The 2011 Cloud Networking Report

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

Produced by:

Webtorials

Part I: Executive Summary & The Emergence of Cloud Computing and Cloud Networking

A10 Networks
Alcatel·Lucent Enterprise
AVAYA The Power of We™
BlueCoat
ca
certeon Accelerate Your Business
CISCO
CITRIX
DELL Force10
exinda
Extreme networks Make Your Network Mobile
hp
IBM
Infoblox
ipanema Technologies
LSI
NETSCOUT
Packet Design
radware
riverbed
Silver Peak
TALARI NETWORKS
VISUAL NETWORK SYSTEMS
VYATTA Open Networking

Table of Contents

Executive Summary

The **2011 Cloud Networking 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.

The Emergence of Cloud Computing and Cloud Networking

Introduction and Forward to the 2011 Edition

The majority of IT organizations have either already adopted, or are in the process of evaluating the adoption of one or more classes of cloud computing. Gartner, for example, estimates that between 2010 and 2015 that enterprises will spend $112 billion cumulatively on Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS), combined[1].

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 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. 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 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 2010 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 challenges and solutions that are associated with cloud networking. A secondary goal of this report is to identify how IT organizations are currently approaching 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 and to the attendees of the Interop conferences. Throughout this report, those two groups of respondents will be respectively referred to as The Webtorials Respondents and The Interop Respondents. In some cases, the results of the surveys given to The Webtorials Respondents and The Interop Respondents will

[1] http://www.gartner.com/it/page.jsp?id=1389313
[2] http://www.webtorials.com/content/2010/12/2010-cloud.html

be compared to the results of surveys given to these two groups in 2010. The purpose of these comparisons is to quantify the ongoing changes that are occurring.

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 isn't a universally accepted definition of what is meant by cloud computing. This 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, most of the SLAs that are associated with public cloud computing services 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 SLA that includes a meaningful commitment to critical network performance metrics such as delay, jitter and packet loss.

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.

Figure 1: Distribution Models for Cloud-Based Solutions

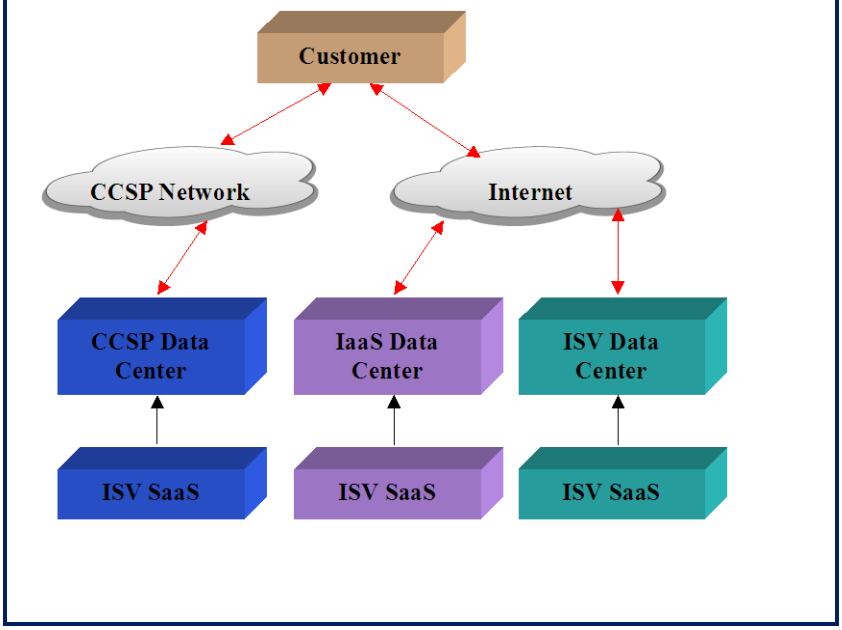

An approach to providing public cloud computing services that does lend itself to offering more meaningful SLAs is based on a CCSP providing these solutions to customers from the CCSP's data center and over a network that is provided by the CCSP and based on a technology such as MPLS.

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*. As previously noted there is not a litmus test to determine which characteristics have to be in a solution for the solution to be deemed to be a cloud computing solution. As a result, an IT organization that has centralized some of all of its servers into their data centers or into a collocation site, virtualized some of all of those servers, implemented some additional automation and that also moves virtual machines (VMs) between servers can reasonable claim that they have implemented a private cloud.

Public Cloud Computing

CCSPs that provide their services either over the public Internet or over other WAN services are offering a class of solution that is often referred to as the *public cloud* or *public cloud computing.* The research report entitled *Cloud Computing: A Reality Check and Guide to Risk Mitigation* presented the results of a survey in which 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 2.

Figure 2: The Drivers of Public Cloud Computing

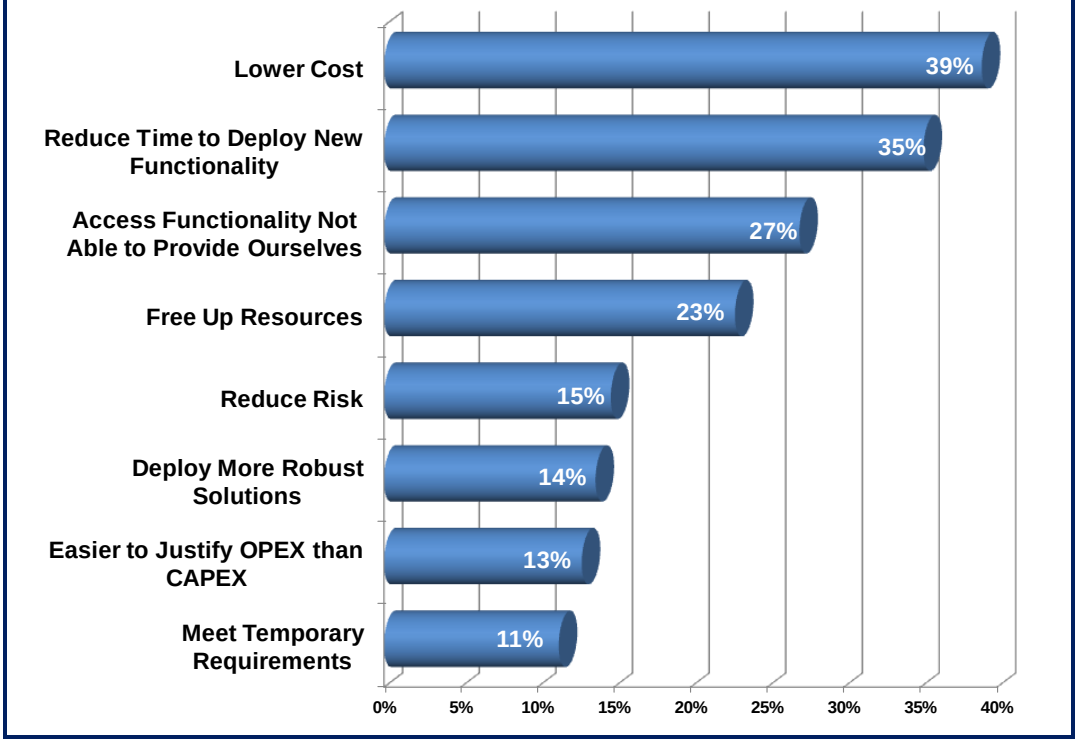

One of the observations that can be drawn from Figure 2 is that:

The primary factors that are driving the use of public cloud computing solutions are the same factors that drive any form of out-tasking.

That research report also pointed out that the primary factor that inhibits IT organizations from acquiring public cloud computing solutions is the concern over the security and confidentiality of data. Hence, it appears to be counter intuitive that almost 15% 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 respondent's 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.

In some cases, the use of a public cloud computing solution reduces risk.

As described in the report entitled A Guide for Understanding Cloud Computing[3], the two primary types of services provided by CCSPs are Software-as-a-Service (SaaS) and Infrastructure-as-a-Service (IaaS)[4].

Software-as-a-Service

According to IDC[5], the Software as a Service (SaaS) market had worldwide revenues of $13.1 billion in 2009 and is projected to reach $40.5 billion by 2014. 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 research report entitled *Cloud Computing: A Reality Check and Guide to Risk Mitigation*[6] reported on the results of a survey in which the survey respondents were asked about their company's use of SaaS-based applications. Figure 3 shows the percentage of respondents whose company either currently acquires, or is likely to acquire within the next year, various categories of applications from a SaaS provider.

