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 3: The Wide Area Network (WAN)

Featured Sponsors of Part 3:

A10 Networks
BlueCoat
certeon Accelerate Your Business
CITRIX
exinda
ipanema Technologies
radware
riverbed
Silver Peak
TALARI NETWORKS
VYATTA Open Networking

Additional Sponsors:

Alcatel·Lucent Enterprise
AVAYA The Power of We™
ca
CISCO
DELL Force10
extreme networks Make Your Network Mobile
hp
IBM
Infoblox
LSI
NETSCOUT
Packet Design
VISUAL NETWORK SYSTEMS

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 third of the serial publications. One goal of this publication is to describe the wide area networking (WAN) environment and to contrast it with the data center LAN environment. Another goal is to identify the factors that are causing IT organizations to rethink their approach to wide area networking and to provide insight into the WAN technologies and design choices that IT organizations are making. The third and primary goal of this publication is to describe the WAN architecture and technology options that either are currently available in the market or are likely to be available within two years.

In order to quantify the technology and design choices that IT organizations are making, this publication 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 publication, those two groups of respondents will be respectively referred to as The Webtorials Respondents and The Interop Respondents.

The Wide Area Network (WAN)

Introduction

Background

The modern WAN got its start in 1969 with the deployment of ARPANET which was the precursor to today's Internet. The technology used to build the Internet began to be commercialized in the early 1970s with the development of X.25 based packet switched networks.

In addition to the continued evolution of the Internet, the twenty-year period that began in 1985 saw the deployment of four distinct generations of enterprise WAN technologies[1]. For example, in the mid to late 1980s, it became common for enterprise IT organizations to deploy integrated TDM-based WANs to carry both voice and data traffic. In the early 1990s, IT organizations began to deploy Frame Relay-based WANs. In the mid to late 1990s, some IT organizations replaced their Frame Relay-based WANs with WANs based on ATM (Asynchronous Transfer Mode) technology. In the 2000s, many IT organizations replaced their Frame Relay or ATM-based WANs with WANs based on MPLS. Cost savings was the primary factor that drove the adoption of each of the four generations of WAN technologies. The cost savings, however, were very modest when compared to the price performance improvements that are associated with local area networking.

However, in contrast to the volatility of this twenty-five year period:

Today there is not a fundamentally new generation of WAN technology in development.

Relative to the deployment of new WAN services, what sometimes happens in the current environment is that variations are made to existing WAN technologies and services. An example of that phenomenon is Virtual Private LAN Service (VPLS)[2]. As described later in this section of the report, within VPLS an Ethernet frame is encapsulated inside of MPLS. While creating variations on existing services can result in significant benefits, it does not produce fundamentally new WAN services.

Contrasting the LAN and the WAN

As noted, the WAN is notably different than the data center LAN. These differences include the fact that:

- After a lengthy period in which there was little or no fundamental innovation, the LAN is experiencing broad fundamental change. In contrast, after a lengthy

[1] An enterprise WAN is designed to provide for connectivity primarily within the enterprise and between the enterprise and key contacts such as partners. This is in contrast to the Internet that is designed to provide universal connectivity.
[2] http://vlt.me/vpls-0810

period in which the WAN underwent repeated fundamental change, there are currently no fundamental changes in store for the WAN.

- In the vast majority of instances, the latency, jitter and packet loss that the LAN exhibits doesn't have an appreciable impact on application performance. In many instances, the latency, jitter and packet loss that the WAN exhibits has an appreciable impact on application performance. This is particularly true of 3G/4G networks.

- One of the primary design criteria for designing a data center LAN is scalability. A manifestation of the ongoing improvements in LAN scalability is that over the last fifteen years the speed of a data center LAN has increased from 10 Mbps to 10 Gbps – which is a factor of a thousand. In contrast, in many cases the primary design criterion for designing a WAN is to minimize cost. For example, in many parts of the world it is possible to get high-speed WAN links such as an OC-192 link. These links, however, are usually not affordable.

- The LAN follows Moore's Law. In contrast, the price/performance of WAN services such as MPLS tends to improve by only a couple of percentage points per year.

The WAN doesn't follow Moore's Law.

WAN Budgets

Both in 2010 and again in 2011, The Webtorials Respondents were asked how their budget for the forthcoming year for all WAN services compares to what it is in the current year year. Their responses are contained in Table 1.

Table 1: WAN Budget Increases

	Responses in 2010	Responses in 2011
Reduced by more than 10%	5.9%	3.2%
Reduced by 1% to 10%	17.0%	11.1%
Basically static	36.3%	34.9%
Increased by 1% to 10%	30.4%	32.8%
Increased by more than 10%	10.4%	18.0%

The change in the budget for WAN services in 2011 bears a lot of similarities to what the change in the budget for WAN services was in 2010. The biggest differences are at the extremes. For example, in 2011 the percentage of The Webtorials Respondents that expect that their WAN budgets will decrease has been cut almost in half when compared to what it was in 2010. In addition, in 2011 the percentage of The Webtorials Respondents that expect that their WAN budgets will increase by more than 10% is almost double what it was in 2010.

Over the next year, roughly forty percent of IT organizations will increase their WAN budget and in many cases, the increase will be significant.

As is explained in the next subsection, the adoption of cloud computing will increase the rate of growth in the amount of traffic that transits the WAN. As such,

IT organizations must either make changes to how they use WAN services, or else accept ongoing increases in their WAN budget due to the increased traffic generated by the use of cloud computing.

Drivers of Change

As explained in the section of this report entitled *The Emergence of Cloud Computing and Cloud Networking,* one of the characteristics of cloud computing is increased reliance on the network. The increased reliance on the WAN in particular stems from the fact that the resources that support cloud computing solutions are centralized in a small number of data centers and the vast majority of users access these solutions over the WAN. Hence, the more use that organizations make of cloud computing, the more traffic transits the WAN.

Below are some of the specific factors that are putting more traffic onto the WAN and hence, driving the need for IT organizations to change their approach to wide area networking.

Virtual Machine Migration

The section of this report entitled *The Emerging Data Center LAN* quantified the great interest that IT organizations have in server virtualization in general and in moving virtual machines (VMs) between data centers in particular. That section of the report also discussed the fact that one of the requirements associated with moving VMs between data centers is that the data storage location, including the boot device used by the VM being migrated, must be accessible by both the source and destination physical servers at all times. If the servers are at two distinct locations and the data is replicated at the second site, the two data sets must be identical. One approach to enabling data access is to extend the SAN to the two sites and to maintain a single data source. Another option is migrate the data along with the VM to the secondary site. In either case, it is necessary to coordinate VM and storage migrations and to be able to move large data sets efficiently between data centers, which will have a significant impact on the WAN.

Virtual Desktops

Another form of virtualization that will drive a further increase in WAN traffic is desktop virtualization. In order to quantify the interest that IT organizations have in desktop virtualization, The Webtorials Respondents were asked to indicate the percentage of their company's desktops that have either already been virtualized or that they expected would be virtualized within the next year. Their responses are shown in Table 2.

Table 2: Deployment of Virtualized Desktops					
	None	**1% to 25%**	**26% to 50%**	**51% to 75%**	**76% to 100%**
Have already been virtualized	55%	36%	3%	1%	4%
Expect to be virtualized within a year	30%	51%	8%	4%	7%

The data in Table 2 indicates the growing interest that IT organizations have in desktop virtualization. For example,

> ***Over the next year, the percentage of IT organizations that have not implemented any desktop virtualization will be cut roughly in half.***

Part of the challenge in supporting virtualized desktops is that the implementation of virtualized desktops puts more traffic on the WAN, which typically leads to the need for more bandwidth. In addition to the bandwidth challenges, as explained in The 2011 Application and Service Delivery Handbook[3], there are performance challenges associated with each of the two primary form of desktop virtualization; e.g., client side (a.k.a., streamed desktops) and server side (a.k.a., hosted desktops).

In the case of client side desktop virtualization, the code for streamed applications is typically transferred via a distributed file system protocol, such as CIFS, which is well known to be a chatty protocol[4]. Server side protocols such as ICA and RDP, tend to work relatively well when supporting traditional applications[5]. However, these protocols can behave badly when supporting graphics and video. Some newer protocols, such as Teradici's PC-over-IP (PCoIP)[6], consume considerable WAN bandwidth and are latency sensitive.

Collaboration

As was described in the section of this report that is entitled *The Emergence of Cloud Computing and Cloud Networking,* many organizations are beginning to acquire services such as collaboration from a cloud computing service provider (CCSP). Independent of whether the collaboration service is provided by a CCSP or by the IT organization, it stresses the WAN. This stress comes in part from the fact that the performance of applications such as video and telepresence is very sensitive to delay, jitter and packet loss. The stress also comes in part because video and telepresence consume considerable WAN bandwidth. It is common, for example, to allocate several megabits per second of WAN bandwidth to a single telepresence session.

[3] http://www.webtorials.com/content/2011/07/2011-application-service-delivery-handbook.html
[4] A chatty protocol requires hundreds, if not thousands of round trips to complete a transaction.
[5] Even though they work relatively well in native mode, many IT organizations choose to implement WAN optimization in order to improve the performance of these protocols.
[6] http://en.wikipedia.org/wiki/PCoIP

The current conventional wisdom in the IT industry is that organizations are increasing their use of video. In order to evaluate that assertion, The Webtorials Respondents were asked to indicate how much change in the use of all forms of video they anticipated that their organization would make over the next year. Their responses are shown in Table 3.

Table 3: Anticipated Change in the Use of Video

Down by More than 25%	Down by 1% to 25%	No Change	Up 1% to 25%	Up 26% to 50%	Up 51% to 75%	Up 76% to 100%	Up more than 100%
0.0%	0.5%	19.9%	34.6%	24.6%	9.9%	4.7%	5.8%

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

> ***Over the next year almost 80% of IT organizations will increase their use of video, and in many cases the increased use of video will be substantial.***

Mobile Workers

In the last few years there has been an explosive growth in the number of mobile workers. There are a number of key concerns relative to supporting mobile workers. One such concern is the number and types of devices that mobile workers use. As recently as a couple of years ago, many IT organizations tried to control the types of devices that their users could utilize. In the current environment the majority of IT organizations are in a position where they have to support a large and growing set of mobile devices from a range of vendors. It most cases mobile workers have two mobile devices (i.e., a laptop and a smartphone) and in a growing number of cases, mobile workers have three mobile devices; i.e., a laptop, a smartphone and a tablet.

Another key concern relative to supporting mobile workers is how the applications that these workers access has changed. At one time, mobile workers tended to primarily access either recreational applications or applications that are not delay sensitive; e.g., email. However, in the current environment mobile workers also need to access a wide range of business critical applications, many of which are delay sensitive. This shift in the applications accessed by mobile workers was highlighted by SAP's recent announcement[7] that it will leverage its Sybase acquisition to offer access to its business applications to mobile workers.

One of the technical issues associated with supporting mobile workers' access to delay sensitive, business critical applications is that because of the way that TCP functions, even the small amount of packet loss that is often associated with wireless networks results in a dramatic reduction in throughput. A related issue is that typically there is a large amount of delay associated with 3G and 4G networks.

[7] Wall Street Journal, May 17, 2011, page B7

Traditional WAN Services

Background

The Webtorials Respondents were given a set of eleven WAN services and asked to indicate the extent to which they currently utilize each WAN service. The survey question included Frame Relay and ATM among the set of WAN services. In the not too distant past, these services were widely deployed. However, over half of The Webtorials Respondents don't have any Frame Relay in their networks and almost two thirds of The Webtorials Respondents don't have any ATM in their networks. In addition, few IT organizations are increasing their use of these technologies[8], while many IT organizations are decreasing their use of these technologies[9].

One of the observations that can be drawn from the response to this survey question is that:

The primary WAN services used by IT organizations are MPLS and the Internet.

The Webtorials Respondents were also asked to indicate the change that they anticipated that their organization would make over the next year relative to their usage of MPLS, VPLS and the Internet. Table 4 shows the percentage of The Webtorials Respondents that indicated that their organization would increase their use of those services.

Table 4: Increase in the use of key WAN Services

WAN Service	Percentage of Organizations Increasing their Use of this Service
MPLS	50.0%
VPLS	37.5%
Internet traffic to external sites	83.5%
Internet traffic to internal sites	74.3%

One of the conclusions that can be drawn from the data in Table 4 is that:

While IT organizations will increase their reliance on both MPLS and the Internet, they will make a relatively greater increase in their reliance on the Internet.

[8] Roughly 2% of IT organizations are increasing their use of Frame Relay and 6% of IT organizations are increasing their use of ATM.
[9] Roughly 34% of IT organizations are decreasing their use of Frame Relay and 22% of IT organizations are decreasing their use of ATM.

WAN Design Criteria and Challenges

The Webtorials Respondents were given a list of possible concerns and were asked to indicate which two were their company's primary concerns relative to its use of MPLS and the Internet. The set of concerns that were presented to The Webtorials Respondents is shown in the left hand column of Table 5. The second and third columns from the left in Table 5 show the percentage of The Webtorials Respondents who indicated that the concern is one of their company's two primary concerns with MPLS and the Internet respectively. The right hand column is the difference between the second and third columns from the left. This column will be referred to as the delta column.

