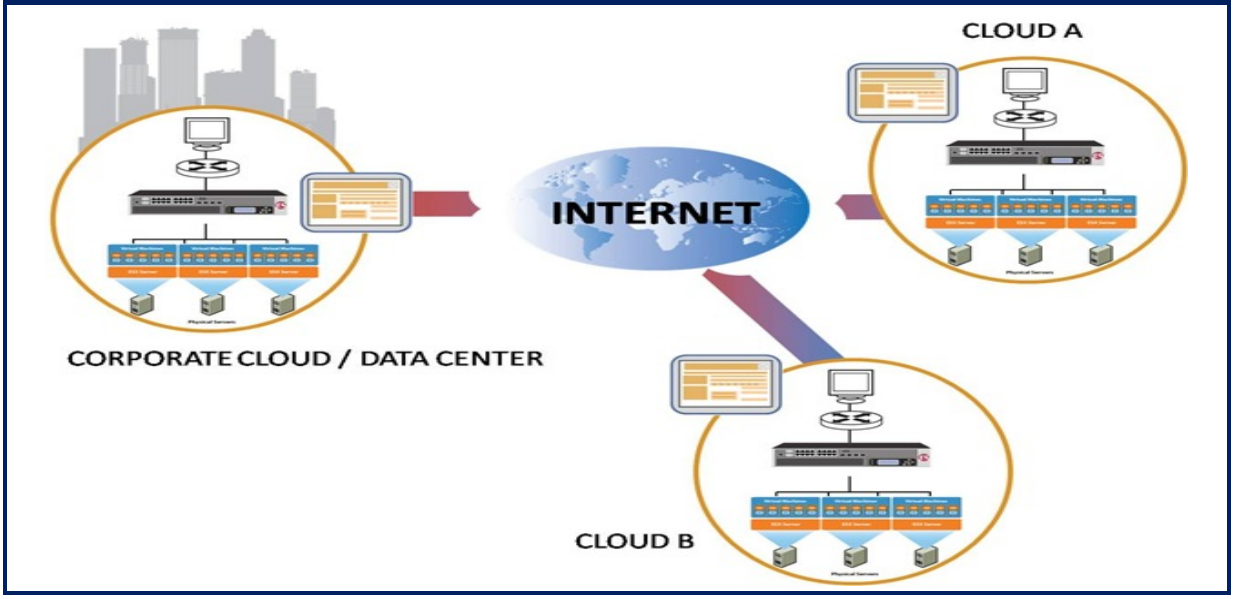

[9] http://en.wikipedia.org/wiki/Cloud_computing#Hybrid_cloud

The goal of a GSLB solution is to support high availability and maximum performance. In order to do this, a GSLB solution typically makes routing decisions based on criteria such as the application response time or the total capacity of the data center. A cloud balancing solution may well have as a goal supporting high availability and maximum performance and may well make routing decisions in part based on the same criteria as used by a GSLB solution. However, a cloud balancing solution extends the focus of a GSLB solution to a solution with more of a business focus. Given that extended focus, a cloud balancing solution includes in the criteria that it uses to make a routing decision the:

- Performance currently being provided by each cloud
- Value of the business transaction
- Cost to execute a transaction at a particular cloud
- Relevant regulatory requirements

Some of the benefits of cloud balancing include the ability to:

Maximize Performance
Routing a service request to a data center that is close to the user and/or to one that is exhibiting the best performance results in improved application performance.

Minimize Cost
Routing a service request to a data center with the lowest cost helps to reduce the overall cost of servicing the request.

Minimize Cost and Maximize Service
Cloud balancing enables a service request to be routed to a data center that provides a low, although not necessarily the lowest cost while providing a level of availability and performance that is appropriate for each transaction.

Regulatory Compliance
For compliance with regulations such as PCI, it may be possible to partition a web services application such that the PCI-related portions remain in the PCI-compliant enterprise data center, while other portions are cloud balanced. In this example, application requests are directed to the public cloud instance unless the queries require the PCI-compliant portion, in which case they are directed to the enterprise instance.

Manage Risk
Hosting applications and/or data in multiple clouds increases the availability of both. Balancing can be performed across a number of different providers or it can be performed across multiple independent locations of a single cloud service provider.

Emerging Public Cloud Computing Services

Data Center Services

Most of the IaaS providers do not want to compete entirely based on providing commodity services such as basic compute and storage. As such, many IaaS providers are implementing higher value-added data center services such as the ones described below.

