The 2010
Cloud Networking Report

Dr. Jim Metzler, Ashton Metzler & Associates

The Wide Area Network

A10 Networks

certeon® Accelerate Your Business

CiTRiX®

ipanema Technologies

riverbed

STREAMCORE MAKE YOUR NETWORK CONSCIOUS

TALARI NETWORKS

VYATTA® Open Networking

Webtorials

The Wide Area Network (WAN)

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-year period:

Today there is not a 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). As described below, 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.

The Webtorials Respondents were given a set of eleven WAN services[2] and asked to indicate the extent to which they currently utilize each WAN service. They were given a five-point scale:

1. None
2. Minimal
3. Some
4. Quite a bit
5. Extensive

[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] The eleven WAN services are listed in column one of Table 5.2.

The survey results indicate that The Webtorials Respondents utilize on average 4.8 WAN services.

The typical IT organization utilizes a wide range of WAN services.

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 Survey Respondents don’t have any Frame Relay in their networks and almost two thirds of The Survey Respondents don’t have any ATM in their networks.

Looking just at the WAN services that The Webtorials Respondents utilize either quite a bit or extensively, they average 2.3 WAN services.

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

While IT organizations make extensive use of the Internet, quality issues in the public Internet and in consumer-class ISP services generally prevent Internet VPNs from meeting the reliability standards of enterprise IT departments. As a result, Internet VPNs are most 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. In cases where Internet-based VPNs are deployed today, businesses typically prefer an expensive T1/E1 for access, since a single xDSL often results in a lower level of availability and performance predictability.

In many cases, Internet-based VPNs that use DSL for access are ‘good enough’.

Drivers of Change

The Webtorials Respondents were asked to indicate the anticipated change in their organization's use of the eleven WAN services. For each WAN service, Table 5.1 indicates the percentage of The Webtorials Respondents that expect to reduce their use of that service; to make no change in the use of that service; and to increase their use of that service.

Service	% Decrease	% Stay the Same	% Increase
Private lines between your sites	26.2%	62.7%	11.1%
Private lines to access network services	20.7%	62.8%	16.5%
Frame Relay	33.9%	64.4%	1.7%
ATM	22.3%	71.4%	6.3%
MPLS	4.8%	43.2%	52.0%
VPLS	8.0%	68.1%	23.9%
Internet-based VPNs with T1/T3/OC-3/OC-12 for access	12.4%	64.5%	23.1%
Internet-based VPNs with DSL or cable for access	10.4%	52.0%	37.6%
IP VPN	4.9%	61.5%	33.6%
An Internet overlay from a company like Akamai or CDNetworks	6.3%	83.0%	10.7%
Internet traffic to external sites	2.4%	57.3%	40.3%

Table 5.1: Expected Change in the Use of WAN Services

As previously noted, the primary WAN services used by IT organizations are MPLS and the Internet. As shown in Table 5.1, the majority of IT organizations are expecting to increase their use of MPLS. In addition, the majority of IT organizations are also increasing their use of one or more forms of Internet services.

One form of centralization of resources that is likely to drive a further increase in WAN traffic is desktop virtualization. To put the challenge of desktop virtualization into perspective, The Interop Respondents were asked about their organization's current and planned deployment of desktop virtualization. Their responses are shown in Table 5.2.

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

Table 5.2: The Percentage of Desktops that Already Have or Will be Virtualized

The data in Table 5.2 indicates that within the next year:

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

As pointed out in ***Virtualization: Benefits, Challenges and Solutions***, desktop virtualization results in a number of new protocols, such as Teradici's PC-over-IP (PCoIP), transiting the WAN. As that report details, in many instances these new protocols consume considerable WAN bandwidth and are latency sensitive. The performance challenges associated with desktop virtualization are described in detail in a subsequent section of this report.

As previously noted, given that storage tends to follow Moore's Law, an IT organization may well be able to support a significant increase in storage requirements without a dramatic increase in cost. Unfortunately, WAN services do not follow Moore's Law.