The delta column contains positive and negative numbers. A positive number means that that concern was mentioned more often relative to MPLS than it was mentioned relative to the Internet. For example, The Webtorials Respondents mentioned cost as one of their primary concerns about the use of MPLS 22.1% more often than they mentioned cost as one of their primary concerns about the use of the Internet. Analogously, a negative number means that that concern was mentioned more often relative to the Internet than it was relative to MPLS. For example, The Webtorials Respondents mentioned latency as one of their primary concerns about the use of the Internet 19.3% more often than they mentioned latency as one of their primary concerns about use of MPLS.

Table 5: Concerns about MPLS

Concern	MPLS	Internet	Delta
Cost	60.1%	38.0%	22.1%
Lead time to implement new circuits	32.2%	11.4%	20.8%
Uptime	30.1%	46.3%	-16.2%
Latency	27.0%	46.3%	-19.3%
Lead time to increase capacity on existing circuits	23.5%	13.1%	10.4%
Jitter	14.8%	18.8%	-4.0%
Packet Loss	12.2%	26.2%	-14.0%

The primary concerns that IT organizations have with the use of MPLS are cost, the lead time to implement new circuits and uptime. The primary concerns that IT organizations have with the use of the Internet are uptime, latency and cost.

IT organizations typically design their WAN based on the following criteria:

1. Minimize cost
2. Maximize availability
3. Ensure appropriate performance

As shown in Table 5, MPLS is regarded by The Webtorials Respondents as doing a good job at ensuring appropriate performance because it exhibits relatively small amounts of delay, jitter and packet loss. Unfortunately, MPLS is regarded poorly relative to the goal of minimizing cost. In contrast, the Internet is regarded relatively well on the goal of minimizing cost but is regarded relatively poorly on the goal of ensuring appropriate performance. In addition, The Webtorials Respondents expressed concerns about both MPLS and the Internet relative to the goal of maximizing availability.

As was pointed out in the section of this report entitled *The Emergence of Cloud Computing and Cloud Networking*, 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*. As that section also pointed out, in order to support a small number of business critical services and applications, a cloud network that is *good enough* will have to provide the highest possible levels of availability and performance. However, in a growing number of instances, a cloud network is *good enough* if it provides a best effort level of service at a reduced price. Hence, independent of the concerns that IT organizations have about the Internet:

> ***In a growing number of instances, Internet-based VPNs that use DSL for access are 'good enough' to be a cloud network.***

Some of the concerns that IT organizations have with the use of the Internet such as uptime, stem from the fact that in many cases IT organizations access the Internet over a single DSL link. The availability of DSL is somewhat lower than the availability of access technologies such as T1/E1 links. One impact of this reduced availability is that Internet VPNs based on DSL access are often used only as a backup connection to a primary private WAN circuit. This is unfortunate because the shortfall in quality is fairly small when compared to the dramatic cost savings and additional bandwidth that can be realized by using broadband connections such as DSL and cable. One technology that addresses this issue is referred to as an *aggregated virtual WAN*.

> ***The key concept behind an aggregated virtual WAN is that it simultaneously utilizes multiple enterprise WAN services and/or Internet connections in order to optimize reliability and minimize packet loss, latency and jitter.***

Aggregated virtual WANs and other types of alternate WAN services are discussed later in this section of the report. As that discussion highlights, aggregated virtual WANs have the potential to maximize the benefits of the Internet and possibly MPLS while minimizing the negative aspects of both.

Local Access to the Internet

The traditional approach to providing Internet access to branch office employees has been to carry their Internet traffic on the organization's enterprise network (e.g., their MPLS network) to a central site where the traffic was handed off to the Internet. The advantage of this approach is that it enables IT organizations to exert more control over their Internet traffic and it simplifies management in part because it centralizes the complexity of implementing and managing security policy. One disadvantage of this approach is that it results in extra traffic transiting the enterprise's WAN, which adds to the cost of the WAN. Another disadvantage of this approach is that it usually adds additional delay to the Internet traffic. The fact that centralized Internet access exhibits these disadvantages is significant because as highlighted in Table 5, cost and delay are two of the primary concerns that IT organizations have relative to the use of the Internet.

Some of the concerns that IT organizations have about the use of the Internet are exacerbated by backhauling Internet traffic to a central site.

The Webtorials Respondents were asked to indicate how they currently route their Internet traffic and how that is likely to change over the next year. Their responses are contained in Table 6.

Table 6: Routing of Internet Traffic

Percentage of Internet Traffic	Currently Routed to a Central Site	Will be Routed to a Central Site within a Year
100%	39.7%	30.6%
76% to 99%	24.1%	25.4%
51% to 75%	8.5%	13.4%
26% to 50%	14.2%	14.2%
1% to 25%	7.1%	6.7%
0%	6.4%	9.7%

Driven in part to save money and in part to improve application performance:

Over the next year, IT organizations will make an increased use of distributed access to the Internet from their branch offices.

Cloud Networking Without the Internet

There is a temptation to associate the WAN component of *cloud networking* either exclusively or primarily with the traditional Internet[10]. However, due to a variety of well-known issues, such as packet loss at peering points, BGP's inability to choose the path with the lowest delay, the TCP Slow start algorithm, the Internet often exhibits

[10] Throughout this report, the phrase "traditional Internet" will refer to the use of the Internet with one access link and not optimization functionality.

performance problems. As such, the Internet is not always the most appropriate WAN service to use to access cloud computing solutions. To put the use of the Internet into context, The Webtorials Respondents were asked to indicate which WAN service their users would most likely use when accessing public and private cloud computing services over the next year. Their responses are shown in Table 7.

Table 7: WAN Services to Access Cloud Computing Services

	The Internet	An Internet overlay from a company such as Akamai	A traditional WAN service such as MPLS	WAN Optimization combined with a traditional WAN service; e.g. MPLS
Public Cloud Computing Services	58.8%	10.3%	17.7%	13.2%
Private Cloud Computing Services	27.2%	4.1%	37.8%	30.9%

The Webtorials survey base was asked the same question in 2010 and the answers they provided in 2010 are remarkably similar to what they provided in 2011. This implies that few IT organizations are making a significant change to their WAN in order to support cloud computing.

The data in Table 7 indicates that IT organizations understand the limitations of the traditional Internet relative to supporting cloud computing. In particular:

> ***In somewhat less than half of the instances that business users are accessing public cloud computing services, the primary WAN service that they intend to use is not the traditional Internet.***
>
> ***In almost three quarters of the instances that business users are accessing private cloud computing services, the primary WAN service that they intend to use is not the traditional Internet.***

However, techniques that IT organizations can use to mitigate their concerns about the use of the Internet are discussed later in this section of the report.

Service Level Agreements

As previously stated, the majority of IT organizations utilize MPLS and the usage of MPLS is expected to increase significantly. One of the reasons for the popularity of MPLS is that the major suppliers of MPLS services offer a number of different classes of service (CoS) designed to meet the QoS requirements of the varying types of applications that transit a WAN. For example, real-time applications are typically placed in what is often referred to as a Differentiated Services Code Point (DSCP) Expedited Forwarding class that offers minimal latency, jitter, and packet loss. Mission

critical business applications are typically relegated to what is often referred to as a DSCP Assured Forwarding Class.

Each class of MPLS service is typically associated with a service level agreement (SLA) that specifies contracted ranges of availability, latency, packet loss and possibly jitter. Unfortunately, in many cases the SLAs are weak. In particular, it is customary to have the SLAs be reactive in focus; i.e., the computation of an outage begins when the customer opens a trouble ticket. In most cases, the carrier's SLA metrics are calculated as network-wide averages rather than for a specific customer site. As a result, it is possible for a company's data center to receive notably poor service in spite of the fact that the network-wide SLA metrics remain within agreed bounds. In addition, the typical level of compensation for violation of service level agreements is quite modest.

To gauge the effectiveness of SLAs that IT organizations receive from their network service providers (NSPs), The Webtorials Respondents were asked to indicate which of the following best describes the SLAs that they get from their NSPs for services such as MPLS.

- The SLAs go a long way towards ensuring that we get a quality service from the network service provider.
- The SLAs are better than nothing, but not by much.
- The SLAs are not worth the paper they are written on.

Their responses are shown in Figure 1.

Figure 1: The Effectiveness of SLAs

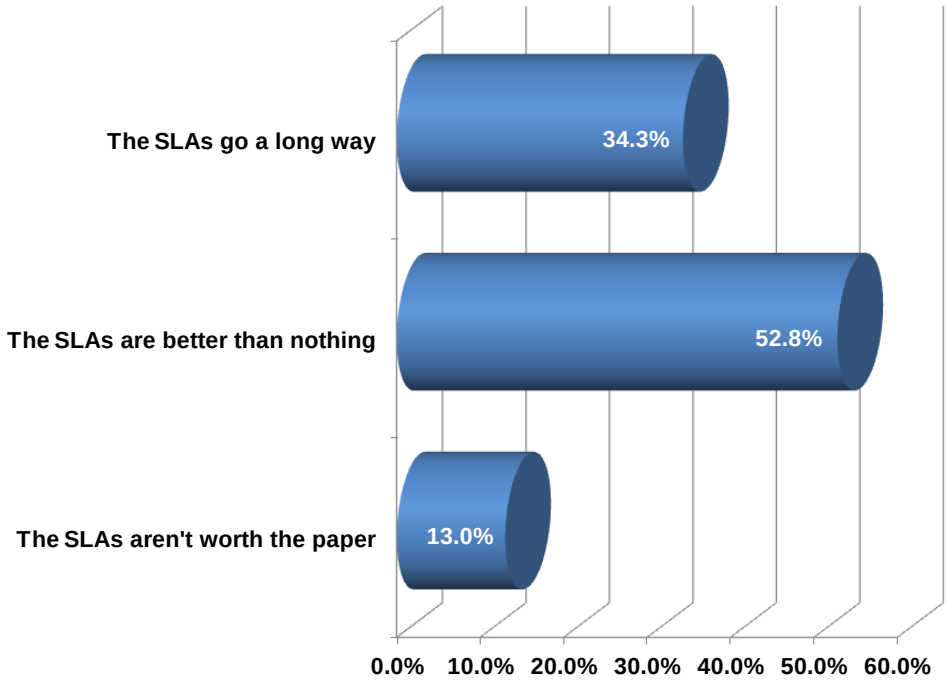

The fact that two thirds of The Webtorials Respondents indicated that the SLAs that they receive from network service providers are either not worth the paper they are written on, or that the SLAs they receive are not much better than nothing, demonstrates the weak nature of most SLAs.

The majority of IT organizations don't regard the SLAs that they receive from their network service providers as being effective.

Optimizing the Performance of IT Resources

Background

This subsection of the report will discuss techniques that IT organizations can implement to overcome the limitations of protocols and applications and to optimize the use of their servers. The focus of this subsection is on how these techniques enable IT organizations to ensure acceptable application and service delivery over a WAN. The discussion in this subsection will focus on two classes of products: WAN Optimization Controllers (WOCs) and Application Delivery Controllers (ADCs).

The introduction to this section of this report discussed how the adoption of cloud computing in general is impacting the WAN and also discussed some of the specific factors that are driving change in the WAN. These factors included both the increasing number of mobile workers and the impact of multiple forms of virtualization. In order to gauge the effect that these factors have on the ability of an IT organizations to ensure acceptable application and service delivery, The Webtorials Respondents were asked "How will the ongoing adoption of mobile workers, virtualization and cloud computing impact the difficulty that your organization has with ensuring acceptable application performance?" Their responses are shown in Figure 2.

Figure 2: Impact of Change on Service and Application Delivery

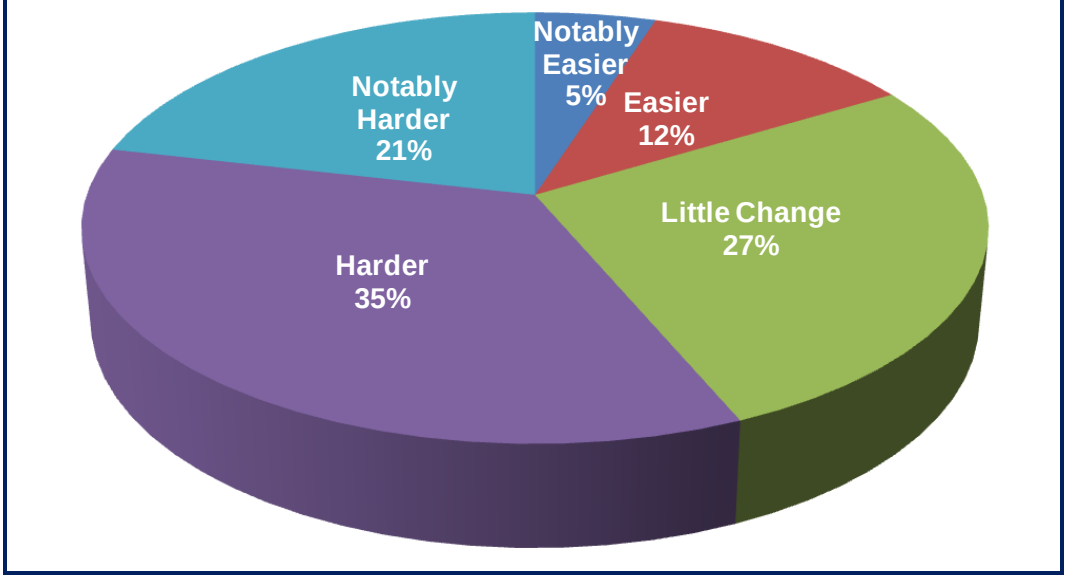

One conclusion that can be drawn from Figure 2 is that:

> ***The majority of IT organizations believe that factors such as the growth in the number of mobile workers and the increase in the use of virtualization and cloud computing will make ensuring acceptable service and application delivery either harder or notably harder.***

WAN Optimization Controllers (WOCs)

Goals of a WOC

The goal of a WOC is to improve the performance of applications and services that are delivered across a WAN from the data center either to a branch office, a home office or directly to a mobile user. In some cases the data center is owned and managed by the enterprise IT organization and in other cases it is owned and managed by a cloud computing service provider (CCSP). The WOC accomplishes this goal by implementing techniques to overcome the limitations of the WAN such as constrained bandwidth, delay and packet loss.