Private Cloud Data Center Services

These services are based on outsourcing the enterprise's multi-tier private data center to a service provider. The data center could be located at either a site controlled by the enterprise or at a service provider's site. In most cases service providers will structure these services so that the customers receive the highest levels of support, as well as assurances written into the corresponding SLA for high levels of availability, performance and security. A private WAN service would typically be used to provide access to these services.

Virtual Private Data Center (VPDC)

These services provide an instance of an entire data center hosted on a service provider's infrastructure that is optimized to provide a high level of security and availability for multiple tenants. From the service provider's perspective, the data center architecture for the VPDC would be similar to the architecture used for a private cloud data center except that the resources would be shared among a number of customers rather than being dedicated to a single customer or tenant. The service provider's architecture needs to effectively leverage virtualization in order to maximize the efficient usage of a shared pool of resources. The architecture also needs to allow for a high degree of flexibility in providing a broad range of required network capabilities. This includes WAN optimization, load balancing and firewall services. Service management software should be in place to enable the co-management of the VPDC by customers and providers.

The hybrid cloud computing model works best in those instances in which the VPDC and the private cloud data center are based on the same hypervisors, hypervisor management systems and cloud controllers. This maximizes the enterprise's control over the hybrid cloud and allows application and server management to remain the responsibility of the enterprise. Access to a VPDC could be provided either over the Internet or a private WAN service.

Cloud Networking Services

With the exception of collaboration, the applications that organizations have historically acquired from CCSPs have typically been enterprise applications such as CRM. Recently, a new class of solutions has begun to be offered by CCSPs. These are solutions that have historically been provided by the IT infrastructure group itself and include network and application optimization, VoIP, Unified Communications (UC), security, network management and virtualized desktops. Within The Report, this new class of solutions will be referred to as Cloud Networking Services (CNSs).

The Survey Respondents were given a set of 7 CNSs and were asked to indicate the CNSs that their organization currently acquires from a CCSP and the services that their organization will likely acquire from a CCSP during the next year. Their responses are shown in **Table 12**.

Table 12: Current and Planned Adoption of CNSs *N = 142*

	Currently Acquire	Will Likely Acquire
VoIP	20.4%	17.6%
Network Management	19.7%	8.5%
Security	18.3%	9.9%
Unified Communications	15.5%	23.2%
Application Performance Management	10.6%	10.6%
Network and Application Optimization	8.5%	9.2%
Virtual Desktops	7.0%	19.0%

The data in **Table 12** shows that the interest in CNS is quite broad, as over twenty-five percent of the survey respondents indicated that over the next year that five of the seven services listed in the table would either likely be acquired, or would be acquired.

Cloud Networking Services represents the beginning of what could be a fundamental shift in terms of how IT services are provided.

Since CNS solutions are just one more form of public cloud computing, when evaluating these solutions IT organizations need to understand the degree to which these solutions overcome the factors that impede the use of any public cloud computing solution. Since concerns about security are typically one of the primary impediments to the adoption of public cloud computing solutions, evaluating the security of the CNS provider's facilities is a critical component of evaluating a CNS solution.

However, just as important as whether or not the CNS solution provides adequate security is whether or not the solution actually provides the benefits that drive IT organizations to use public cloud computing solutions. As previously discussed, the primary benefit of using a public cloud computing solution is typically lower cost. While it can be tricky to compare the usage sensitive pricing of the typical CNS solution with the fully loaded cost of a premise based solution, the cost information provided by the CCSP should give the IT organization all the information it needs to do that analysis. Another key benefit of using a public cloud computing solution is being able to reduce the time it takes to deploy new functionality. Evaluating the agility of a CCSP is notably more difficult than evaluating their cost structure.

One way for an IT organization to evaluate the agility of a CCSP is to identify the degree to which the CCSP has virtualized their infrastructure.

This follows because a virtual infrastructure is notably easier to initialize, scale and migrate than a physical infrastructure is. Since the vast majority of CCSPs implement virtualized servers, server virtualization is unlikely to distinguish one CCSP from another. What can distinguish one CCSP from another is the degree to which they have virtualized other components of their infrastructure, most notably their network. That is one of the reasons why a subsequent section of The Report will discuss network virtualization.