[3] http://www.webtorials.com/content/2009/11/a-guide-for-understanding-cloud-computing.html
[4] A third form of service provided by a CCSP is Platform-as-a-Service (PaaS). Because it hasn't been widely adopted, it will not be included in this report.
[5] http://www.businesswire.com/news/home/20100726005135/en/SaaS-Revenue-Grow-Times-Faster-Traditional-Packaged
[6] http://www.webtorials.com/content/2009/12/cloud-computing-a-reality-check-guide-to-risk-mitigation.html

Figure 3: Popular Categories of SaaS-Based Applications

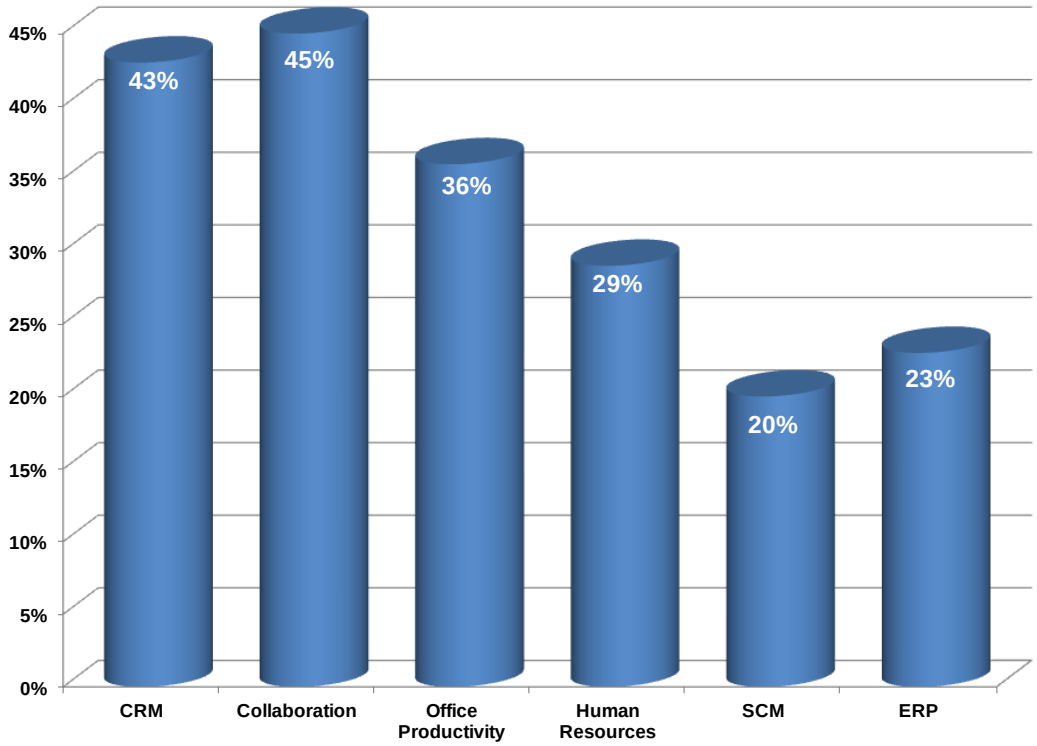

The functionality provided by each of the six categories of applications listed in Figure 3 can be quite extensive and is sometimes overlapping. ERP, for example, can encompass myriad functionality including product lifecycle management, supply chain management (e.g. Purchasing, Manufacturing and Distribution), warehouse management, customer relationship management (CRM), sales order processing, online sales, financials, human resources, and decision support systems.

For each category of application shown in Figure 3, there are tens, and sometimes hundreds, of SaaS-based solutions currently available[7]. Table 1 contains a listing of some representative SaaS providers for each category.

Table 1: Representative SaaS Providers

CRM	Collaboration	Office Productivity	Human Resources	SCM	ERP
Salesforce.com	WebEx	Google Docs	Subscribe-HR	ICON-SCM	SAP
NetSuite	Zoho	Microsoft's Office Web Apps	ThinMind	E2open	Workday
Update	clarizen	feng office	Greytip Online	Northrop Grumman	Lawson Software

[7] http://www.saas-showplace.com/saasproviderdirectory/saasapplicationcategory.html

One of the key challenges facing IT organizations that use SaaS-based applications is improving the performance, management and security of those applications.

Infrastructure as a Service (IaaS)

Over the last few years, IaaS solutions have been comprised primarily of the basic compute and storage resources that are required to run applications. 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 (recently acquired by Time Warner), Savvis (recently acquired by Century Link), Terremark (recently 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[8] estimates that the IaaS market will grow from $3.7 billion in 2011 to $10.5 billion in 2014.

Table 2 provides a high level overview of some of the services offered by IaaS vendors. The data in Table 2 is for illustration purposes only. That follows because it is extremely difficult, if not impossible, to correctly summarize in a table the intricate details of an IaaS solution; e.g., how the solution is priced, the SLAs that are provided and the remedies that exist for when the SLAs are not met. For example, consider the availability of an IaaS solution. On the surface, availability appears to be a well-understood concept. In fact, vendors often have differing definitions of what constitutes an outage and hence, what constitutes availability. For example, within Amazon's EC2 offering an outage is considered to have occurred only when an instance[9] is off line for 5 minutes and a replacement instance cannot be launched from another Availability Zone[10] within Amazon's geographical region. Not all IaaS providers have a similar definition of availability.

Table 2: Representative IaaS Providers

	Amazon AWS	RackSpace	GoGrid
Cloud Server (Virtual Machine (VM) with 2-4 vCPUs and ~8 GB RAM)	34¢/hour	40¢/hour	40¢-$1.53//hour *
Data Transfer	In 10¢/GB Out 15¢/GB	In 8¢/GB Out 18¢/GB	In free Out 7-29¢/GB
Load Balancer	2.5¢//hour 0.8¢/GB in/out	1.5¢/hour/LB 1.5¢/hour/100 connections	Included with server

8 http://www.gas.com/company/data-quality-news/iaas_market_to_record_strong_growth__7178.htm
9 http://aws.amazon.com/ec2/instance-types/
10 http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/index.html?Welcome.html

Table 2: Representative IaaS Providers			
	Amazon AWS	**RackSpace**	**GoGrid**
VM Storage	(Elastic Block Store) 10₡/GB/month 10₡/million I/O requests/month	320 GB included with server	Included with server 400GB per 8 GB RAM
Cloud Storage	5.5-14₡/GB/month	15₡/GB/month	15₡/GB/month over 10 GB
Hypervisors	Xen plus VMware import	Xen (Linux) CitrixXenServer (Windows)	Xen
Server availability SLA	99.95%	100%	100%
Server SLA Remedy	10% of monthly charge/incident	5% of monthly charge/30 minutes downtime	100x hourly rate for downtime period

*=includes O/S licenses and some other items and depends on a variety of pre-payment plans

Table 2 illustrates that:

There are significant differences amongst the solutions offered by IaaS providers, especially when it comes to the SLAs they offer.

It is important to realize that the value of an availability SLA is only partially captured by the number of 9s it features. A number of factors can cause an SLA that promises four or more 9s of availability to become notably less meaningful. One such factor was previously mentioned – how the vendor defines what constitutes an outage. Another such factor is the remedy that the vendor provides for those instances in which the service it offers doesn't achieve the promised availability. In those cases in which the SLA remedies are weak, the IaaS provider can provide a fairly low level of availability and not suffer a significant loss of revenue. This can have the affect of minimizing the incentive that the vendor has to take the necessary steps to ensure high availability. A related factor is the degree of difficulty that an IT organization has in gathering the documentation that is required to establish that the service was unavailable and to apply for the service credits that are specified in the SLA. As the difficulty of this process increases, the meaningfulness of the SLA decreases.