WOCs are often referred to as *symmetric solutions* because they typically require complementary functionality at both ends of the connection; i.e., a WOC in the data center and another WOC at the branch office. However, the requirement to improve the performance of applications and services acquired from a CCSP has been the impetuous for the deployment WOCs in an asymmetric fashion. As shown in Figure 3, in an asymmetric deployment of a WOC content is downloaded from a CCSP to a WOC in a branch office. Once the content is stored in the WOC's cache for a single user, subsequent users who want to access the same content will experience accelerated application delivery. Caching can be optimized for a range of cloud content, including Web applications, streaming video (e.g., delivered via Flash/RTMP or RTSP) and dynamic Web 2.0 content.

As previously described, IT organizations are moving away from a WAN design in which they backhaul their Internet traffic from their branch offices to a central site prior to handing it off to the Internet. Also, as is described in the next section of this report, there are a variety of techniques that enable IT organizations to improve both the price-performance and the availability of distributed Internet access. As a result of these factors, asymmetric WOC deployment as described in the preceding paragraph will increasingly be supported by a network design that features distributed Internet access. However, for this network design to be effective, IT organizations need to ensure that the design includes appropriate security functionality.

Figure 3: Asymmetric WOC Deployment

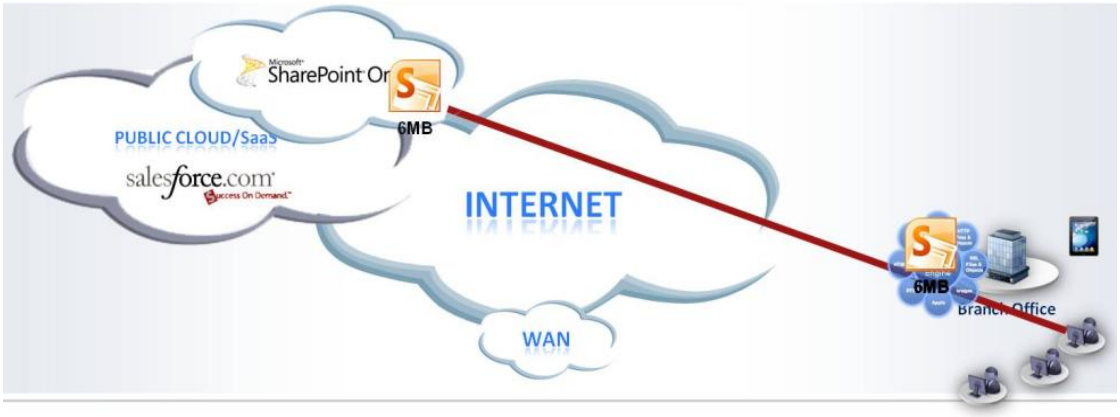

Modeling Application Response Time

A model is helpful to illustrate how the performance of a WAN can impact the performance of an application and it also serves to illustrate how a WOC can improve application performance. The following model (Figure 4) is a variation of the application response time model created by Sevcik and Wetzel[11]. Like all mathematical models, the following is only an approximation. For example, the model shown in Figure 4 doesn't account for the impact of packet loss.

As shown below, the application response time (R) is impacted by amount of data being transmitted (Payload), the WAN bandwidth, the network round trip time (RTT), the number of application turns (AppTurns), the number of simultaneous TCP sessions (concurrent requests), the server side delay (Cs) and the client side delay (Cc).

Figure 4: Application Response Time Model

$$R \approx \frac{Payload}{Goodput} + \frac{(\#\ of\ AppsTurns * RTT)}{Concurrent\ Requests} + Cs + Cc$$

In order to improve the performance of applications that are delivered over the WAN, WOCs implement a variety of techniques. For example, to mitigate the impact of a large payload, WOCs implement techniques such as compression and de-duplication. These techniques are explained in detail in The 2011 Application Delivery Handbook. The handbook also details criteria that IT organizations can use to evaluate WOCs as well as specific techniques that WOCs need to support in order to optimize:

- The rapidly growing amount of traffic that goes between data centers
- Desktop virtualization
- Delay sensitive applications such as voice, video and telepresence

[11] Why SAP Performance Needs Help, NetForecast Report 5084, http://www.netforecast.com/ReportsFrameset.htm

The 2011 Application Delivery Handbook also describes techniques that can optimize the delivery of applications to mobile workers. Many IT organizations, however, resist putting any additional software on the user's device. In addition, many users resent having multiple clients (e.g., WOC, SSL VPN, IPSec VPN, wireless/cellular access) on their access device that are not integrated. One option for IT organizations on a going forward basis is to implement WOC software on mobile devices that is integrated with the other clients used by mobile workers. As is explained below, an alternative way that IT organizations can improve the performance of applications and services delivered to mobile users is to utilize an optimization service from a CCSP.

Application Delivery Controllers (ADCs)

ADCs provide some functionality, such as compression, that optimizes the delivery of bulk data over the Internet. However, the primary goal of an ADC is to improve the performance of servers.

The current generation of ADCs evolved from the earlier generations of Server Load Balancers (SLBs) that were deployed in front of server farms. While an ADC still functions as a SLB, the ADC has assumed, and will most likely continue to assume, a wider range of sophisticated roles that enhance server efficiency and provide asymmetrical functionality to accelerate the delivery of applications from the data center to individual remote users.

An ADC provides more sophisticated functionality than a SLB does.

Referring back to Figure 4, one of the factors that increase the application response time is server side delay. An ADC can reduce server side delay and hence can reduce the application response time. In particular, the ADC can allow a number of compute-intensive functions, such as SSL processing and TCP session processing, to be offloaded from the server. Server offload can increase the transaction capacity of each server, reducing the number of servers required for a given level of business activity.

The 2011 Application Delivery Handbook describes the primary techniques implemented by ADCs and identifies criteria that IT organizations can use to evaluate ADCs

Virtual Appliances

The section of this report entitled *The Emerging Data Center LAN* used the phrase *virtual switch* in two fundamentally different ways. One way referred to making two or more physical switches appear to be a single logical switch. The other way referred to the switching functionality that resides inside of a virtualized server.

In similar fashion, it is possible to look at a *virtual appliance* in a variety of fundamentally different ways. For example, two or more appliances, such as ADCs, can be combined to appear as a single logical ADC. Alternatively, a single physical ADC can be partitioned into a number of logical ADCs or ADC contexts. Each logical ADC can be configured individually to meet the server-load balancing, acceleration and security requirements of a single application or a cluster of applications.

However, the most common use of the phrase *Virtual Appliance* refers to what is typically appliance-based software, together with its operating system, running in a VM. Virtual appliances can include WOCs, ADCs, firewalls, routers, IDS, IPS and performance monitoring solutions. As explained in the next subsection of this report, virtual appliances make it easier for an IT organization to deploy network and application optimization functionality at a CCSP's data center. That, however, is not the only advantage of a virtualized appliance.

> ***One of the compelling advantages of a virtualized appliance is that the acquisition cost of a software-based appliance can be notably less than the cost of a hardware-based appliance with same functionality.***

In many cases the acquisition cost of a software-based appliance can be a third less than the cost of a hardware-based appliance[12]. In addition, a software-based solution can potentially leverage the functionality provided by the hypervisor management system to provide a highly available system without having to pay for a second appliance[13].

In addition to cost savings, another advantage of a virtual appliance is that it offers the potential to alleviate some of the management burdens because most of the provisioning, software updates, configuration, and other management tasks can be automated and centralized at the data center. An example of this is that if virtualized appliances have been deployed, then it is notably easier than it is in a more traditional environment for various networking functions (WOC, ADC, firewall, router, etc.) to be migrated along with VMs in order to replicate the VMs's networking environment in its new location.

> ***In many instances the benefits of the dynamic movement of a VM from one server to another are maximized if the supporting infrastructure is virtualized and can also be dynamically moved.***

A virtualized ADC also makes it easy for an IT organization to package and deploy a complete application. One example of this packaging is the situation in which an entire application resides on VMs inside a physical server. The virtualized ADC that supports the application resides in the same physical server and it has been tuned for the particular application. This makes it easy to replicate or migrate that application as needed. In this case, a virtualized ADC also provides some organizational flexibility. For example, the virtual ADC might be under the control of a central IT group or it might be under the control of the group that supports that particular application. The later is a viable option from an organizational perspective because any actions taken by the application group relative to their virtual ADC will only impact their application.

A virtual firewall appliance can also help IT organizations meet some of the challenges associated with server virtualization. That follows because virtual firewall appliances can be leveraged to provide isolation between VMs on separate physical servers as

[12] The actual price difference between a hardware-based appliance and a software-based appliance will differ by vendor.

[13] This statement makes a number of assumptions, including the assumption that the vendor does not charge for the backup software-based appliance.

well as between VMs running on the same physical server. Through tight integration with the virtual server management system, virtual firewall appliances can also be dynamically migrated in conjunction with VM migration where this is necessary to extend a trust zone to a new physical location. In addition, hypervisor APIs, such as VMware's Vsafe, can allow physical/virtual firewall consoles to monitor servers for abnormal CPU, memory, or disk activity without the installation of special agent software.

The research report entitled Virtualization: Benefits, Challenges and Solutions, contains more detail on virtual appliances. Included in that report is a discussion of the challenges associated with virtual appliances, as well as suggested evaluation criteria.

Optimizing Access to Public Cloud Computing Solutions

As noted in the section of this report entitled *The Emergence of Cloud Computing and Cloud Networking*, one of the key challenges facing IT organizations that use either SaaS or IaaS solutions is improving the performance of those solutions. In order to quantify what IT organizations are doing to improve the performance of those solutions, The Webtorials Respondents were asked "If your company either currently acquires services from an Infrastructure-as-a-Service (IaaS) provider or you expect that they will within the next year, which of the following best describes the primary approach that your company will take to optimizing the performance of those solutions?" They were asked a similar question about their use of SaaS solutions.

The responses to these two questions are summarized in Table 8. The leftmost column in Table 8 lists the approaches that The Survey Respondents had to choose from. The middle column shows the percentage of The Survey Respondents that indicated that that would be their primary approach for optimizing services acquired from an IaaS provider. The rightmost column shows the percentage of The Survey Respondents that indicated that that would be their primary approach for optimizing services acquired from a SaaS provider.

Table 8: Optimizing CCSP Services

Approach	IaaS Services	SaaS Services
Leverage optimization functionality provided by the CCSP	29.3%	31.9%
Place a WOC on the service provider's site and on our premise	22.4%	13.9%
Use an optimization service from a company such as Akamai or Virtela	4.1%	3.6%
Do nothing	19.0%	17.5%
Don't know	25.2%	33.1%

A number of conclusions can be drawn from the data in Table 8, including:

There is significant interest in placing a WOC on premise at an IaaS provider's data centers.

Between a quarter and a third of IT organizations don't know how they will optimize the performance of services that they acquire from an IaaS or a SaaS provider.

In addition, referencing back to the discussion in the previous subsection, IT organization will have a notably easier time placing an optimization device, whether that is a WOC or an ADC, at an IaaS provider's data center if the device is virtualized. That follows because if the device is virtualized, the IT organization can control the deployment of the functionality. If the device is physical, then the IT organization needs to get the IaaS provider to offer space for the device and to install it.

Alternative WAN Services

As noted, there is not a new generation of fundamentally new WAN technology currently under development. However, as is described below, there are a number of WAN service alternatives that are variations on existing WAN technologies and services that better enable IT organizations to meet their WAN design goals. A number of these alternatives are either complementary to the WAN optimization technologies previously discussed or they depend partially on WAN optimization technologies to deliver acceptable levels of service quality.

An Internet Overlay

As described in the preceding subsection, IT organizations often implement WOCs and ADCs in order to improve network and application performance. However, these solutions make the assumption that performance characteristics within the WAN itself can't be optimized because they are determined by the relatively static service parameters controlled by the WAN service provider. This assumption is reasonable in the case of WAN services such as MPLS. However, this assumption doesn't apply to enterprise application traffic that transits the Internet because there are significant opportunities to optimize performance within the Internet itself based on implementing an Internet overlay. An Internet overlay leverages service provider resources that are distributed throughout the Internet in order to optimize the performance, security, reliability, and visibility of the enterprise's Internet traffic. As shown in Figure 5, all client requests to the application's origin server in the data center are redirected via DNS to a server in a nearby point of presence (PoP) that is close to users of the application, typically within a single network hop. This edge server that is close to the users then optimizes the traffic flow to the server closest to the data center's origin server. Throughout this section, the Internet overlay that is depicted in Figure 5 will be referred to as an Application Delivery Service (ADS).