The Culture of Cloud Computing

The rest of The Report will discuss the networking technologies that enable cloud computing. However, as much as cloud computing is about technologies it is also about changing the culture of the IT organization. One such cultural shift was described in the preceding subsection entitled "The Goal of Cloud Computing".

To put this cultural shift into perspective, it is important to realize that it is implicit in the traditional IT culture to implement ongoing enhancements to make the network and the IT services that are delivered over the network, increasingly resilient. The adoption of cloud computing changes that model and as previously described, in some instances it is becoming acceptable for IT services to be delivered on a best effort basis. A clear indication of that change is the success of Salesforce.com. Salesforce.com has three million customers who use their solutions to support critical sales processes. Yet in spite of the importance of the application, in virtually all cases Salesforce.com will not give a customer an availability guarantee and since the application is typically accessed over the Internet, it doesn't come with an end-to-end performance guarantee.

One of the other cultural shifts that is associated with the adoption of cloud computing is that IT organizations become less of a provider of IT services and more of a broker of IT services. In the traditional IT environment, the IT organization is the primary provider of IT services. Part of the challenge that is associated with the IT organization being the primary provider of IT services is that sometimes the IT organization can't meet the needs of the business units in a timely fashion. In the past the way that business unit managers have dealt with this lack of support is by having their own shadow IT organization whereby the business unit managers have some people on their staff whose role is to provide the IT services that the business unit manager can't get from the IT organization[10]. In the current environment, public cloud providers often play the role of a shadow IT organization by providing a company's business unit managers services or functionality that they either can't get from their IT organization or they can't get in a timely manner. In some instances the IT function is in a position to stop the non-sanctioned use of public cloud computing once they find out about it. However, in many other instances they aren't.

Instead of trying to prevent business unit managers from acquiring public cloud services, a better role for an IT organization is to modify their traditional role of being the primary provider of IT services and to adopt a role in which they provide some IT services themselves and act as a broker between the company's business unit managers and cloud computing service providers for other services. In addition to contract negotiations, the IT organization can ensure that the acquired application or service doesn't create any compliance issues, can be integrated with other applications as needed, can scale, is cost effective and can be managed.

IT organizations provide considerable value by being the broker between the company's business unit managers and cloud computing service providers.

Another cultural change that is associated with the adoption of cloud computing is the implementation of more usage sensitive chargeback. Usage sensitive chargeback is not new. Many IT organizations, for example, allocate the cost of the organization's network to the company's business unit managers based on the consumption of that network by the business

[10] The data in Table 11 provides some insight into how often this occurs.

units. Since there has traditionally been a lot of overhead associated with usage sensitive chargeback, usage sensitive chargeback has only made sense in those situations in which the IT organization is in a position both to explain to the business unit managers in easily understood language, what they are paying for and to provide suggestions as to how the business unit managers can reduce their cost. In the current environment, roughly fifty percent of all IT organizations implement usage sensitive chargeback for at least some components of IT. However, relatively few implement it broadly. Input from the Survey Respondents indicates that over the next two years IT organizations will make increased use of usage sensitive chargeback. Most of this increased use will come from having the business unit managers pay the relevant cloud computing service providers for the services that their organization consumes. The movement to implement more usage sensitive chargeback over the next two years will not be dramatic because:

The culture of an IT organization changes very slowly.

About the Webtorials® Editorial/Analyst Division

The Webtorials® Editorial/Analyst Division, a joint venture of industry veterans Steven Taylor and Jim Metzler, is devoted to performing in-depth analysis and research in focused areas such as Metro Ethernet and MPLS, as well as in areas that cross the traditional functional boundaries of IT, such as Unified Communications and Application Delivery. The Editorial/Analyst Division's focus is on providing actionable insight through custom research with a forward looking viewpoint. Through reports that examine industry dynamics from both a demand and a supply perspective, the firm educates the marketplace both on emerging trends and the role that IT products, services and processes play in responding to those trends.

For more information and for additional Webtorials® Editorial/Analyst Division products, please contact Jim Metzler at jim@webtorials.com or Steven Taylor at taylor@webtorials.com.

Published by
Webtorials
Editorial/Analyst
Division
www.Webtorials.com

Division Cofounders:
Jim Metzler
jim@webtorials.com
Steven Taylor
taylor@webtorials.com

Professional Opinions Disclaimer

All information presented and opinions expressed in this publication represent the current opinions of the author(s) based on professional judgment and best available information at the time of the presentation. Consequently, the information is subject to change, and no liability for advice presented is assumed. Ultimate responsibility for choice of appropriate solutions remains with the reader.