Insight into the availability of a number of IaaS solutions was provided by Cedexis at the Interop conference in May, 2011[11]. Cedexis presented data that represented roughly 17 billion measurements that were taken between March 15, 2011 and April 15 2011. As shown in Figure 4, none of the IaaS providers that were monitored delivered availability that was greater than 95%,

[11] Comparing Public Clouds: The State of On-Demand Performance, Marty Kagan, President and Co-Founder, Cedexis

Figure 4: Availability of Server Instances at Various IaaS Providers
(source: Cedexis)

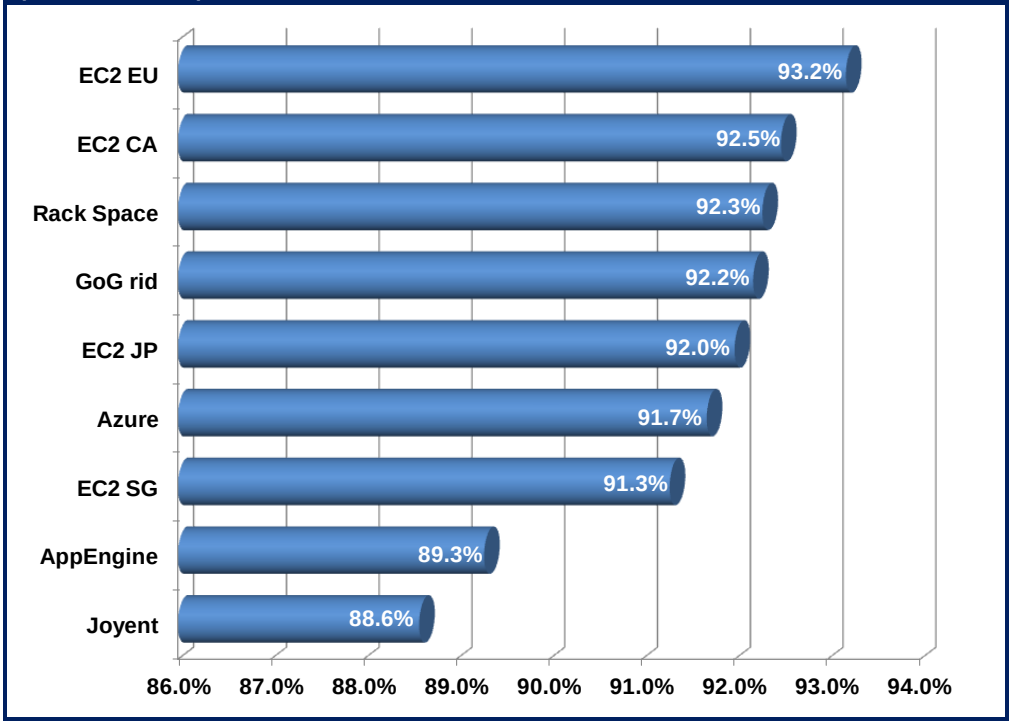

Figure 4 illustrates that:

The availability of IaaS solutions can vary widely.

In addition, similar to the situation with SaaS-based applications,

One of the key challenges facing IT organizations that use IaaS-based solutions is improving the performance, management and security of those solutions.

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[12], "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 5**, 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 5: Cloud Balancing

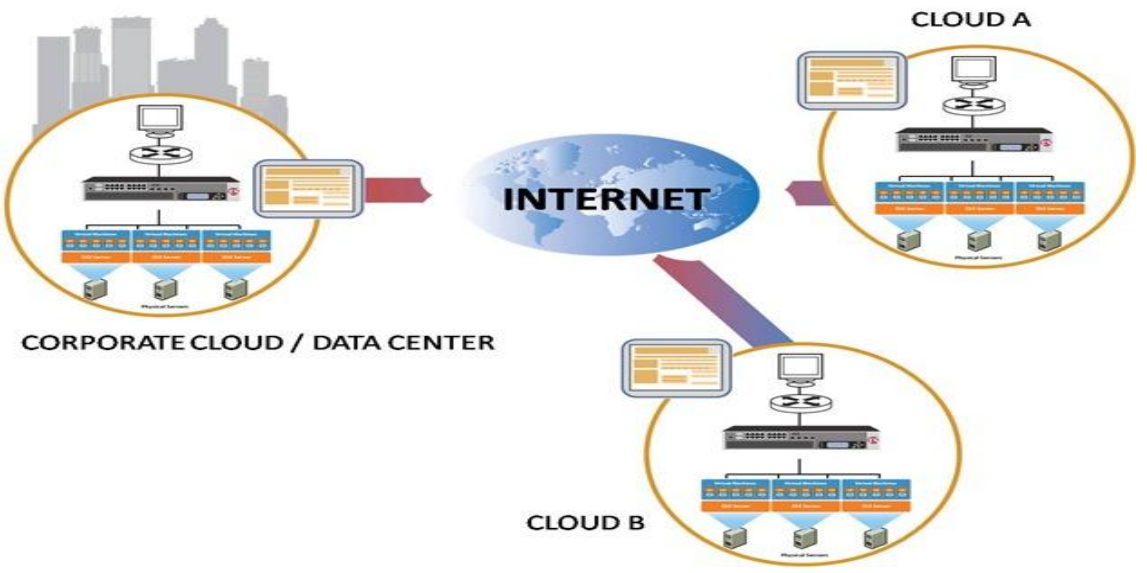

The goal of a GSLB solution is to support high availability and maximum performance. In order to do this, a GSLB solution typically makes routing decisions based on criteria such as the application response time or the total capacity of the data center. A cloud balancing solution may well have as a goal 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:

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

- 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

As shown in Figure 3, with the exception of collaboration, the applications that organizations have acquired from CCSPs have typically been enterprise applications such as CRM. As was previously mentioned, over the last few years IaaS solutions have been comprised primarily of the basic compute and storage resources that are required to run applications. 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. This new class of solutions will be referred to in this report as Cloud Networking Services (CNS).

A recent research report entitled Cloud Networking Services[13] presented the results of a survey in which the survey respondents were asked to indicate how likely it was over the next year that their company would acquire a CNS. Their responses are shown in Table 3.

Table 3: Interest in Cloud Networking Services

	Will Not Happen	Might Happen	50/50 Chance	Will Likely Happen	Will Happen
VoIP	34.3%	17.5%	12.6%	15.4%	20.3%
Unified Communications	26.1%	26.8%	16.9%	14.8%	15.5%
Network and Application Optimization	33.8%	22.1%	14.7%	14.0%	15.4%
Disaster Recovery	30.8%	23.8%	20.0%	11.5%	13.8%
Security	39.0%	16.9%	16.9%	14.0%	13.2%
Network Management	38.8%	26.6%	7.2%	17.3%	10.1%
Application Performance Management	35.8%	28.4%	15.7%	12.7%	7.5%
Virtual Desktops	40.7%	24.4%	18.5%	9.6%	6.7%
High Performance Computing	41.9%	24.8%	16.3%	10.1%	7.0%

The data in Table 3 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 each of the services listed in the top six rows of Table 3 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.

As noted, the two primary forms of public cloud computing are SaaS and IaaS. It would be possible to make a technical argument that at least some CNS solutions are SaaS solutions and that some others are IaaS solutions. While technology is one way to classify CNS solutions, a more compelling way is to look at how the typical IT organization is structured. Most IT organizations have an applications organization whose primary role is to develop, acquire and maintain enterprise applications such as CRM, ERP and SCM. Most IT organizations also have an infrastructure organization whose primary role is to provide, manage, secure and optimize the networks and servers that support the applications that enable the company's business processes. In most cases, services such as voice, collaboration, disaster recovery, management, security, optimization and virtual desktops are provided by the infrastructure organization – not the applications organization. Because of the way that IT organizations are

[13] http://www.webtorials.com/content/2011/09/2011-cloud-networking-services.html

typically structured, throughout this report CNS solutions will be considered to be the next wave of IaaS solutions.

Since CNS solutions are just one more form of public cloud computing, when evaluating these solutions IT organizations also need to understand the degree to which these solutions overcome the factors that impede the use of any public cloud computing solution. As previously mentioned, concerns about security is the primary impediment to the adoption of public cloud computing solutions and hence 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 (Figure 2) that drive IT organizations to use public cloud computing solutions. The primary benefit of using a public cloud computing solution is 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. The second most important 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. One such component is networking. By implementing routing software that runs on top of the most common hypervisors, CCSPs increase their ability to quickly provision and configure capacity. This approach to providing routing functionality also maps more closely to the usage sensitive pricing that most CCSPs offer.