Figure 5: An Application Delivery Service

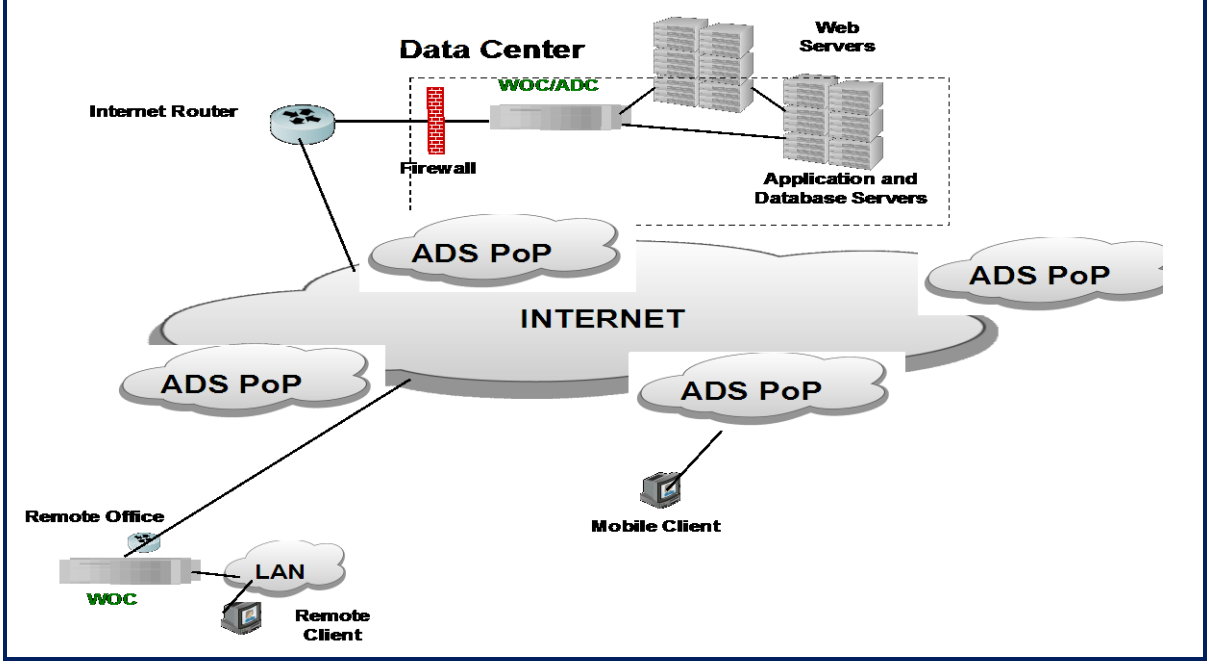

An ADS provides a variety of optimization functions that generally complements the functionality provided by WOCs and ADCs. One such function that is often provided by an ADS is content offload. This calls for taking static content out of a data-center and placing it in caches in servers and in replicated in-cloud storage facilities that are close to the users. Because the content is close to the users, IT organizations that offload content and storage improve response time and simultaneously reduce both their server utilization as well as the bandwidth utilization of their data center access links.

Some of the other functionality that is often associated with an ADS includes:

- Route optimization
- Transport optimization
- HTTP protocol optimization
- Visibility

In addition to the functionality listed above, some ADSs incorporate Web application firewall functionality.

One use case for an ADS that is growing in importance stems from that fact that not all CCSPs will support virtual WOC instances in their data centers. This is particularly true of SaaS providers. Access to services provided by a CCSP can be accelerated via an ADS. Even greater improvement in application delivery over the Internet can be achieved by migrating WOC functionality into ADS PoPs and by migrating ADS

functionality into enterprise WOCs. The result places WOC functionality very close[14] to the CCSP's data center and moves the ADS PoP right to the on-premise edge of the enterprise network. This helps optimize application and service delivery over the Internet between the branch office and the CCSP's site.

Another approach to the asymmetrical acceleration of access to CCSP services over the Internet was mentioned in the preceding subsection: caching content at the branch office or other enterprise site. As mentioned, once the content is stored in the cache for a single user, subsequent users will experience accelerated application delivery.

Dual ISP Internet VPN with Policy Based Routing

The preceding subsection of this report identified the concerns that IT organizations have with the use of the Internet. The two primary concerns are uptime and latency. Another approach to overcoming the limitations of the Internet is to connect each enterprise site to two ISPs. Having dual connections can enable IT organizations to add inexpensive WAN bandwidth and can dramatically improve the reliability and availability of the WAN.

For example, Figure 6 depicts a system that is composed of two components that are connected in parallel.

Figure 6: Two Components Connected in Parallel

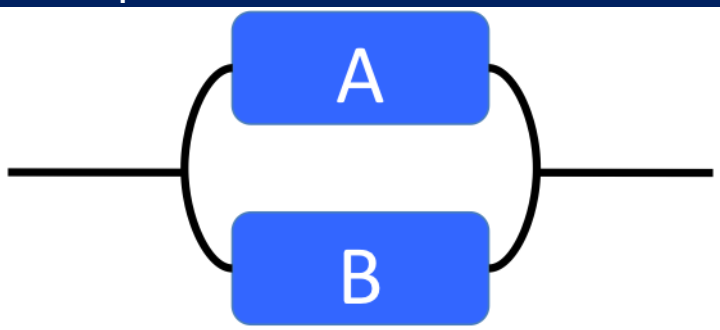

The system depicted in Figure 6 is available unless both of the two components are unavailable. Assuming that each component is a diversely routed DSL or cable access line and that one of the access lines has an availability of 99% and the other has an availability of 98%, then the system has an availability of 99.98%. Alternatively, if both access lines have an availability of 99%, then the system is available 99.99% of the time[15]. This level of availability is equal to or exceeds the availability of most MPLS networks.

Traffic can be shared by the two connections by using Policy Based Routing (PBR). When a router receives a packet, it normally decides where to forward it based on the destination address in the packet, which is then used to look up an entry in a routing

[14] While the closeness to the CCSP will vary based on the CCSP and the provider of the ADS, ideally the WOC functionality will be one hop away from the CCSP.
[15] If, as described later, 4G is added as a third access technique and if each access technique has an availability of 99%, then the system as a whole has an availability of 99.9999%.

table. Instead of routing by the destination address, policy-based routing allows network administrators to create routing policies to select the path for each packet based on factors such as the identity of a particular end system, the protocol or the application.

Perhaps the biggest limitation of the PBR approach it that it creates a static allocation of traffic to multiple links and it doesn't have the ability to reallocate the traffic when the quality of one of the links degrades. The static nature of the policies means unless there is an outage of one of the links, that a given class of traffic will always be allocated to the same network connection.

Dual ISPs and PBR can be used in conjunction with WOCs to further alleviate the shortcomings of Internet VPNs, bringing the service quality more in line with MPLS at a much lower cost point. For example, a WOC can classify the full range of enterprise applications, apply application acceleration and protocol optimization techniques, and shape available bandwidth in order to manage application performance in accordance with enterprise policies. As a result,

In many situations, a dual ISP-based Internet VPN with PBR can deliver a level of CoS and reliability that is comparable to that of MPLS at a significantly reduced price.

Part of the cultural challenge that IT organizations have relative to migrating traffic away from their MPLS network and onto an Internet based network is that Internet based networks don't provide a performance based SLA. However, as previously described, the majority of IT organizations don't place much value in the SLAs that they receive from their network service providers.

Hybrid WANs with Policy Based Routing

As noted, some IT organizations are reluctant to abandon traditional enterprise services such as MPLS. An alternative design that overcomes their concerns is a hybrid WAN that leverages multiple WAN services, such as traditional enterprise WAN services and the Internet, and which uses PBR for load sharing. The advantage of a hybrid WAN is that the CoS of MPLS can be leveraged for delay sensitive, business critical traffic with the Internet VPN used both for other traffic and as a backup for the MPLS network. As in the case of the dual ISP based Internet VPN, the major disadvantage of this approach is the static nature of the PBR forwarding policies. Since PBR cannot respond in real time to changing network conditions, it will consume more costly bandwidth than would a dynamic approach to traffic allocation. A second drawback of Hybrid WANs based on PBR is that they can prove to be overly complex for some IT departments. As with many other types of WAN services, hybrid WANs can also be used in conjunction with WOCs and ADCs.

Aggregated Virtual WANs

A relatively new class of device has emerged to address the shortcomings of PBR-based hybrid WANs. WAN path controller (WPC) is one phrase that is often used to describe devices that work in conjunction with WAN routers to simplify PBR and to

make selections of the best WAN access link or the best end-to-end WAN path from a number of WAN service options.

Some members of this emerging class of products are single-ended solutions whereby a device at a site focuses on distributing traffic across the site's access links on a per-flow basis. Typical capabilities in single-ended solutions include traffic prioritization and bandwidth reservation for specific applications. These products, however, lack an end-to-end view of the available paths and are hence limited to relatively static path selections.

In contrast, symmetrical or dual-ended solutions are capable of establishing an end-to-end view of all paths throughout the network between originating and terminating devices and these solutions can distribute traffic across access links and specific network paths based on either a packet-by-packet basis or a flow basis. These capabilities make the multiple physical WAN services that comprise a hybrid WAN appear to be a single *aggregated virtual WAN*.

Aggregated virtual WANs (avWANs) represent another technique for implementing WANs based on multiple WAN services (e.g., MPLS, Frame Relay and the Internet) and/or WANs based on just multiple Internet VPN connections. An aggregated virtual WAN transcends simple PBR by dynamically recognizing application traffic and allocating traffic across multiple paths through the WAN based on real-time traffic analytics, including:

- The instantaneous end-to-end performance of each available network: This allows the solution to choose the optimal network path for differing traffic types. One differentiator among virtual WAN solutions is whether the optimal path is chosen on a per packet basis or on a per flow basis. Per packet optimization has the advantage of being able to respond instantaneously to short term changes in network conditions.

- The instantaneous load for each end-to-end path: The load is weighted based on the business criticality of the application flows. This enables the solution to maximize the business value of the information that is transmitted.

- The characteristics of each application: This includes the type of traffic (e.g., real time, file transfer); the performance objectives for delay, jitter and packet loss; as well as the business criticality and information sensitivity.

One of the primary reasons why IT organizations backhaul their Internet traffic to a central site over an enterprise WAN service is because of security concerns. In order to mitigate those concerns when using an avWAN for direct Internet access, the avWAN should support security functionality such as encryption.

Like other hybrid WANs, an avWAN (Figure 7) allows IT organizations to add significant amounts of additional bandwidth to an existing MPLS-based WAN at a relatively low incremental cost. In addition to enabling the augmentation of an MPLS WAN with inexpensive Internet connectivity, aggregated virtual WANs also give IT organizations the option to reduce its monthly ongoing expense by either eliminating or reducing its MPLS connections while simultaneously providing more bandwidth than the original network design provided.

Figure 7: Aggregated Virtual WANs

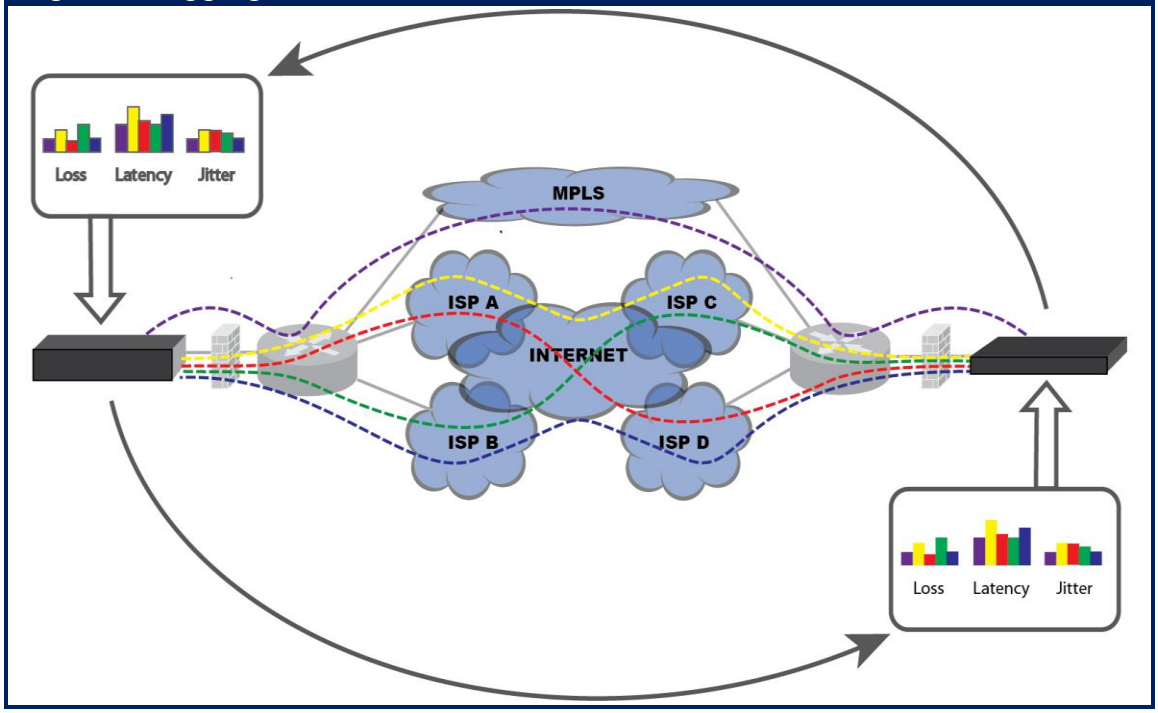

As shown in Figure 7, because the two avWAN appliances work together to continuously measure loss, latency, jitter and bandwidth utilization across all of the various paths between any 2 locations, an aggregated virtual WAN can rapidly switch traffic away from a path that is exhibiting an unacceptable level of performance. This capability, combined with the availability advantages of parallel systems as depicted in Figure 6, means that all of the bandwidth in each of the paths can be used most of the time, and that most of the bandwidth can be used virtually all of the time. This combination of capabilities also underscores the ability of aggregated virtual WANs to deliver performance predictability that equals, and in many cases exceeds, that of a single MPLS network.