> ***The price/performance of MPLS tends to improve by only a couple of percentage points per year.***

As such, IT organizations will not be able to support a significant increase in bandwidth requirements without a significant increase in cost. To put the challenge of supporting significant increases in WAN bandwidth into context, The Webtorials Respondents were asked how their budget this year for all WAN services compares to what it was last year. Their responses are contained in Figure 5.1

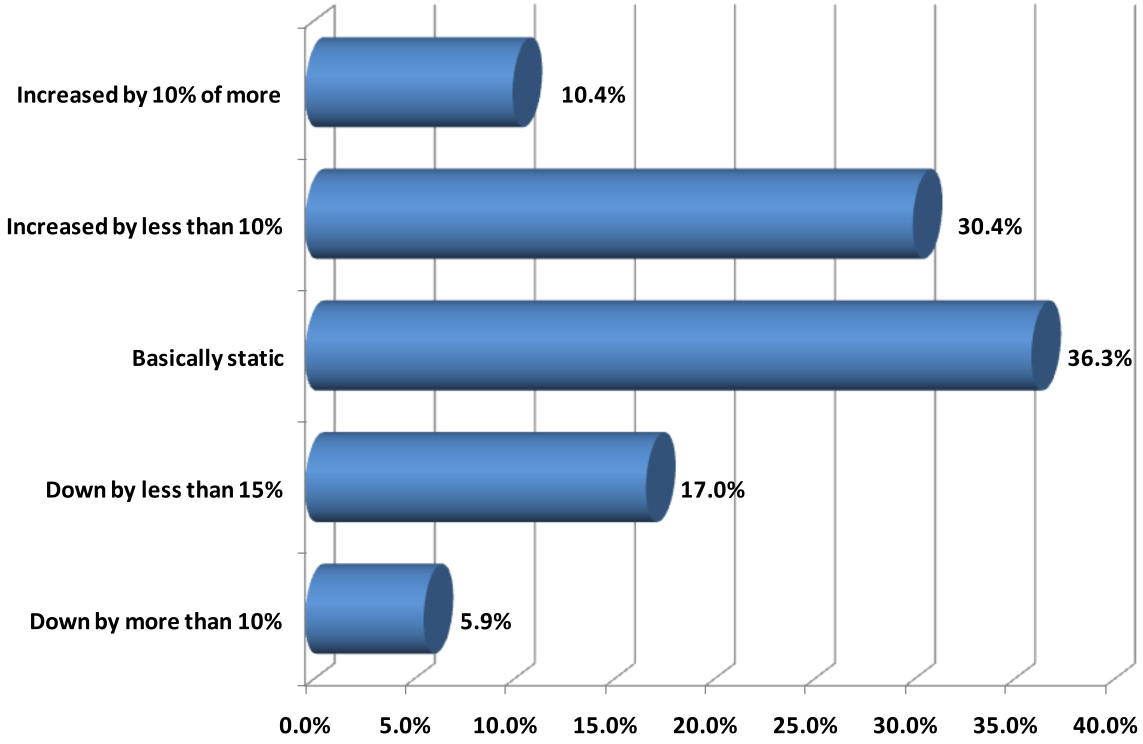

Figure 5.1: Anticipated Change in WAN Budgets

The data in Figure 5.1 shows that while over forty percent of IT organizations are experiencing an increase in their WAN budget, few are experiencing a significant increase. In addition, almost a quarter of IT organizations are experiencing a decrease in their WAN budget.

> ***IT organizations will not be able to support a significant increase in the use of WAN services with constrained budgets unless they make changes to how they use WAN services.***

Cloud Networking Without the Internet

There is a temptation to associate the WAN component of *cloud networking* exclusively with the Internet. For example, some definitions of cloud computing[3] assume that cloud services are always delivered over the Internet. Due to a variety of well-known issues (e.g., packet loss at peering points, BGP's inability to choose the path with the lowest delay, the TCP Slow start algorithm) the Internet often exhibits performance problems. As such, the Internet is not always the most appropriate WAN service to use to access cloud computing solutions. To put the use of the Internet into context, The Webtorials Respondents were asked to indicate which WAN service they would most likely use when accessing public and private

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

cloud computing services over the next year. Their responses are shown in Table 5.3.

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

Table 5.3: WAN Services to Access Cloud Computing Services

The data in Table 5.3 indicates that IT organizations understand the limitations of the Internet relative to supporting cloud computing.

> ***In many cases, the WAN service that IT organizations plan to use to support cloud computing is not the Internet.***

WAN Design

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

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

One of the key WAN design challenges is that WAN traffic is typically comprised of widely varying traffic types. This includes traffic generated by enterprise applications that are business critical and delay-sensitive (e.g., SCM, ERP); highly-visible, delay-sensitive real-time applications (e.g., voice, video and telepresence); and the growing use of public cloud computing services (e.g., Salesforce.com, Amazon's EC2) and social networking sites; e.g., LinkedIn and Facebook.

As previously demonstrated, the majority of IT organizations utilize MPLS. 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 types of applications discussed in the preceding paragraph. 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 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 5.2.