The Culture of Cloud Computing

The rest of this 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 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. 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 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 Webtorials 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.

A10 Networks

The Fastest Growing Application Networking Company

64-bit AX Series

Application Delivery

- Advanced Application Delivery Controller (ADC)
- New Generation Server Load Balancer (SLB)

IPv6 Migration

- LSN, CGN, NAT444
- DS-Lite, 6rd
- NAT64 & DNS64
- SLB-PT, SLB-64

Cloud Computing & Virtualization

- SoftAX & AX-V
- AX Virtual Chassis
- AX Virtualization (Multi-tenancy)

Advanced Core Operating System (ACOS)

AX Series Advantage

- All inclusive pricing for hardware appliances, no performance or feature licenses
- Most scalable appliances in the market with unique modern 64-bit ACOS, solid-state drives (SSD) and multiple hardware acceleration ASICs
- Faster application inspection with aFleX TCL rules
- aXAPI for custom management

Application Solutions

The AX Series increases scalability, availability and security for enterprise applications. Visit A10's web site for deployment guides, customer usage scenarios and to participate in the Application Delivery Community.

Blackboard · Microsoft · Microsoft Exchange · Microsoft Lync · IMPERVA · Infoblox · Juniper Networks · ORACLE · vmware READY

Cloud Networking – the not-so-quiet revolution

Avaya's vision for the Enterprise calls for a new level of synergy between people, the collaborative real-time applications they use, and the underlying network. A key building block for this vision is the foundational networking technology. As real-time communications continue the evolution to IP, the data network becomes completely integrated into the delivery of communications-enabled business services and mission critical business applications.

Avaya Networking provides advanced enterprise-class reliability, performance, and security that organizations throughout the world depend on to run their businesses. Because our solutions are streamlined to better utilize and manage networking resources, an Avaya data network can uniquely deliver both mission critical dependability and superior return on investment.

Virtualization within the Data Center is now taken for granted, with some declaring that 'Cloud Computing' will be the choice of most enterprises and that applications and information will become commodities. Experience has proved one thing; the Data Center of the future cannot be built on the technology of the past. General-purpose products, outmoded techniques, and legacy designs cannot be re-packaged as 'Data Center-ready'. The industry will take the best and leave the rest. Ethernet is readily available, cost-effective, extensible, and – as the 40/100 Gigabit developments prove – seamlessly without limitation of scale, however many of the underlying deployment methodologies are no longer an option.

Today's Enterprise network must be flatter, less tree centric, and able to support sustained east-west flows between multiple servers, in addition to traditional client/server transactions. Factors driving the transformation of enterprise networks include the transition to composite application architectures, an adoption of business operations intelligence applications (based on communications-enabled business processes and complex-event processing), and an increase in live virtual machine migrations. With each factor creating a unique challenge for the Data Center network, ranging between sensitivity to latency and loss, increased traffic levels (background noise), and risk of extended saturation of the common I/O connection, what's required is an agile, high-performance, latency-optimized networking solution that delivers exceptionally high performance.

To support the transition to a multi-dimensional environment the underlying network also needs to change. Provisioning needs to be simpler, and availability and performance need to scale seamlessly. Empowering a truly commoditized approach to service delivery requires a solution that is characterized by simplification, and a standards-based approach will help ensure an open architecture that avoids costly or inflexible lock-in.

Avaya is able to clearly demonstrate a set of differentiating benefits:

- Reduction in the configuration burden by up to 25X over the techniques traditionally implemented in large Data Centers
- Simplification of application implementation and number of devices affected, thereby reducing chances for configuration errors; it's these human-errors that account for up to 40% of all network downtime
- Data Center resiliency that delivers millisecond convergence times during failover and recovery

Enabling Enterprises to build a Private Cloud infrastructure that is extensible from Data Center to Campus and ultimately to the Branch Office; end-to-end network virtualization is an important element of the Avaya Virtual Enterprise Network Architecture (VENA). Designed for next-generation networking, Avaya VENA is a flexible solution that can be tailored to fit current business needs while providing a smooth migration path that accommodates business evolution. Addressing crucial Data Center requirements, Avaya VENA creates self-aware network infrastructures that simplify the logical provisioning of network services and provide the components required to create an Ethernet fabric featuring active/active connectivity for all attached servers, and service-orientated networking from Top-of-Rack to Core. Chief among Avaya VENA components are our innovative Switch Clustering and the IEEE's 802.1aq Shortest Path Bridging virtualization technologies – enhanced with enterprise-friendly, Layer 3 functionality, authenticated network access, and a network management toolset that simplifies deployment, monitoring, and troubleshooting.

Avaya, uniquely positioned based on decades of networking experience, helps ensure that the transition to the next-generation of fabric-based infrastructure is low-risk, seamless, and evolutionary. Avaya's pedigree of proven, ground-breaking innovation delivers a truly fit-for-purpose Cloud-ready solution that encompasses both the Data Center and the Campus; ensuring simplified yet optimized end-to-end connectivity between users and their content.

Cisco Unified Fabric

Converged. Scalable. Intelligent.

Cisco Unified Fabric is a flexible, innovative, and proven platform for physical, virtual or cloud deployments. It provides the foundational connectivity within and across data centers so resources are highly available wherever and whenever they are needed.

A key building block for cloud-based environments and virtualized data centers, the Cisco Unified Fabric brings unmatched architectural flexibility and scale to meet the diverse requirements of massively scalable data centers, bare-metal infrastructures, high performance and big data applications.

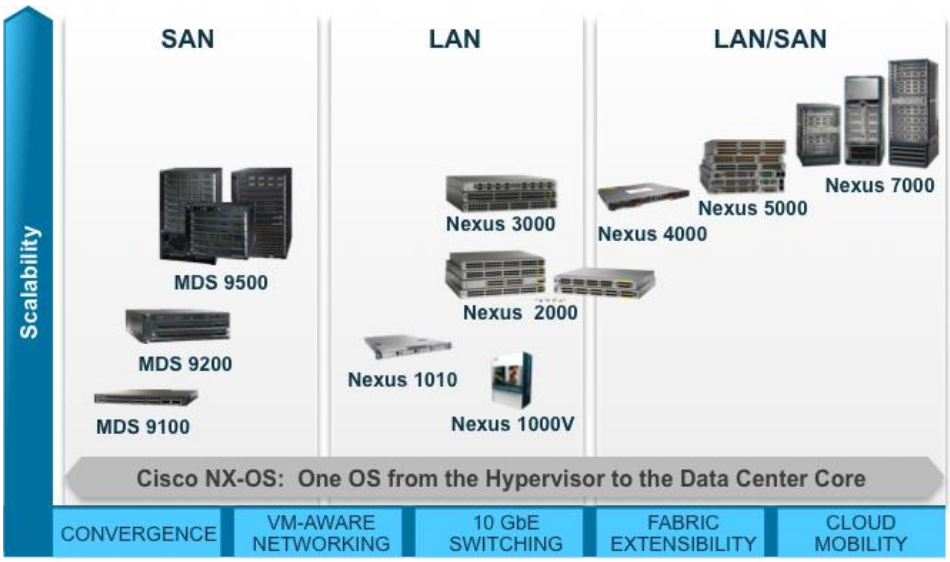

- Revolutionary fabric scale with over twelve thousand 10 GbE server connectivity with Cisco Nexus
- Highest 10Gb Ethernet density in the industry with Cisco Nexus 7000
- High performance and ultra-low latency networking at scale with Cisco Nexus
- Network services delivered in virtual and physical form factors with Cisco ASA, ASA 1000v, WAAS, vWAAS, VSG and more
- Virtual networking from the hypervisor layer on up with Cisco Nexus 1000v, VSS, VDC, and more
- High availability within and across devices with ISSU, VSS, vPC, and more.
- Flattened and scalable networking at Layer 2 and Layer 3 with Cisco FabricPath, TRILL, L3 ECMP, and more
- Overcome the challenges of expanding networks across locations and the limitations of network segmentation at scale with Cisco OTV, LISP, VXLAN, and more
- Unified operational, control, and management paradigms across the entire fabric with Cisco NX-OS, DCNM and open APIs
- Converged networking to carry every kind of traffic on a single fabric with DCB and FCoE with Cisco Nexus and MDS

Cisco Unified Fabric is a flexible, innovative, and proven platform for physical, virtual or cloud deployments with a non-disruptive, evolutionary approach to create future-proofed, service- and cloud-ready data centers and prevent 'rip and replace' for existing data centers. For more info: http://www.cisco.com/go/unifiedfabric

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Consolidation and Cloud Computing Without Compromise

Citrix virtualization and cloud networking solutions accelerate, optimize, and secure application and service delivery from both the enterprise datacenter and the Cloud.