The Application Fluent Data Center Fabric

Introduction

The rise of virtualization and cloud computing requires the selection of a best-of-breed data center switching solution as part of an enterprise's overall data center strategy. And at the heart of this strategy is the need to deliver a high quality user experience with new virtualized applications, including video, on new devices such as smart phones and tablets. However, the traditional 3-layer networks designed for a client/server communication model cannot meet the requirements of these new applications and devices, nor can it address the new requirements of virtualized servers and desktops.

Application Fluency for the Data Center

Resilient Architecture

- Simplified 10 & 40 GigE network with low latency and ready for 100 GigE
- Multi-path data center network extends between data center sites and to public cloud
- Supports definition of virtual data centers
- Ready for storage convergence with lossless Ethernet

Automatic Controls

- Application profiles ensure that the network is aware of application provisioning, security and QoS requirements
- The network will automatically sense virtual machine location and movement
- The network will automatically adjust to VM motion within and between data center sites

Streamlined Operations

- Applications are automatically provisioned
- Core switches automatically configure top of rack switches
- Converged management for data center network and virtual machine mobility
- Low power consumption

The Alcatel-Lucent Mesh

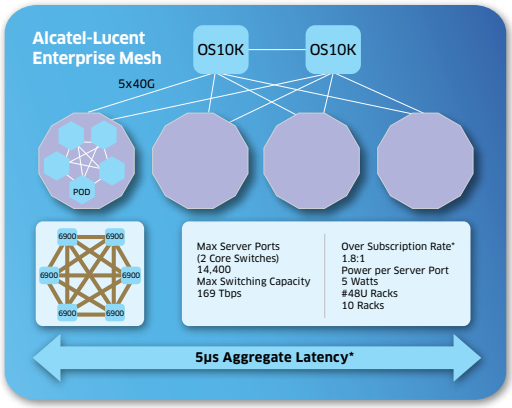

Alcatel-Lucent provides a unique Application Fluent approach to maximize the benefit from virtualization technologies for servers, the desktop, as well as the network. Alcatel-Lucent's application fluent data center fabric can scale from several hundred to over 14,000 server facing ports while keeping aggregate latency at 5ms, and can automatically adapt to virtual machine movement no matter which server virtualization platform is used.

The Alcatel-Lucent Virtual Network Profile (vNP), embedded in the Alcatel-Lucent Mesh, includes the critical information the fabric needs to understand each application, including provisioning requirements, security profiles, and expected quality of service levels. With this knowledge, the network can manage applications as services, including automatically discovering the location of each virtual machine, modifying the network configuration to follow virtual machine moves and providing an integrated view on visibility on VM movement and current location from a network perspective.

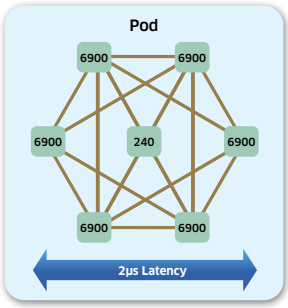

Application fluency in the corporate data center includes its transformation into a multi-site private cloud by extending layer 2 connectivity between data center sites and allowing for seamless delivery of public cloud-based services on the corporate network.

The Alcatel-Lucent Mesh enables enterprises to provide a high quality user experience with mission critical, real-time applications, and to improve agility in deploying new applications while significantly reducing data center costs.

Open Ecosystems and Market Success

Alcatel-Lucent Enterprise is committed to open standards, allowing enterprises to select best-of-breed suppliers for their complete data center solution: servers, storage, data center fabric, and data center interconnect.

- Winner: Best of Interop 2011 for Data Center Switching and Storage
- Data center ecosystem partners include Emulex, NetApp, VMware, Citrix, and QLogic
- Participant in IEEE sponsored Shortest Path Bridging interoperability test with Avaya, Huawei, Solana and Spirent
- Over 20 million Ethernet ports shipped

For More Information

Alcatel-Lucent Data Center Switching Solution
Alcatel-Lucent Application Fluent Networks
Alcatel-Lucent Enterprise

**Assuming Server to Server Traffic 70% within a Pod, 20% between Pods and 10% Via Core*

Alcatel·Lucent
Enterprise