Because of the high availability and performance predictability of aggregated virtual WANs, IT organizations can now leverage a number of WAN services that are dramatically lower in cost than traditional MPLS services. This includes DSL and cable Internet access from branch offices and fiber access to the Internet from data centers. It also positions IT organizations to take advantage of the huge volumes of very inexpensive Internet access bandwidth that are typically available at co-location facilities.

While the preceding discussion focused on DSL and cable access to the Internet it is important to realize that over the next year or two, there will be a broad scale deployment of 4G services on the part of most wireless service providers. There will be some variability in the effective bandwidth of 4G services based in part on the fact that the wireless service providers will not all implement the same technologies. It should generally be possible, however, for users of these services to realize throughput in the range of three to four megabits per second, which is roughly equivalent to two T1 or E1 access lines. This will make 4G services a viable access service for some

branch offices. For example, a 4G service could be combined with Internet access via DSL as part of a virtual WAN. In addition to providing cost savings, due to the inherent diverse routing associated with 4G and DSL, this design would provide a very high level of reliability.

There are three scenarios in which an avWAN offers significant benefits for accessing cloud computing services. The first scenario is for enterprises that have implemented centralized Internet access and who want to access a wide range of public cloud network services on the Internet, such as those offered by SaaS and IaaS providers. The implementation of an avWAN makes Internet backhaul far less expensive, higher capacity and more scalable than if the backhaul was done entirely using traditional enterprise WAN services. Between the remote sites and the central data center, all the backhauled traffic will benefit from the reliability, security, and QoS features of an avWAN, as shown in Figure 8. Between the central data center and public cloud services sites reliability will be improved because of the dual ISP connections (ISPs C and D), but the unique benefits of WAN virtualization will not generally be available on this portion of the end-to-end path because of the impracticality of the SaaS or IaaS CCSP provisioning a dedicated APN appliance for each service subscriber.

Figure 8: Scalable and Economical Internet Backhaul

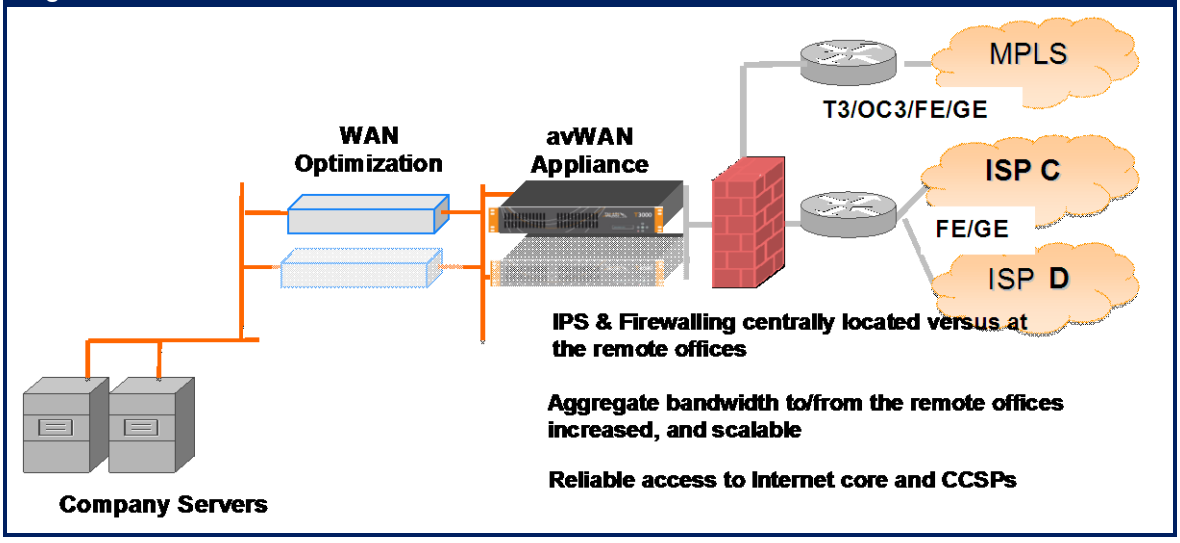

The second scenario is the situation in which the enterprise has implemented a private cloud computing data center, either at a central site managed by the enterprise or at a co-location facility. Here the full benefits of an avWAN can be derived for all of the traffic, including the traffic that goes between enterprise users and the private cloud resources as well as the server-to-server communications between the data centers. As shown in Figure 9, the economic advantages of an avWAN are further enhanced when the private cloud data center is built at a co-location facility where Internet access costs are generally considerably lower than at the enterprise central sites.

In this scenario, all Internet traffic can be backhauled to the enterprise's data center or to the co-location facility. If there are a number of geographically dispersed co-location sites, directing backhauled Internet traffic to the nearest site can minimize the extra propagation latency associated with centralized Internet access.

Figure 9: Private Cloud Data center at Colo Facility

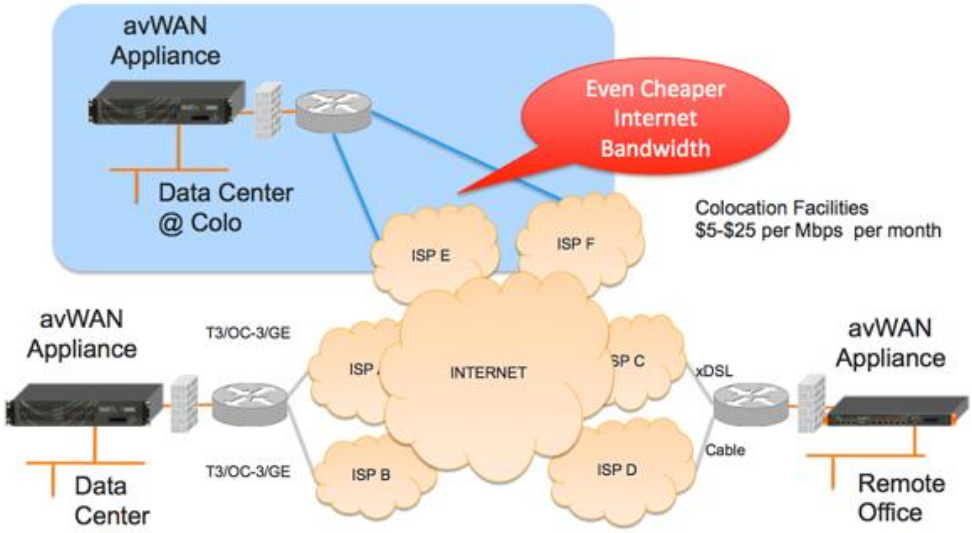

The third scenario is the situation in which the enterprise has subscribed to a public cloud based service, such as an outsourced private data center located on the CCSP premises or a Virtual Private Data Center hosted in a CCSP multi-tenant data center. With both types of data center services it should be possible to extend the avWAN to include the CCSP's data center by having avWAN appliances provisioned at these data center sites. Figure 10 shows one way that a hybrid cloud computing environment could be supported by an avWAN that is comprised of dual ISP connections at all sites plus an MPLS connection at some or all of the remote sites and at the enterprise's data center. As before, all general purpose Internet traffic could be economically backhauled to the enterprise data center, while all intra-enterprise cloud traffic can flow directly via the virtual WAN.

Figure 10: Hybrid Cloud Computing based on Aggregated Virtual WAN

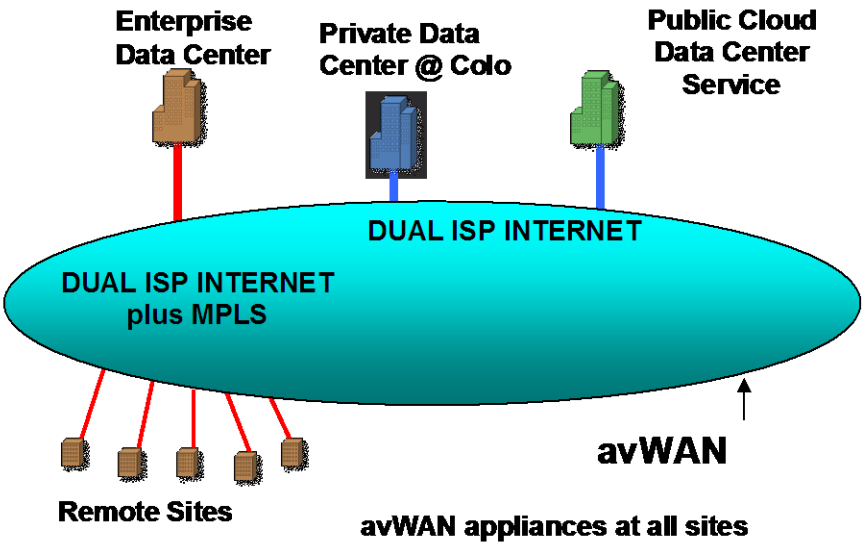

Cloud-Based Network and Application Optimization

As mentioned in the section of this report entitled *The Emergence of Cloud Computing and Cloud Networking*, network and application optimization has become available from CCSPs as a Cloud Networking Service (CNS). In this situation, instead of a physical or virtual WOC at each site, the WOC functionality is provided at the CCSP's cloud data centers or POPs, which ideally are in close proximity to the enterprise users, the data centers and the providers of other cloud services. As shown in Figure 11, the PoPs are interconnected by the CCSP's core network with customer access to each PoP provided via the Internet or via an enterprise WAN service. The CNS core network could be an Internet overlay, a private IP network or possibly a multi-carrier MPLS/IP network that uses intelligent routing capabilities similar to an aggregated virtual WAN or ADS in order to provide high levels of performance and reliability.

Figure 11: Network and Application Optimization CNS

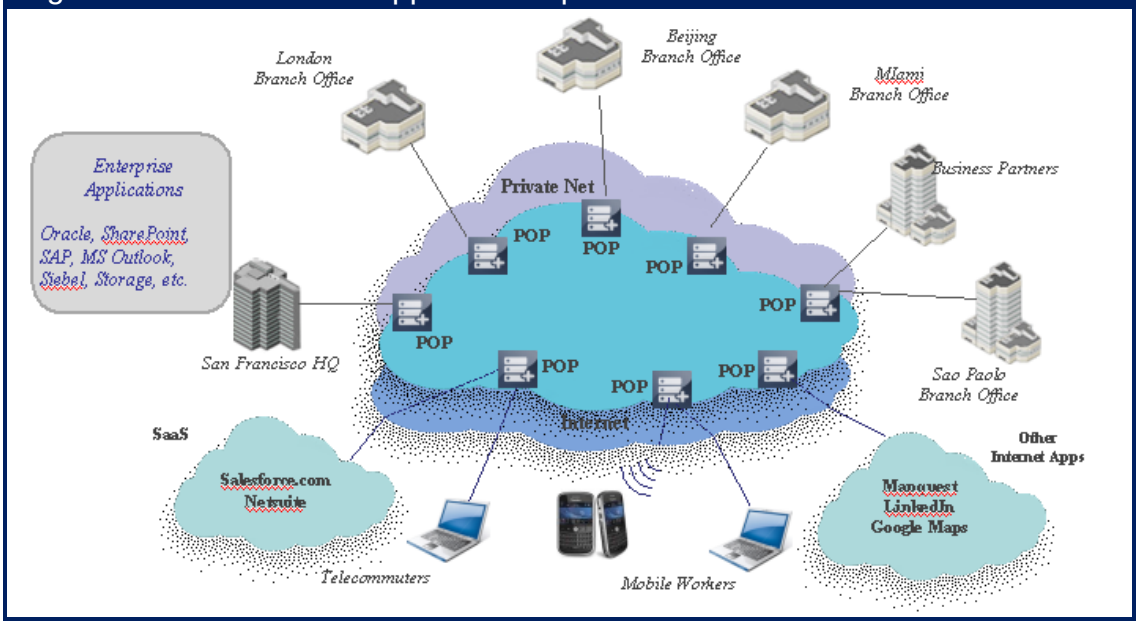

Therefore, a network and application optimization CNS can be considered to be another form of alternative WAN service that is layered on top of the Internet or existing private WAN services from one or more carriers. The key differentiation is that this class of WAN service has WAN Optimization built into the PoPs. The form of a CNS can be used for the traditional challenge of optimizing communications between users in a branch office and IT services that are provided at the enterprise's data centers. This form of a CNS can also be used in those situations in which installing WOC functionality is either not economical or it is problematical; e.g., at home offices, small branch offices, IaaS data centers, SaaS data centers or on mobile devices.

VPLS

As previously mentioned:

VPLS represents the combination of Ethernet and MPLS.

While VPLS is not widely implemented today, the data in Table 4 indicates that more than a third of IT organizations will increase their use of VPLS over the next year.