Starting Point

Server, storage, and other virtualization technologies are enabling organizations to consolidate infrastructure and transform to a dynamic, cloud computing model of IT service delivery. The result is a substantial reduction in capital and operating costs, *plus* a highly scalable and agile approach to meeting the computing needs of the business.

Next Step

To maximize gains, organizations should also extend virtualization and cloud computing principles to crucial networking components, including application delivery controllers (ADCs). Taking advantage of the flexibility and cost effectiveness of virtual appliance ADCs to more thoroughly ensure the performance, availability, and security of business-critical applications and services is a significant next step. Ideally, though, it should also be possible to consolidate numerous standalone ADCs to help reduce datacenter complexity and further control costs.

No Compromises

A new service delivery platform from Citrix, NetScaler SDX addresses this need by enabling multiple, independent instances of the NetScaler ADC to run on a single physical appliance. With NetScaler SDX, organizations gain the opportunity to reduce ADC footprint and total cost of ownership by maximizing consolidation of standalone ADC devices, across both different applications (i.e., horizontally) and different network zones (i.e., vertically).

> NetScaler SDX is a true multi-tenant platform that enables consolidation of core data center services. It delivers full functionality and meets the most demanding availability, security and performance SLAs.

Unique NetScaler Strengths

- **High consolidation density** – Up to 40 ADC instances can run independently on a single NetScaler SDX platform —more than double what competitors offer.
- **Complete isolation of ADC resources** – All critical system resources, including memory, CPU and SSL processing capacity, are assigned to individual NetScaler instances. Performance SLAs can thus be maintained on a per tenant basis.
- **Full ADC functionality** – Support for 100 percent of the NetScaler application delivery capabilities enables consolidation of all existing ADC deployments without any policy constraints or compromises.
- **Pay-As-You-Grow scalability** – An innovative, software-based Pay-As-You-Grow option provides essential elasticity, enabling organizations to scale performance and capacity on-demand without the need for expensive hardware upgrades.

For more information and a free NetScaler VPX download, please visit www.citrix.com/netscaler.

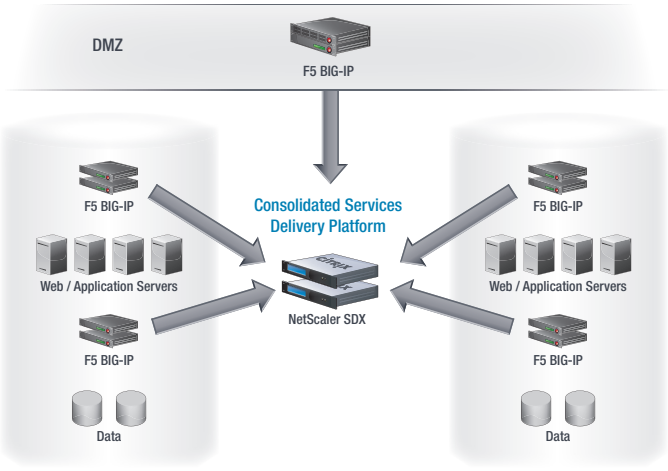

Network Services Consolidation

The power to do more

Unlock the full potential of your data center.

Servers
Intelligent compute architecture

Storage
Fluid data architecture

Networking
Open Cloud Networking

Application layer
Virtual Integrated System

Open Cloud Networking from Dell Force10 gives you the ability to achieve new levels of flexibility, performance and automation without having to rip and replace systems.

When we say "open" we mean it.

Open Architectures - build upon your existing infrastructure
Open Automation - simplify data tasks and vm management
Open Ecosystems - maximum choice and true architectural freedom

Is your network "open" for business? Find out more at **dell.com/OCN**

Scan this tag to unlock the full potential of your data center Dell Force10 Open Cloud Networking.

Get a free QR tag reader at your favorite app store.

Dell Force10
Data center switching, Top-of-rack, and next generation distributed core networks

HP FLEXNETWORK ARCHITECTURE

Meet the stringent performance, security, and agility demands of cloud computing

Enterprises are turning to the cloud to accelerate business innovation, improve business agility, and contain costs. Cloud computing reshapes the way applications are deployed and consumed and influences data center network designs. HP helps organizations build unified, virtualization-optimized networks that meet the rigorous performance, scalability, availability, and agility demands of the cloud.

HP FlexNetwork—an architectural blueprint for cloud-optimized networking

HP FlexNetwork architecture—HP's blueprint for cloud-optimized networking—lets enterprises securely deploy and centrally orchestrate cloud-optimized architectures that scale from the data center to the network edge.

HP FlexNetwork: Industry's only network architecture converging data center, campus, and branch office

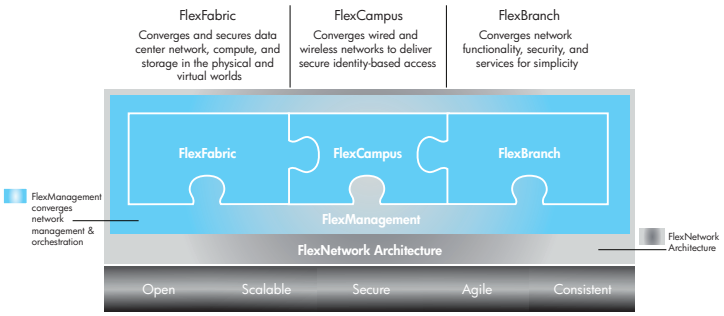

HP FlexFabric and **HP FlexCampus** enable the construction of flat, low-latency data center and campus networks with fewer layers, less equipment and cabling, and greater port densities.

HP FlexBranch includes comprehensive WAN optimization and routing solutions for delivering dynamic cloud-based services to geographically distributed enterprises.

HP FlexManagement provides a unified view into the virtual and physical network infrastructure, which accelerates application and service delivery, simplifies operations and management, and boosts network availability.

HP CloudSystem—a single platform for private, public, and hybrid clouds

HP CloudSystem is the industry's most complete, integrated, and open system for building and managing cloud services. Based on proven, market-leading HP Cloud Service Automation and Converged Infrastructure, HP CloudSystem combines servers, storage, networking, and security together with automated system and hybrid service delivery management. It enables organizations to build, manage, and consume cloud services across private clouds, public clouds, and traditional IT environments—without having to know, or care, whether those services come from HP CloudSystem's own "on-premises" resources or from the public domain.

HP CloudSystem and HP FlexNetwork networking solutions deliver:

- **Flatter and more efficient data center networks** with fewer layers, less equipment and cabling, and greater port densities
- **High-performance, low-latency intra-data-center connectivity** for virtual machine migration and bandwidth-intensive server-to-server communications
- **Virtualization-aware security** to partition multi-tenant environments and isolate virtual resources and intra-server communications flows
- **Optimal WAN performance** for the highest-quality end-user and application experiences and most efficient use of WAN resources
- **Unified administration and service orchestration** to accelerate the delivery of cloud-based applications and services
- **Multi-site, multi-vendor management** to connect and control thousands of physical and virtual resources from a single pane of glass

For more information

HP Networking Solutions: **www.hp.com/networking**

HP Cloud Solutions: **www.hp.com/go/cloud**

HP CloudSystem: **www.hp.com/go/cloudsystem**

To learn more about how HP can help you build a cloud-optimized data center network, please contact your HP account manager or reseller.

Created November 2011

EXECUTIVE VIEWPOINT

Maximizing Your IT Resources

Network Service Automation Rightsizes IT Staff and Delivers "Time to Value"

Steve Nye

EXECUTIVE VICE PRESIDENT, PRODUCT STRATEGY AND CORPORATE DEVELOPMENT, INFOBLOX, INC.