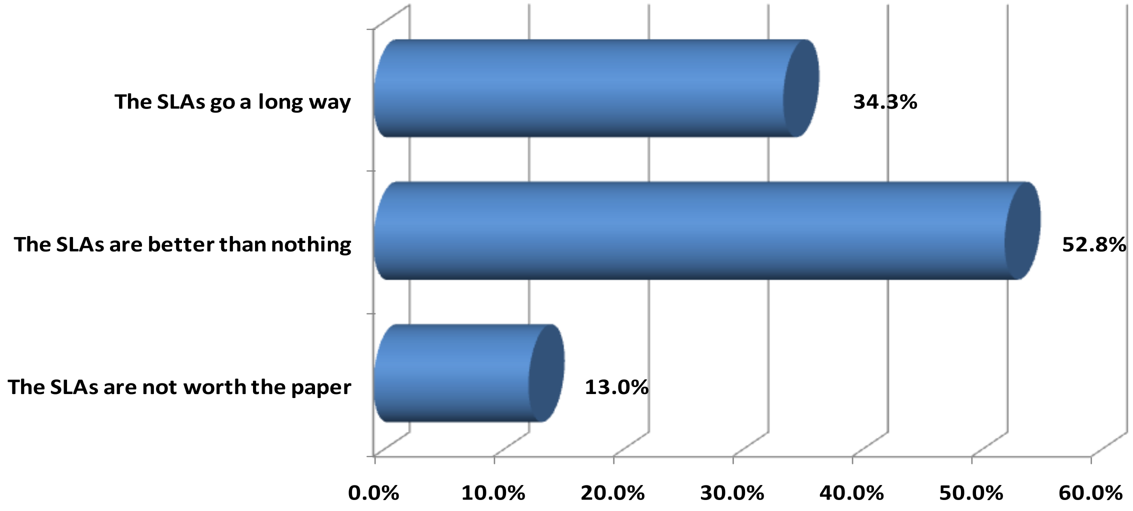

Figure 5.2: The Effectiveness of SLAs

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 NSPs as being effective.***

While the pricing and other characteristics of MPLS and Internet services vary on an international basis, Table 5.4 depicts how these services typically compare.

	MPLS	**Internet**
Cost	High	Low
Availability	High	High
Performance	Predictable	Non-Predictable
QoS	Sophisticated	None

Table 5.4: Characteristics of MPLS and Internet Services

One obvious conclusion that can be drawn from Table 5.4 is that IT organizations won't be able to satisfy the primary WAN design criteria with a simple network design that supports all traffic types. For example a network design that uses MPLS to interconnect an organization's large sites and uses a traditional approach to Internet access (e.g., a single DSL circuit) to connect their small sites would have two key limitations. This design would overpay for WAN services at the organization's large sites while not supporting any QoS functionality at their small sites.

WAN Service Alternatives

As noted, there is no new generation of 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.

VPLS

As previously mentioned:

> ***VPLS represents the combination of Ethernet and MPLS.***

While VPLS is not widely implemented today, the data in Table 5.1 indicates that roughly one quarter 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 in a single bridged (multipoint-to-multipoint) domain over a provider's IP/MPLS network, as shown in Figure 5.3. 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.

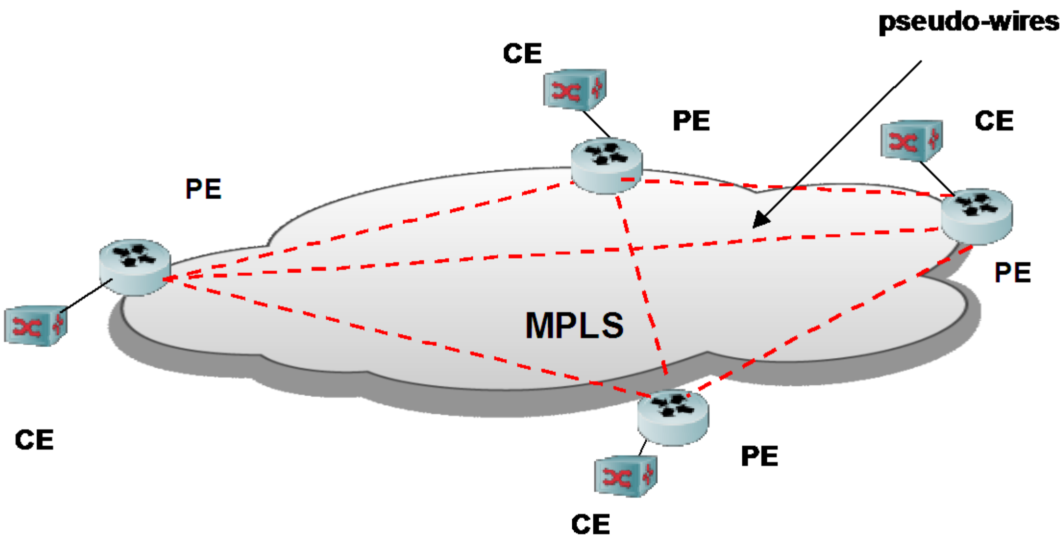

Figure 5.3: A VPLS Service Linking Four Customer Sites

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.