VPLS is a class of VPN that supports the connection of customer edge (CE) Layer 2 switches at multiple sites into a single bridged, multipoint-to-multipoint domain over a service provider's IP/MPLS network, as shown in Figure 12. VPLS presents an Ethernet interface to customers that simplifies the LAN/WAN boundary for Service Providers and customers, and enables rapid and flexible service provisioning. All sites in a VPLS appear to be on the same LAN, regardless of location. A companion technology, Virtual Private Wire Services (VPWS), provides point-to point services.

Figure 12: A VPLS Service Linking Four Customer Sites

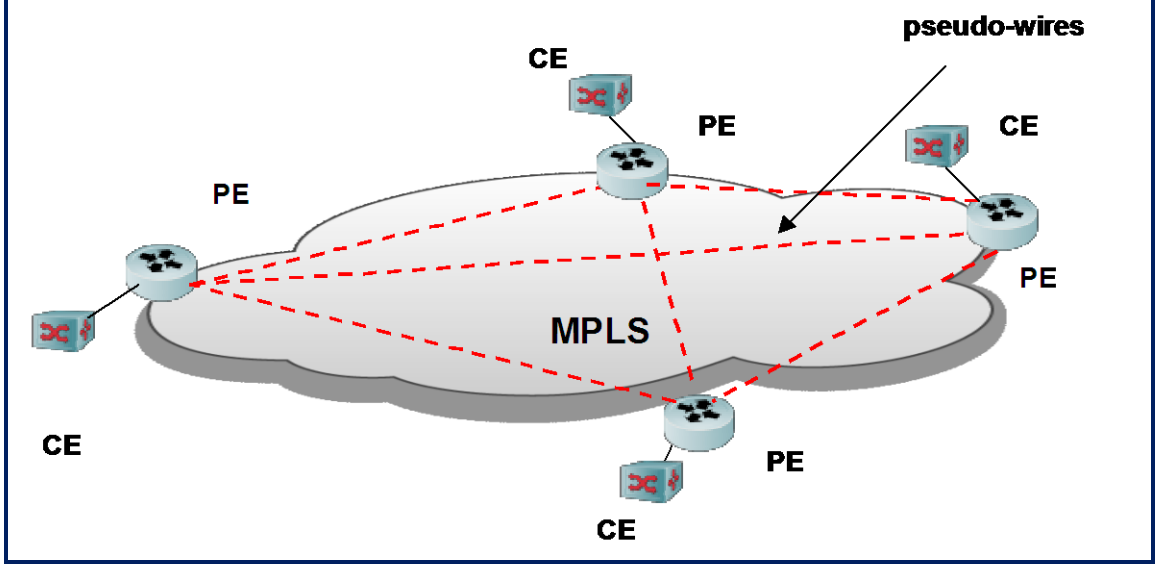

With VPLS, either the Border Gateway Protocol (BGP) or the Label Distribution Protocol (LDP) is used to create the required pseudo-wires to fully mesh the provider edge (PE) devices serving the customer sites. Meshed pseudo-wires support the multipoint-to-multipoint nature of the virtual LAN and improve reliability. Reliability is enhanced because in case of failure in the MPLS network, traffic will automatically be routed along available backup paths, providing very short failover times.

VPLS MPLS packets have a two-label stack. The outer label is used for normal MPLS forwarding in the service provider's network. If BGP is used to establish the VPLS, the inner label is allocated by a PE as part of a label block. If LDP is used, the inner label is a virtual circuit ID assigned by LDP when it first establishes a mesh between the participating PEs. Every PE keeps track of assigned inner label, and associates these labels with the VPLS instance.

Table 9 provides a high level comparison of the different types of Ethernet WAN services available for LAN extension between data centers . It should be noted that there are other options for LAN extension, such as Ethernet over leased dark fiber and Ethernet over GRE tunneling through a private IP network. As described previously, VXLAN is VM-specific overlay solution for LAN extension.

Table 9: Ethernet WAN Service Types

Service Topology	Access Link	Provider Core	Service Type	Tunneling
Ethernet end-end	Ethernet	Ethernet	Pt-Pt or Mpt-Mpt	802.1Q or Q in Q
Ethernet/IP	Ethernet	IP	Pt-Pt or Mpt-Mpt	L2TPv3
VPLS/VPWS	Ethernet	MPLS	Pt-Pt or Mpt-Mpt	EoMPLS

Emerging Cloud Networking Specific Solutions

The preceding discussion of WAN services provided some insight into the interplay between the general requirements of cloud computing and the capabilities of WAN services to meet those requirements. One of the goals of this section of the report is to describe the functionality that is required to support a particular form of hybrid cloud computing – cloud balancing. Another goal of this section of the report is to describe some of the optimization functionality that is being developed specifically to support cloud computing.

Cloud Balancing

As previously described, a hybrid cloud relies on a WAN to provide the connectivity between the enterprise's locations, including the enterprise's data center(s) and its remote sites, and the public cloud data center(s) that is providing the IaaS or other cloud service. One of the goals of cloud balancing is to have the collection of individual data centers appear to both users and administrators as a single cloud data center, with the physical location of application resources as transparent as possible. The goal of having the location of application resources be transparent creates a number of requirements. This includes:

- **VLAN Extension**
  As is the case for private clouds, hybrid clouds depend heavily on VM migration among geographically dispersed servers connected by a WAN in order to ensure high availability and dynamic response to changes in user demand for services. The VLANs within which VMs are migrated must be extended over the WAN between and amongst the private and public data centers. This involves the creation of an overlay network that allows the Layer 2 VLAN traffic to be bridged or tunneled through the WAN.

- **Secure Tunnels**
  These tunnels must provide an adequate level of security for all the required data flows over the Internet. For the highest level of security, this would typically involve both authentication and encryption, such as that provided by IPsec tunnels.

- **Universal Access to Central Services**
  All application services, such as load balancing, DNS, and LDAP, should be available and function transparently throughout the hybrid cloud. This enhances security as well as transparency by allowing these application services to be provisioned from the private enterprise data center and eliminating manual intervention to modify server configurations as the application and its VM are transferred from the private cloud to the public cloud.

- **Application Performance Optimization**
  Application performance must meet user expectations regardless of the location of the users or the IT resources that the users are accessing. This means that the public cloud data centers need to offer the same WAN optimization and application acceleration capabilities that are deployed within the enterprise. In addition, WOCs may well be needed between the enterprise's private cloud data center(s) and the public cloud data center(s) in order to accelerate VM migration, system backups, and other bulk data transfers between these data centers.

- **Interoperability Between Local and Global ADC Functions**
  Cloud balancing is based on making routing decisions based on a combination of local and global variables. This requires interoperability between local and global ADC functions.

- **Synchronizing Data between Cloud Sites**
  In order for an application to be executed at the data center that is selected by the cloud balancing system, the target server instance must have access to the relevant data. In some cases, the data can be accessed from a single central repository. In other cases, the data needs to co-located with the application. The co-location of data can be achieved by migrating the data to the appropriate data center, a task that typically requires highly effective optimization techniques. In addition, if the data is replicated for simultaneous use at multiple cloud locations, the data needs to be synchronized via active-active storage replication, which is highly sensitive to WAN latency.

WAN Optimization and Application Delivery for Cloud Sites

One of the most significant trends in the WAN optimization market is the development of new products and new product features that are designed to enable IT organizations to leverage public and hybrid clouds as extensions of their enterprise data centers. Some recent and anticipated developments include:

Cloud Optimized WOCs: These are purpose-built virtual WOC appliances for deployment in public cloud environments. Cloud Optimized features include compatibility with cloud virtualization environments, SSL encryption and acceleration, and automated migration or reconfiguration of virtual WOCs in conjunction with VM provisioning or migration. As previously mentioned, WOCs can either be deployed in a

symmetric fashion, with a WOC on each end of the WAN link; or in an asymmetric fashion, with a WOC deployed just in a branch office.

Cloud Storage Optimized WOCs: These are purpose-built virtual or physical WOC appliances for deployment in the enterprise's data center(s) and also at public cloud Storage as a Service environments that are used for backup and archival storage. Cloud optimized features can include support for major backup and archiving tools, de-duplication to minimize the required data transfer bandwidth and the storage capacity that is required, and support for SSL and AES encryption.

Data Mobility Controller Enhancements: Data Mobility Controllers (DMCs) facilitate the transfer of high volume data between enterprise data centers or private cloud data centers. DMC products are still in a early stage of evolution and a number of developments can be expected in this space, including enhanced hardware support for various functions including encryption and higher speed WAN and LAN interfaces at 10 GbE and higher in order to support a combination of highly efficient data reduction and high bandwidth WAN services.

Cloud Optimized Application Delivery Controllers: One trend in the evolution of ADCs is increasing functional integration with more data center service delivery functions. As organizations embrace cloud computing models, service levels need to be assured irrespective of where the applications are hosted. As is the situaion with WOCs, ADC vendors are in the process of adding enhancements that support the various forms of cloud computing, including:

- **Hypervisor–based Multi-tenant ADC Appliances:** Partitioned ADC hardware appliances have for some time allowed service providers to support a multi-tenant server infrastructure by dedicating a single partition to each tenant. Enhanced tenant isolation in cloud environments can be achieved by adding hypervisor functionality to the ADC appliance and by dedicating an ADC instance to each tenant. Each ADC instance is then afforded the same type of isolation as a virtualized server instance, with protected system resources and address space. A combination of hardware appliances, virtualized hardware appliances and virtual appliances provides the flexibility for a cloud service provider to offer highly customized ADC services that are a seamless extension of an enterprise customer's IT environment.

- **Cloud Bursting and Cloud Balancing ADCs:** Cloud bursting refers to directing user requests to an external cloud when the enterprise private cloud is at or near capacity. Cloud balancing refers to routing user requests to application instances deployed in the various different clouds within a hybrid cloud. Cloud balancing requires a context-aware load balancing decision based on a wide range of business metrics and technical metrics characterizing the state of the extended infrastructure. By comparison, cloud bursting can involves smaller set of variables and may be configured with a pre-determined routing decision. However, cloud bursting may require rapid activation of instances at the remote cloud site or possibly the transfer of instances among cloud sites. Cloud bursting and balancing can work well where there is consistent application delivery architecture that spans all of the clouds in question. This basically means that the enterprise's application delivery solution is replicated in the public cloud. One way to achieve this is with virtual appliance implementations of GSLBs and ADCs that support the range of

variables needed for cloud balancing or bursting. If these virtual appliancess support the IaaS cloud hypervisors, they can be deployed as VMs at each cloud site. The architectural consistency insures that each cloud site will be able to provide the information needed to make global cloud balancing routing decisions. When architectural consistency extends to the hypervisors across the cloud, integration of cloud balancing/bursting ADCs with the hypervisors management systems can help the routing of application traffic synchronized with private and public cloud resource availability and performance. Access control systems integrated within the GSLB and ADC make it possible to maintain control of applications wherever they reside in the hybrid cloud.

The following table summarizes the applicability of the various WAN services and optimization solutions to different types of cloud computing.

Table 10: Applicability of WAN Technologies in the Cloud

	Private Cloud	Hybrid Cloud	IaaS	ISV SaaS
MPLS	Yes	Yes	Depends on the Service Provider	Depends on the Service Provider
Internet	Yes	Yes	Yes	Yes
Internet Overlay	Yes	Yes	Yes	Yes
Hybrid WAN	Yes	Yes	Depends on the Service Provider	Depends on the Service Provider
WAN Virtualization	Yes	Yes	No	No
VPLS	Yes	Yes	Depends on the Service Provider	Depends on the Service Provider
Cloud Bridging	No	Yes	No	No
WOC/ADC	Yes	Yes	No	No
VA WOC/ADC	Yes	Yes	Yes	No
Cloud WOC	Yes	Yes	Yes	Yes
Storage Service WOC	No	No	Yes	No
DMC	Yes	Yes	Yes	No
Hypervisor ADC	In multi-tenant data centers	Yes	Yes	No
Cloud Bursting/ Balancing ADC	Yes	Yes	No	No
Cloud Balancing	Yes	Yes	No	No

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.

For editorial and sponsorship information, contact Jim Metzler or Steven Taylor. The Webtorials Editorial/Analyst Division is an analyst and consulting joint venture of Steven Taylor and Jim Metzler.

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.

Giving Cloud Applications a Lift

Software-as-a-service (SaaS) enables businesses to quickly innovate and compete in worldwide markets while lowering IT costs. From ubiquitous Microsoft SharePoint for robust content management, to customer relationship management by Salesforce.com, to hosted Windows Media Services pushing out multimedia communications and training; SaaS powers business productivity anywhere in the world.

As SaaS platforms, however, these applications are beyond the bounds of IT control. As a result, latency, chatty protocols, and packet loss easily impact SaaS performance, inhibiting business productivity and competiveness. WAN optimization is essential to assure reliable user experiences, yet the conventional symmetric approach of deploying appliances at both ends of the transaction does not apply for cloud-based applications.

Asymmetric Advantage

Cloud-based SaaS utilizes common protocols (e.g. HTTP, SSL, TCP, FTP) to deliver applications and data over WAN/Internet connections. Protocols, such as HTTP/SSL that secures data over the internet/WAN, can be chatty. When combined with the latency, limited bandwidth and packet loss associated with the WAN, these protocols can reduce SaaS performance. Add to this the bandwidth demands caused by the rapid growth in the use of video and other multi-megabyte files within the SaaS infrastructure; it's easy to see how performance degradation compounds particularly for remote sites with limited bandwidth.