Steve Nye is the Executive Vice President of Product Strategy and Corporate Development for Infoblox, Inc. He is responsible for formulating the Company's longer-term strategy for portfolio and market expansion. Within his organization he directs all product management, marketing and business development activities. He oversees corporate development, which includes strategic alliances, both technical and marketing, as well as M&A activity.

Infoblox®

www.infoblox.com
1-866-463-6256
info@infoblox.com

WHAT IS THE BIGGEST CHALLENGE YOU SEE DRIVING IT DEPARTMENTS THESE DAYS?

Our customers and business partners say complexity is on the rise, which is putting more demands on IT to respond faster to business changes. However, because their budgets and staff are constrained, most companies cannot move quickly. They need help with scaling in an environment in which technology is moving faster than IT talent. We think new solutions that help manage the growing chaos surrounding IP initiatives will help increase network availability by reducing errors or delays in rolling out new services.

WHAT IS THE IMPACT OF VIRTUALIZATION ON NETWORK STAFF?

Virtualization breaks the traditional "one server, one application" architecture, and that creates new management challenges. For example, troubleshooting and seeing which virtual machine is connected to which port have become more difficult. Businesses need new discovery and visualization tools that automatically collect configuration information and automate repetitive and high-responserate chores such as assigning IP addresses and server names in a virtual environment. The task of issuing IP addresses and names for virtual machines should happen just as fast as a virtual machine can be provisioned. The network team in a virtualized environment must be as dynamic as the server team's ability to provision new systems. This type of automation is a critical part of any private cloud strategy.

HOW DOES THE INFLUX OF NEW MOBILE CONSUMER DEVICES CORRELATE WITH THE NEED FOR MORE NETWORK AUTOMATION?

IT managers are often not informed when new mobile devices come into the company. Employees bring them to work, or business units buy new systems because they do not want to wait for funds to be allocated to fulfill a critical business need. The IT department needs to know what is being attached to the enterprise network, because the impact of these devices can be significant. This shift to a more mobile and dynamic computing environment puts a strain on mission-critical network services such as Domain Name Service (DNS). As a result, IT needs simple-to-use, intuitive tools that monitor network activity while proactively managing and securing connections from a single central console.

HOW DOES THE MOVEMENT TO IPV6 AFFECT NETWORK STAFF?

The migration has already begun. T-Mobile is delivering IPv6 support in its phones, and these new IPv6 devices still need to connect to IPv4 networks. In the past, address management was done on spreadsheets, but 128-bit-IPv6 addressing brings an entire new set of challenges. When you add virtualization and cloud to this challenge, managing IP addresses with just a spreadsheet becomes impossible. IT teams will need automated network services.

WHERE SHOULD A COMPANY START AND HOW CAN YOU GAUGE SUCCESS?

Automation is a new "big idea." To some, it means ripping and replacing—or making significant investments in professional services and/or integration work. At Infoblox, we strive to make automation compelling by demonstrating that we can make adoption simple. By using automation, companies can reduce a 40-step process to a few clicks of a mouse. As a result, companies can make huge productivity gains and save money—many of our customers see an immediate increase in network availability and savings of millions of dollars annually by embracing automation.

Once companies see such results, they can expand their use of these tools and dramatically increase IT staff productivity. Infoblox's heritage is in automating network services such as DNS and IP address management. We anticipate that both automation and next-generation network services will be key elements powering the next 10 years of IT.

nano|engine

Full application control at 10% of the cost

www.ipanematech.com

A unique technology that breaks the price/performance barrier to guarantee business application performance in branch offices

- For the first time it is possible to guarantee application performance with a device compatible with branch office constraints;
- The nano|engines fully integrate with the other components of Ipanema's ANS solution;
- Plug-and-Play devices, nano|engines are managed under SALSA;
- Real-time changes in network performance and each user's behavior are taken into account in real-time.

Algorithms embedded in the nano|engine automatically adapt to real-time changes as they happen on the network:

- Traffic from private data centers mixed with traffic from external public clouds;
- Hybrid networks combining MPLS and Internet;
- Unified Communications branch-to-branch flows;
- Virtual desktops and rich media delivery…

The nano|engine's ability to guarantee application performance at the branch maximizes productivity, prevents brownouts and protects the business.

Ultra compact **nano|engine** appliances are tailored for providing full application control with unmatched performance/price ratio in broadband branch offices.

The **nano|engine** devices target broadband branch offices and provide:

- Application aware, **per connection Control and dynamic QoS** for public and private application flows to guarantee an excellent and stable Quality of Experience to each user;
- **End-to-end visibility** of application performance of each flow with comprehensive KPIs and application quality scores;
- **Dynamic WAN path selection** among up to 3 networks for optimized control of multi-attached branches, local Internet breakouts and hybrid networks.

Self-managed, nano|engines are installed at the edge locations of the WAN, typically between the CPE router and branch office LAN. Fully "Plug and Play," nano|engines require no on-site configuration. They operate under control of the central management software, SALSA. Customers simply need to plug the nano in, and configuration and provisioning are managed by SALSA.

The nano|engine family fits particularly well in B to C sectors like retail, finance and hospitality, where slow response times to access customer data or delays in processing an order lead to customer dissatisfaction and loss of productivity. Nano|engines' ability to guarantee application performance prevents any brownouts and protects the business.

The nano|2 addresses branch offices with up to 20 users and 4 Mbps while the nano|5 targets branch offices with up to 50 users and 20 Mbps.

Packet Design Solutions:

Packet Design's IP routing and traffic analysis solutions empower network management best practices in the world's largest and most critical enterprise, Service Provider and Government OSPF, IS-IS, BGP, EIGRP and RFC2547bis MPLS VPN networks, enabling network managers to maximize network assets, streamline network operations, and increase application and service up-time.

Packet Design

Route Explorer: Industry-Leading Route Analytics Solution

Optimize IP Networks with Route Explorer

- Gain visibility into the root cause of a signification percentage of application performance problems.
- Prevent costly misconfigurations
- Ensure network resiliency
- Increase IT's accuracy, confidence and responsiveness
- Speed troubleshooting of the hardest IP problems
- Empower routing operations best practices
- Complement change control processes with real-time validation of routing behavior
- Regain network visibility across outsourced MPLS VPN WANs

Deployed in the world's largest IP networks

400+ of the world's largest enterprises, service providers, government and military agencies and educational institutions use Packet Design's route analytics technology to optimize their IP networks.

Overview of Route Explorer

Route Explorer works by passively monitoring the routing protocol exchanges (e.g. OSPF, EIGRP, IS-IS, BGP, RFC2547bis MPLS VPNs) between routers on the network, then computing a real-time, network wide topology that can be visualized, analyzed and serve as the basis for actionable alerts and reports. This approach provides the most accurate, real-time view of how the network is directing traffic, even across MPLS VPNs. Unstable routes and other anomalies – undetectable by SNMP-based management tools because they are not device-specific problems – are immediately visible. As the network-wide topology is monitored and updated, Route Explorer records every routing event in a local data store. An animated historical playback feature lets the operator diagnose inconsistent and hard-to-detect problems by "rewinding" the network to a previous point in time. Histograms displaying past routing activity allow the network engineer to quickly go back to the time when a specific problem occurred, while letting them step through individual routing events to discover the root cause of the problem. Engineers can model failure scenarios and routing metric changes on the as-running network topology. Traps and alerts allow integration with existing network management solutions. Route Explorer appears to the network simply as another router, though it forwards no traffic and is neither a bottleneck or failure point. Since it works by monitoring the routing control plane, it does not poll any devices and adds no overhead to the network. A single appliance can support any size IP network, no matter how large or highly subdivided into separate areas.

Traffic Explorer: Network-Wide, Integrated Traffic and Route Analysis and Modeling Solution

Optimize IP Networks with Traffic Explorer

- Monitor critical traffic dynamics across all IP network links
- Operational planning and modeling based on real-time, network-wide routing and traffic intelligence
- IGP and BGP-aware peering and transit analysis
- MPLS VPN service network traffic analysis
- Network-wide and site to site traffic analysis for enterprise networks utilizing MPLS VPN WANs
- Visualize impact of routing failures/changes on traffic
- Departmental traffic usage and accounting
- Network-wide capacity planning
- Enhance change control processes with real-time validation of routing and traffic behavior

Traffic Explorer Architecture:

Traffic Explorer consists of three components:

- **Flow Recorders:** Collect Netflow information gathered from key traffic source points and summarize traffic flows based on routable network addresses received from Route Explorer
- **Flow Analyzer:** Aggregates summarized flow information from Flow Recorders, and calculates traffic distribution and link utilization across all routes and links on the network. Stores replayable traffic history
- **Modeling Engine:** Provides a full suite of monitoring, alerting, analysis, and modeling capabilities

Traffic Explorer Applications

Forensic Troubleshooting: Traffic Explorer improves application delivery by speeding troubleshooting with a complete routing and traffic forensic history.

Strengthened Change Management: Traffic Explorer greatly increases the accuracy of change management Processes by allowing engineers to model planned changes and see how the entire network's behavior will change, such as if there will be any congestion arising at any Class of Service.

Network-Wide Capacity Planning: Using its recorded, highly accurate history of actual routing and traffic changes over time, Traffic Explorer allows engineers to easily perform utilization trending on a variety of bases, such as per link, CoS, or VPN customer. Traffic Explorer ensures application performance and optimizes capital spending by increasing the accuracy of network planning.

Disaster Recovery Planning: Traffic Explorer can simulate link failure scenarios and analyze continuity of secondary routes and utilization of secondary and network-wide links.

Overview of Traffic Explorer

Traffic Explorer is the first solution to combine real-time, integrated routing and traffic monitoring and analysis, with "what-if" modeling capabilities. Unlike previous traffic analysis tools that only provide localized, link by link traffic visibility, Traffic Explorer's knowledge of IP routing enables visibility into network-wide routing and traffic behavior. Powerful "what-if" modeling capabilities empower network managers with new options for optimizing network service delivery. Traffic Explorer delivers the industry's only integrated analysis of network-wide routing and traffic dynamics. Standard reports and threshold-based alerts help engineers track significant routing and utilization changes in the network. An interactive topology map and deep, drill-down tabular views allow engineers to quickly perform root cause analysis of important network changes, including the routed path for any flow, network-wide traffic impact of any routing changes or failures, and the number of flows and hops affected. This information helps operators prioritize their response to those situations with the greatest impact on services. Traffic Explorer provides extensive "what-if" planning features to enhance ongoing network operations best practices. Traffic Explorer lets engineers model changes on the "as running" network, using the actual routed topology and traffic loads. Engineers can simulate a broad range of changes, such as adding or failing routers, interfaces and peerings; moving or changing prefixes; and adjusting IGP metrics, BGP policy configurations, link capacities or traffic loads. Simulating the affect of these changes on the actual network results in faster, more accurate network operations and optimal use of existing assets, leading to reduced capital and operational costs and enhance service delivery.

For more information, contact Packet Design at:

Web: http://www.packetdesign.com
Email: info@packetdesign.com
Phone: +1 408-490-1000

WANop Anywhere

Highest Capacity, Any Platform, Lowest Cost

Accelerate cloud adoption with the leader in data center class WAN optimization. Silver Peak delivers the highest-capacity, lowest-cost, and most comprehensive portfolio of virtual and physical WAN optimization appliances.

DOWNLOAD FOR FREE TODAY!

www.vx-xpress.com

Silver Peak

www.silver-peak.com

Leveraging Software-based Networking for End-to-End Cloud Infrastructures

VYATTA Open Networking

Industry Praise

Gartner

"Vyatta is certainly the headline name behind open-source networking"
Mark Fabbi
Gartner Inc. Analyst

Tolly

"Vyatta's open system running on standard hardware not only can scale better in enterprise and service provider edge deployments, but it also delivers enough headroom for expansion"
The Tolly Group

DELL

"Vyatta is able to provide the network services and secure connectivity that Dell GIS Cloud requires in a package that addresses the virtualization, commoditization and cost-benefit requirements of cloud computing."
Sanjay Basu
Dell Services

ZDNet

"Vyatta, has taught Cisco and the market that a networking box is really nothing but a computer with software."
Dana Blankenhorn
ZDNet

TechRepublic

"..anything you want to do with a standard Cisco router, you can do with Vyatta for the most part, and you don't have to worry about the various Cisco IOS licenses."
David Davis(CCIE, CCNA, CCNP)
TechRepublic

The Vyatta Network OS

The Vyatta network operating system is a scalable, integrated, enterprise-class networking solution that delivers advanced routing and network security functionality for physical, virtual and cloud networking environments. The Vyatta network OS includes dynamic routing, stateful firewall, VPN support, threat protection, traffic management and more in a package that is optimized to take advantage of multicore x86 processing power, common hypervisor platforms and emerging cloud architectures. All features are configured through Vyatta's familiar, networking-centric CLI, web-based GUI or third party management systems using the Vyatta Remote Access API.

The Power of Open Networking

Open and flexible networking is a requirement for today's evolving network. For the first time in two decades the industry is experiencing platform shifts that are dictating that networking be delivered as a software solution.

- **Datacenter Shifts:** Infrastructure shifting to the cloud requires flexible networking and security.
- **Virtualization:** Server and application consolidation requires virtualization-ready, platform independent application protection.
- **Edge Consolidation:** Special-purpose devices are giving way to multi-function, best-of-breed, multi-vendor integrated solutions.

The New Network Requirements

Features	Vyatta Network OS	Cisco IOS
Multifunction Layer 3+ (Routing, Firewall, VPN, IPS, Web Filter +)	Yes	Yes
Hardware Scalability	Seamless across x86 Cores	Cisco Limited
Software Performance	Unlimited	Platform Limited
Virtual Machine Availability	Yes (VMware, Xen, XenServer, KVM)	No
Open Management API	Yes	No
Integration into Custom Edge Devices	Yes	No
Cloud Readiness	Yes	No

The Vyatta Advantage

Network Right-Sizing: As a single network OS that scales up and down to meet your requirements, Vyatta puts the freedom in your hands to right-size your network as needed. Using readily available off-the-shelf systems and components, Vyatta breaks the "box lock" model of proprietary hardware vendors and allows you to drive as little or as much performance as your network requires.

Hardware Price/Performance: Standards have turned networking into a server workload. Today x86 hardware can easily outperform proprietary network devices at a small fraction of the cost. And the x86 universe means that faster systems at lower price are always on the horizon.

Virtualization: Vyatta gives you the optional power of running networking functions as a virtual machine. Whether it's VMs at the network edge or VMs in the cloud datacenter, Vyatta radically increases your infrastructure flexibility and produces a substantially higher ROI than proprietary solutions.

Deploying Vyatta in the Cloud: Common Use Cases:

As cloud moves from vision to reality, networking quickly moves to the front as a major impediment to meeting these major requirements. The reason is simple: traditional networking infrastructure has not been modernized the way server and storage infrastructure has been over the past decade. While the business promise of cloud computing is broad, there are a few basic enabling themes underlying an effective cloud design:

- Highly dynamic, on-demand infrastructure
- Granular service control levels
- High infrastructure utilization (multi-tenancy)
- Elastic pricing

CLOUD INFRASTRUCTURE

Designing a network infrastructure for cloud computing should deliver the same benefits as the rest of the cloud computing infrastructure in terms of lowered cost, flexibility, scalability and high utilization. Choosing a software-based network OS allows cloud providers to standardize entire infrastructures on x86 server hardware, leverage investments in hypervisor platforms and utilize a single network OS from the network edge to the customer for everything from high-performance BGP routing to per customer firewalling and LAN bridging.

SECURE CONNECTIVITY

Cloud users access their applications and data over the Internet, requiring every user's connection to be encrypted for security. Software-based networking is an exceptionally clean solution for this requirement. Within the cloud a new Vyatta VPN virtual machine can be started in moments, using a small fraction of an existing server. The high cost associated with acquiring and installing a unique physical device is completely eliminated, as is the requirement for more space, power and cooling. The customer can deploy the same software or virtual machine at each access location rapidly and with minimal expense, as a "secure cloud connector."