Blue Coat offers a different WAN optimization approach. With its asymmetric CloudCaching Engine, Blue Coat is able to overcome latency, chatty protocols, limited bandwidth and packet loss with 3x -110x performance improvements on the same WAN/Internet connectivity. Blue Coat's asymmetric WAN Optimization technology is able to reduce transfer times and the bandwidth consumed by accessing SaaS infrastructure.

Accelerate and Optimize Cloud SaaS Applications

Blue Coat MACH5 WAN optimization is a combination of five separate application management and tuning technologies that provide unrivaled improvements in application performance and bandwidth utilization. These technologies include:

1. **Bandwidth Management** assigns priority and network resources based not only on port or device, but on users, applications and content to more accurately reflect corporate policies on the network. Whether alone or integrated with network QoS, bandwidth management provides application intelligence to the packet switching network.
2. **Protocol Optimization** improves the performance of protocols that are inefficient over the WAN through specific enhancements that improve tolerance to the higher latencies occurring there. Blue Coat offers multiple improvements for TCP, CIFS, HTTP, HTTPS, MAPI and most streaming video and IM protocols.
3. **Byte Caching** stores repetitive traffic found in the byte stream and serves it locally to reduce the amount of traffic that traverses the WAN. Requests can be served from a local WAN optimization appliance, accelerating updates to existing files with significantly less bandwidth; leading to dramatic bandwidth savings.
4. **Object Caching** further reduces demands on bandwidth by storing and serving files, videos and web content locally, without the overhead and risk of traditional wide area file services. For content delivery, no technology does more to improve the end user experience by reducing latency and bandwidth.

5. Finally, inline **Compression** can reduce predictable patterns even on the first pass, making it an ideal complement to byte caching technology.

Blue Coat's asymmetric WAN optimization technology can also be combined with industry leading QoS and visibility, to ensure that SaaS application data and files are transferred with the right priority and sufficient bandwidth to avoid competition with recreational or non-critical traffic.

Setting the SaaS Example

To demonstrate these performance gains, tests were conducted using a simulated 1.544 Mbps (T1) WAN link with 100ms latency to simulate a Cloud SaaS (data center) to branch office environment. Results show access times between the branch and Cloud for Salesforce.com are reduced by up to 33x. SharePoint BPOS access via the data center or direct-to-net is reduced by 7x – 110x. WAN bandwidth demand required to stream video from Microsoft Media Services Servers was reduced by 65% - 95%, and enabled up to 100 users to view streaming video simultaneously. Even files saved over WAFS show dramatically reduced transfer times (see Figure 1):

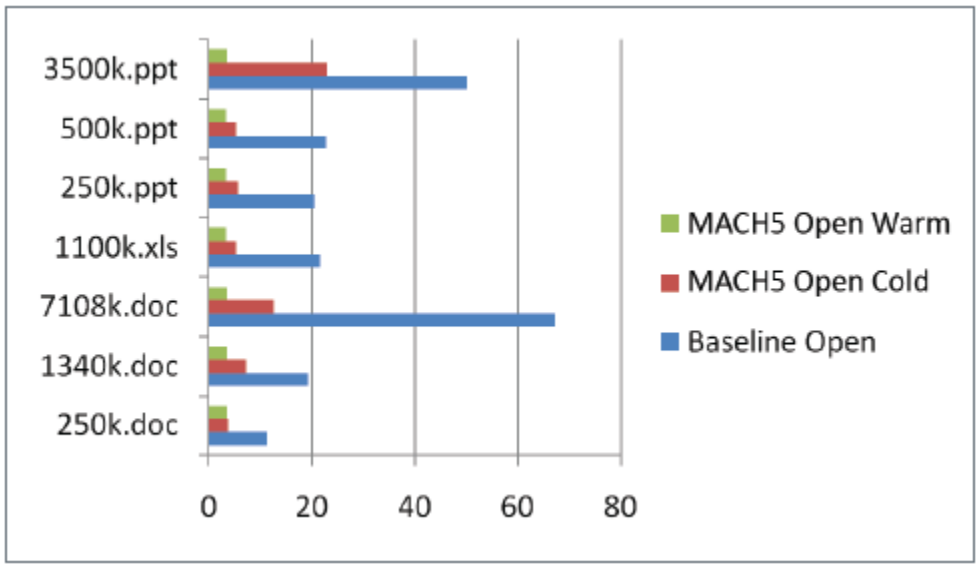

Figure 1: PowerPoint, Excel and Word file open times with and without MACH5 WAN optimization

The right choice for Cloud SaaS

Blue Coat Systems is a leading provider of web security and WAN Optimization solutions that secure and optimize the flow of information to any user, on any network, anywhere. Video and SaaS application delivery are today's IT's challenges, and with the right acceleration strategy you can gain superior business value from your internal and external infrastructure. Find out how Blue Coat can help you at bluecoat.com.

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.

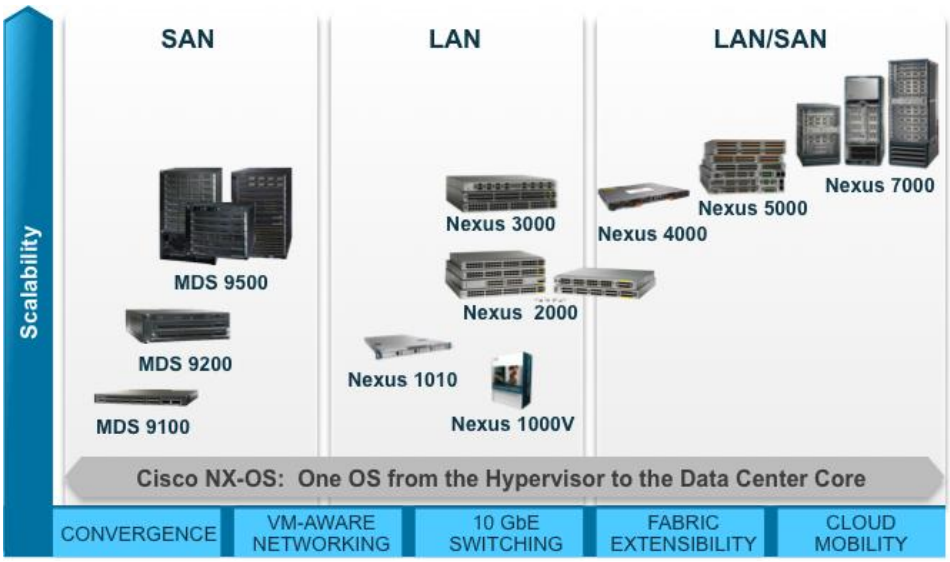

- Revolutionary fabric scale with over twelve thousand 10 GbE server connectivity with Cisco Nexus
- Highest 10Gb Ethernet density in the industry with Cisco Nexus 7000
- High performance and ultra-low latency networking at scale with Cisco Nexus
- Network services delivered in virtual and physical form 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.

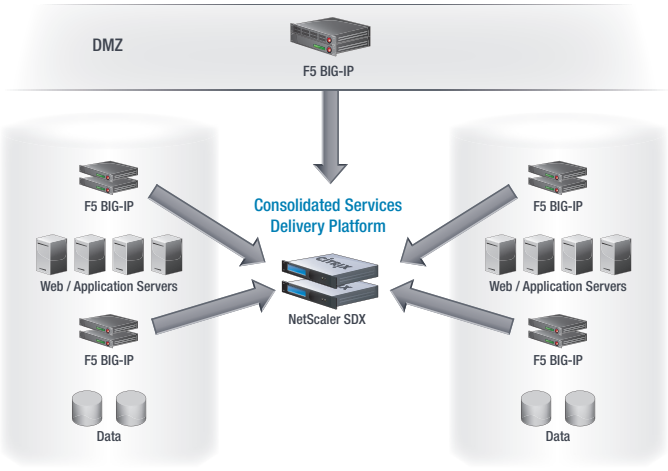

Network Services Consolidation

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

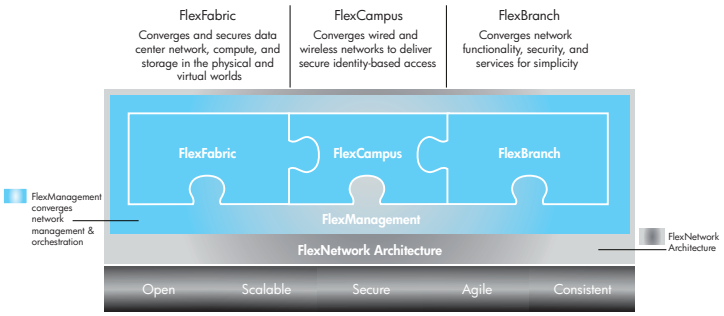

HP FlexFabric and **HP FlexCampus** enable the construction of flat, low-latency data center and campus networks with fewer layers, less equipment and cabling, and greater port densities.

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.

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

IT Organizations Find Key Enabling Technologies for Adopting Cloud Architectures

With engineers and administrators at companies like Google and Amazon redefining the standards for application, infrastructure, and data center efficiency, IT organizations have begun to reexamine their internal operations in order to apply the lessons of cloud computing. What they have discovered is that cloud computing is not a technology that can be applied, but an architecture that is built from many existing components and key enabling technologies.

Those key technologies support the centralization and consolidation of infrastructure, as well as the automation of IT processes, such as provisioning and scaling capacity. It is an architecture that favors economies of scale – such scale that for certain types of workloads, the most attractive and cost-effective deployment option is with third-party cloud providers. While many organizations initially hesitate to deploy their applications and store their data on the shared infrastructure of a public cloud provider, organizations that ultimately adopt third-party services recognize that shifting the burden of infrastructure administration to a provider operating at massive scale not only yields cost savings, but frees IT personnel to focus on more differentiated technology efforts.

As a result, the advent of cloud computing offers new choices in architecting IT infrastructure for the best possible blend of performance, availability, cost, and control. Finding that optimal balance will require both consolidation to fewer data centers and migration of selected applications and data to more cost efficient public cloud services.

> “Regardless of whether an organization chooses a private, public or hybrid cloud approach, they will likely experience performance problems as they encounter challenges caused by distance and the sheer growth of data.”
>
> *Eric Wolford, executive vice president of marketing and business development, Riverbed*

After identifying which applications are candidates to centralize into a consolidated private data center and which are candidates to move to a public cloud service, organizations must consider what their existing infrastructure supports and what new requirements will emerge. For example, centralizing resources and adopting public cloud services inherently requires users to depend on a network connection when accessing data and applications. However, migrating data and accessing applications across the WAN or public Internet is negatively impacted by distance, which introduces latency, as well as bandwidth congestion. For that reason, WAN optimization, with it’s ability to reduce data traffic and accelerate applications, is one of those key enabling technologies of cloud computing, by supporting the movement of infrastructure from inefficient, distributed models, to highly-automated, centralized cloud models. Not all WAN optimization vendors support the full spectrum of deployment scenarios that may make up an organization’s mix of private and public resources, but Riverbed Technology is one such vendor that has made it’s Steelhead product available in cloud deployments as well as traditional private WAN environments. “Regardless of whether an organization chooses a private, public or hybrid cloud approach, they will likely experience performance problems as they encounter challenges caused by distance and the sheer growth of data,” said Eric Wolford, executive vice president of marketing and business development at Riverbed.

Another challenge is in shifting legacy stove-piped application deployments to take full advantage of virtualized, scaled-out cloud architectures. Converting the application into a virtual machine is a critical step, but that alone does not ensure that an application can scale to more capacity and seamlessly migrate across available cloud resources. Application delivery controllers, which encompass traditional load-balancing functionality, can add a point of flexibility in an application’s architecture to allow organizations to seamlessly and automatically add additional server capacity to an application without disrupting its availability. Similarly, organizations can take advantage of hybrid cloud cost efficiencies by deploying an application across multiple public and private cloud data centers and using global load balancing technologies to manage application traffic across these multiple cloud deployments, reducing risk and improving the performance and capacity of applications.

However, a physical application delivery controller appliance tethers an application, even a virtualized one, to a limited set of resources in the data center. Thus, only a software-based virtual application delivery controller provides the flexibility necessary to enable cloud computing. As a virtual appliance, it can seamlessly migrate with a virtual application across available resources, within a single data center or between cloud data centers operated by different entities.

Transitioning to a cloud architecture, whether public or private, means applications run in massive, virtualized data centers. There are necessarily fewer of them and they will be farther apart and farther from end users. Thus, part of the transition to cloud computing is using enabling technologies to overcome the inherent challenges to running in virtual environments across wide distance.

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

WAN Virtualization Reduces Costs by 40% to 90%, Significantly Increases Bandwidth and Improves Reliability

As a CIO or IT manager responsible for network architecture, you may have connected more branch offices over recent years, or consolidated your data centers. As a result, you've witnessed first-hand the phenomenon that as servers move farther away from users, more WAN traffic is generated.

Also adding to your WAN traffic are the increased use of latency-sensitive applications, like VoIP, videoconferencing and desktop virtualization.

Because you don't want to hear unnecessary complaints when VoIP calls drop or applications perform poorly, you've likely purchased very expensive leased lines or MPLS services to ensure scalable, reliable and predictable WAN connectivity. Although alternative connectivity choices (e.g., Internet, DSL, etc.) are extremely attractive from a cost point of view, they simply don't provide the necessary four nines reliability to keep your business-critical applications up and running 24X7.