CLOUD ON-BOARDING - SECURE LAYER 2 BRIDGING

An often overlooked requirement in cloud computing is the need to enable customers to securely migrate data to the cloud from the enterprise datacenter. The Vyatta Network OS combines Layer 2 bridging and IPSec/GRE Tunneling functionality to deliver a cloud bridging solution which allows physically separate networks to securely communicate with each other over the internet as if they were on a single Ethernet network. This capability simplifies the migration of applications and physical servers between data centers, ensures continuity during a phased migration, and enables the moving of virtual machines between physical servers on physically separate networks.

VIRTUAL FIREWALLING

For IT architectures within a customer's own datacenters, it's common for firewalls to be deployed at various places to ensure data security for sensitive databases and transaction systems. Issues related to both internal security (HR databases, financial systems) and external compliance (credit cards, health care, etc) must be clearly addressed. Deploying these IT systems in a cloud environment increases this firewall requirement. The customer not only must firewall its sensitive systems as it had before, but also to ensure security in a multi-tenant environment using a shared connection to the public Internet. Using traditional networking would require a lot of traditional hardware firewalls at a high cost, slow deployment, and with deep inflexibility. Software-based networking allows firewalls to be instantly deployed as virtual machines with no operating cost.

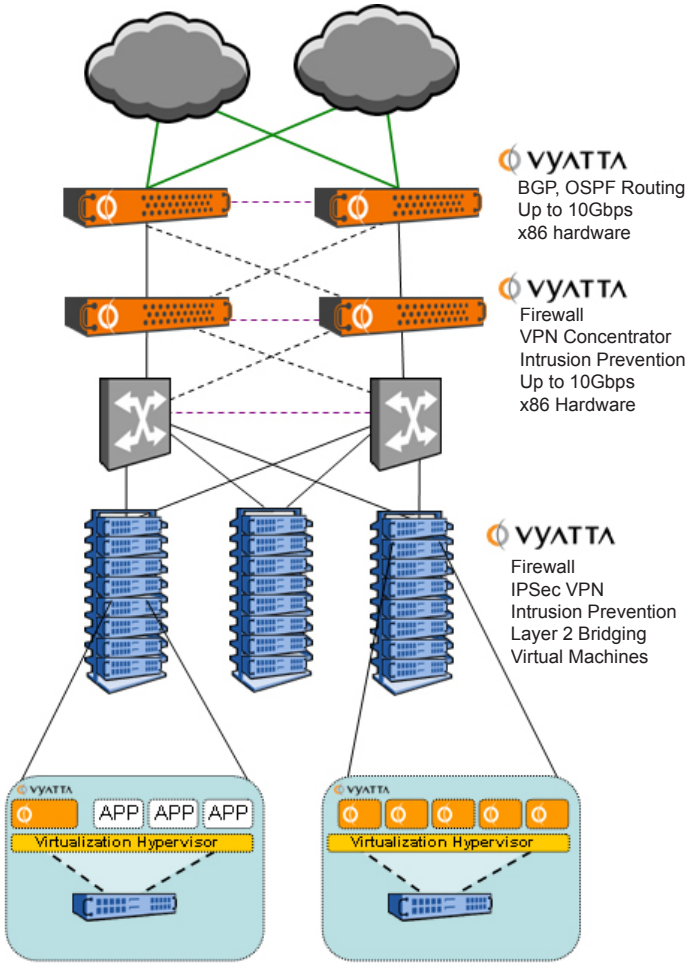

The Vyatta Network OS

The Vyatta network operating system is a scalable, integrated, enterprise-class networking solution that delivers advanced routing and network security functionality for physical, virtual and cloud networking environments. The Vyatta network OS includes dynamic routing, stateful firewall, VPN support, threat protection, traffic management and more in a package that is optimized to take advantage of multicore x86 processing power, common hypervisor platforms and emerging cloud architectures. All features are configured through Vyatta's familiar, networking-centric CLI, web-based GUI or third party management systems using the Vyatta Remote Access API.

Vyatta Software Highlights:

Network Connectivity

At the core of the Vyatta system is a complex routing engine with full support of IPv4 and IPv6 dynamic routing protocols (BGP, OSPF, RIP). Vyatta systems include support for 802.11 wireless, Serial WAN Interfaces and a wide variety of 10/100 thru 10Gb Ethernet NICs.

Firewall Protection

The Vyatta firewall features IPv4/IPv6 stateful packet inspection to intercept and inspect network activity and protect your critical data. Vyatta advanced firewall capabilities include stateful failover, zone and time-based firewalling, P2P filtering and more.

Content and Threat Protection

Vyatta systems offer an additional level of proactive threat protection with integrated secure web filtering and advanced intrusion prevention rules available as subscription-based Vyatta PLUS services.

Secure Connectivity

Establish secure site-to-site VPN tunnels with standards-based IPSec VPN between two or more Vyatta systems or any IPSec VPN device. Or provide secure network access to remote users via Vyatta's SSL-based OpenVPN functionality.

Traffic Management

The Vyatta system provides a wide variety of QoS queuing mechanisms that can be applied to inbound traffic and outbound traffic for identifying and prioritizing applications and traffic flows.

High Availability

Mission critical networks can deploy Vyatta with the confidence that high availability and system redundancy can be achieved through a number of industry standard failover and configuration synchronization mechanisms.

IPv6 Compatibility

Vyatta Subscription Edition software is the only software-based routing and security solution with proven IPv6 functionality and interoperability, ensuring a future-proof investment in a solution that offers a simplified migration path from IPv4 to IPv6.

Administration & Authentication

Vyatta systems can be managed through our familiar network-centric command line interface, web-based GUI or through external management systems using Vyatta's Remote Access API. All network management sessions can be securely managed using SSHv2, RADIUS or TACACS+.

Monitoring and Reporting

Vyatta systems present complete logging and diagnostics information that can be monitored using in industry standard toolsets such as SNMP, Netflow, Syslog, Wireshark and more.

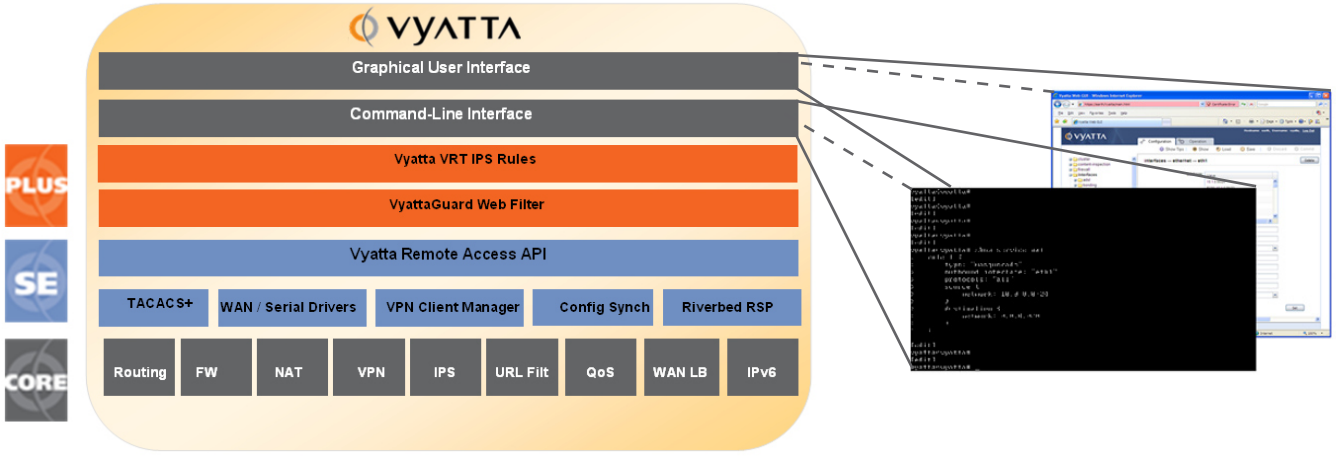

About Vyatta

Vyatta is disrupting the networking industry by delivering a software-based, open-source, network operating system that is portable to standard x86 hardware as well as common virtualization and cloud computing platforms. Vyatta software provides a complete enterprise-class routing and security feature set capable of scaling from DSL to 20Gbps performance at a fraction of the cost of proprietary solutions. Thousands of physical and virtual infrastructures around the world, from small enterprise to Fortune 500 customers, are connected and protected by Vyatta. For more information, please visit http://www.vyatta.com.