Into this carrier-pricing environment where a price/performance factor of 2x is enormous enters WAN Virtualization via Adaptive Private Networking (APN) technology from Talari Networks. WAN Virtualization brings Moore's Law and Internet economics to enterprise WAN buyers for the first time in 15-plus years. Further, Talari's Mercury appliances do this incrementally and seamlessly on top of existing networks – no forklift upgrades required.

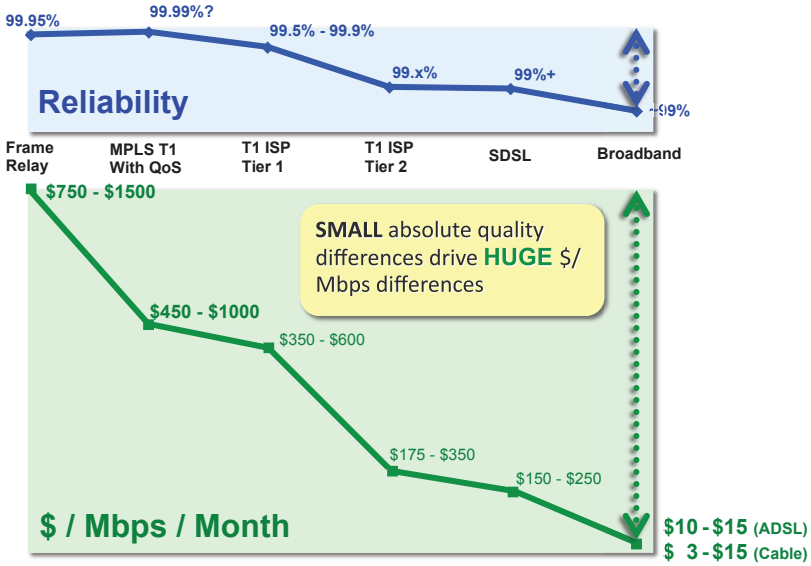

Figure 1: Private / Public WAN Pricing Disparity

Talari Networks Customer's 'AHA' Moment

Tim Hays at Lextron Inc. has used what is now called "cloud computing" in his network for over a decade. After he deployed Talari's solution, he said, "That was an 'aha' moment for me because I thought, 'Somebody finally gets it.' Talari's Adaptive Private Networking technology allows me to route each packet over the best, most reliable route, over multiple paths, including private lines, MPLS, DSL, and cable modem. By using WAN Virtualization, we've essentially created our own, big, private tunnel that aggregates different types of connectivity transparently across the Internet."

Real-Time, Per-Packet Traffic Engineering

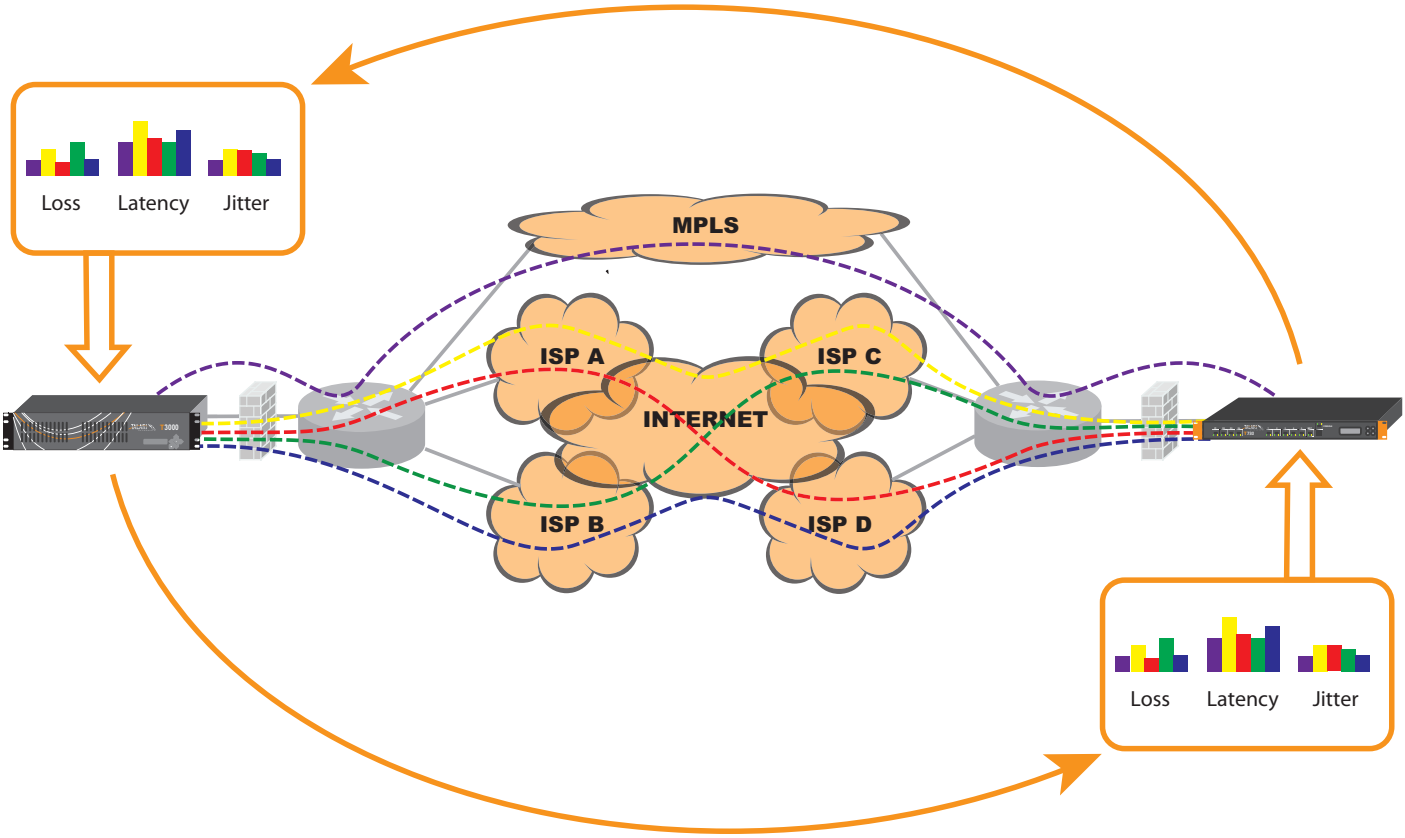

Figure 2: Continuous Measurment and Adaptation to Network Conditions

Requiring only two IP connections at each site which can include an existing private WAN connection, WAN Virtualization combines a variety of networks into a virtual WAN to deliver packets without being lost or excessively delayed 99.99% of the time. All network paths between locations are continually measured to determine current conditions. This allows each and every packet to be sent on the most appropriate path as determined by the type of traffic and available network resources. In addition, sub-second response to any congestion detected ensures predictable performance for all applications.

With this approach, Talari customers are building WANs where:

- **30 to 100 times more bandwidth can be purchased for every dollar spent**
- **Ongoing monthly WAN service charges can be reduced by 40% to 90%**
- **The resulting network is more reliable than any single MPLS private WAN**
- **Public cloud resources can be accessed with high reliability**

An APN Appliance for Every Situation

The Mercury family of APN appliances offer a wide range of perfomance points that span from large data centers to small remote offices and can be seamlessly added to your existing network in an overlay configuration to leave your current routed infrastructure intact. This allows you to introduce WAN Virtualization at your own pace to eventually migrate some or all of your locations off expensive private WAN connections.

Talari's customers see significant reductions in their ongoing monthly WAN expense that results in payback times for their WAN Virtualization deployments in the range of 6 to 12 months.

To learn more about how WAN Virtualization can transform the economics of your WAN please contact Talari Networks: **www.talari.com**.

VISUAL NETWORK SYSTEMS™

Application Performance Management in the Cloud

By 2016, 41% of all enterprise communications application users worldwide will have migrated to the cloud according to a study by ABI Research. That translates into trillions of dollars in business revenue depending on the delivery of these services. Application performance management will become even more critical to daily operations, but a surprisingly small number of cloud-based application users have adequate performance management software monitoring service delivery today.

In today's non-Cloud environment, with the current technology for application performance management, it is possible to instrument and collect run-time metrics, and provide access to management tools to analyze and report on the metrics. It is also possible to obtain a comprehensive view of an application including end user experience, specific transactions, and the supporting delivery infrastructure in order to manage the availability and performance of a business service.

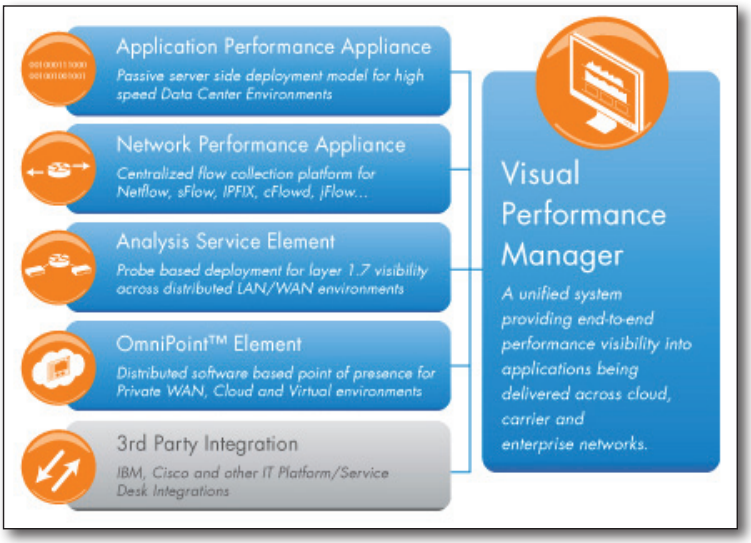

Visual Performance Manager is a unified system providing end-to-end performance visibility into applications being delivered across cloud, carrier and enterprise networks

When parts or all of an application moves to a Cloud, the view into the application is disrupted. One thing that doesn't change for both Cloud and non-Cloud environments is that the users, representing the business, expect the same level of availability, access to the applications and performance. Here in lies the challenge.

At a high level, the Cloud infrastructure includes the application delivery infrastructure which is made up of the applications running in virtualized environments and the network that supports the delivery of the applications to users. While the infrastructure itself is still made up of switches, routers, firewalls WAN optimization devices, VM Hosts and servers; the new part is that there could be more than one owner, such as cloud service providers and the private enterprise, for different parts of the delivery infrastructure.

For an Enterprise IT organization that is building a private cloud and virtualizing applications, or migrating to a hybrid cloud, the following are a few criteria that can help you find the right solution as you evaluate application performance management products.

- **Bridge application and network performance management between cloud and non-cloud environments –** More than likely, you are not moving all applications to the cloud. Whether you are in a transition to migrate applications or simply maintaining both cloud and non-cloud based applications, you are presented with the challenge of managing availability and performance for both sets of applications.

- **Flexible data collection instrumentation -** Within your private cloud, the challenge is visibility of applications in a virtualized environment. It is important that the instrumentation allows you to measure the performance of multi-tier applications as well as providing you with transaction level information for root cause analysis when performance degrades. This requires supporting deployment models to see the intra virtual machines traffic.

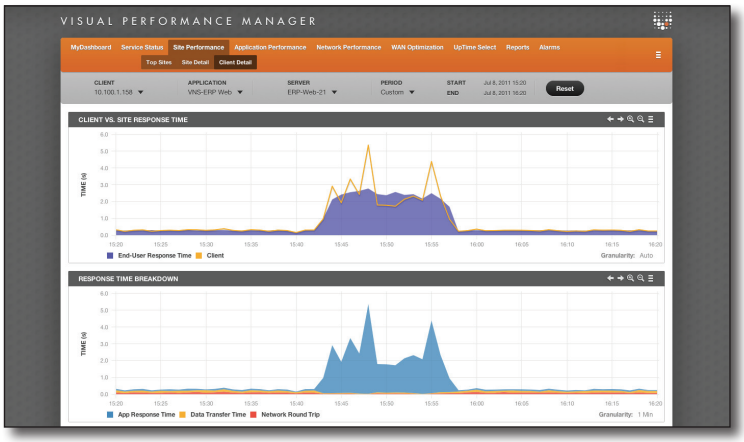

Managing end user experience with Visual Performance Manager. Quickly isolate user problem at a remote site.

- **Future proof and scalable architecture** - As with any new technology, you will need new information that you do not know about today. The chosen solution needs to be extensible to support new relevant performance metrics without having to do a mass rip and replace. A proven scalable architecture is important especially if you are managing many remote offices. For the IT team to be effective, the architecture needs to be able to support mediating a variety of data sources in your delivery infrastructure and correlating performance metrics to provide a comprehensive view of application performance.

- **Establish service level agreements with your Cloud provider** - For Enterprise IT using hybrid cloud environments, in addition to visibility of the application performance in your private cloud, you should be demanding that your Cloud service provider establish service level agreements and prove that application services are delivered according to availability and performance objectives.

SECURE & CONTROL ANY CLOUD ENVIRONMENT

The Vyatta Network OS is industry's most complete virtual networking & security solution. Optimized for VMware, XenServer, Red Hat KVM and in the Amazon Cloud, Vyatta enables you to segment, isolate and secure applications and data in any cloud environment.

- Stateful Firewall
- IPSec and SSL VPN
- Intrusion Prevention
- Dynamic Routing
- Web Filtering
- Management API

...and more

VYATTA®

http://www.vyatta.com