Within 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 the assigned inner labels, and associates these labels with the VPLS instance.

Table 5.5 provides a high level comparison of the different types of Ethernet WAN provider 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. Cisco's Overlay Transport Virtualization (OTV) falls into the latter category, as a proprietary option.

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

Table 5.5: Ethernet WAN Service Types

Hybrid WANs/WAN Virtualization

As previously noted, IT organizations won't be able to cost-effectively satisfy the primary WAN design criteria with a simplistic network design. As was also noted, in the current environment it is somewhat common to create a new WAN service by making minor variations to existing WAN services. One example of that approach is the utilization of Policy Based Routing (PBR) to implement a hybrid WAN that leverages multiple WAN services such as MPLS and the Internet. 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 table. Instead of routing by the destination address, PBR allows network administrators to create routing policies to allow or deny paths based on factors such as the identity of a particular end system, the protocol or the application.

An advantage of the PBR approach to creating a hybrid WAN is that IT organizations already have the functionality to do so. There are, however, some disadvantages of this approach. For example, configuring PBR is complex, time consuming and error prone. There is also not a direct linkage between PBR and other critical WAN functionality, such as the ability to optimize network and application traffic and the ability to have visibility into the traffic that transits the WAN.

Perhaps the biggest limitation of this simple 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 allocation of traffic causes all one type of traffic (e.g., SCM) to always transit an MPLS service while causing all traffic of another type (e.g., FTP) to always transit the Internet. There will be times, however, when sending some SCM traffic over the Internet will result in acceptable application performance and conserve expensive resources. Since PBR cannot respond in real time to changing network conditions, it will consume more costly bandwidth than would a dynamic approach to creating a hybrid WAN.

> ***In order to be effective, a hybrid WAN has to be able to perform adaptive path selection across two or more WAN links in a dynamic, intelligent fashion.***

In order to perform adaptive path selection a dynamic hybrid WAN must be able to select a WAN link in real time based on:

- The instantaneous end-to-end performance of each available network: This allows the solution to choose the optimal network path for differing traffic types.

- 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.

As noted, one option is for a hybrid WAN to balance traffic over MPLS and the Internet. While this reduces cost vs. an approach that puts all traffic over MPLS, additional cost savings are possible. For example, another option is to balance traffic over multiple low cost Internet access services such as DSL and cable. As previously noted, many IT organizations have avoided utilizing these access options for branch offices because the traditional approach to using DSL or cable for Internet access results in an unacceptably low level of availability and performance predictability.

However, because of adaptive path selection, the availability of a hybrid WAN that consists of multiple parallel paths is very high even if the availability of each component path is only moderately high. For example, Figure 5.4 depicts a system that is composed of two components that are connected in parallel.

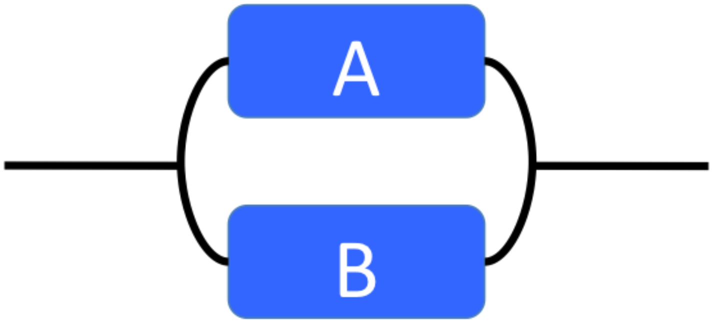

Figure 5.4: Two Components Connected in Parallel

The system depicted in Figure 5.4 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. This level of availability is equal to or exceeds the availability of most MPLS networks.

As described above, one of the principal advantages of a dynamic hybrid WAN is that is allows IT organizations to add significant amounts of additional bandwidth to an existing MPLS-based WAN at a relatively low incremental cost. WAN Virtualization (Figure 5.5) can be thought of as a variation of a hybrid WAN. In addition to enabling the augmentation of an MPLS WAN with inexpensive Internet connectivity, WAN Virtualization also gives IT organizations the option to reduce its monthly ongoing expense by either eliminating or reducing its MPLS connections while simultaneously providing more bandwidth than did the original design.

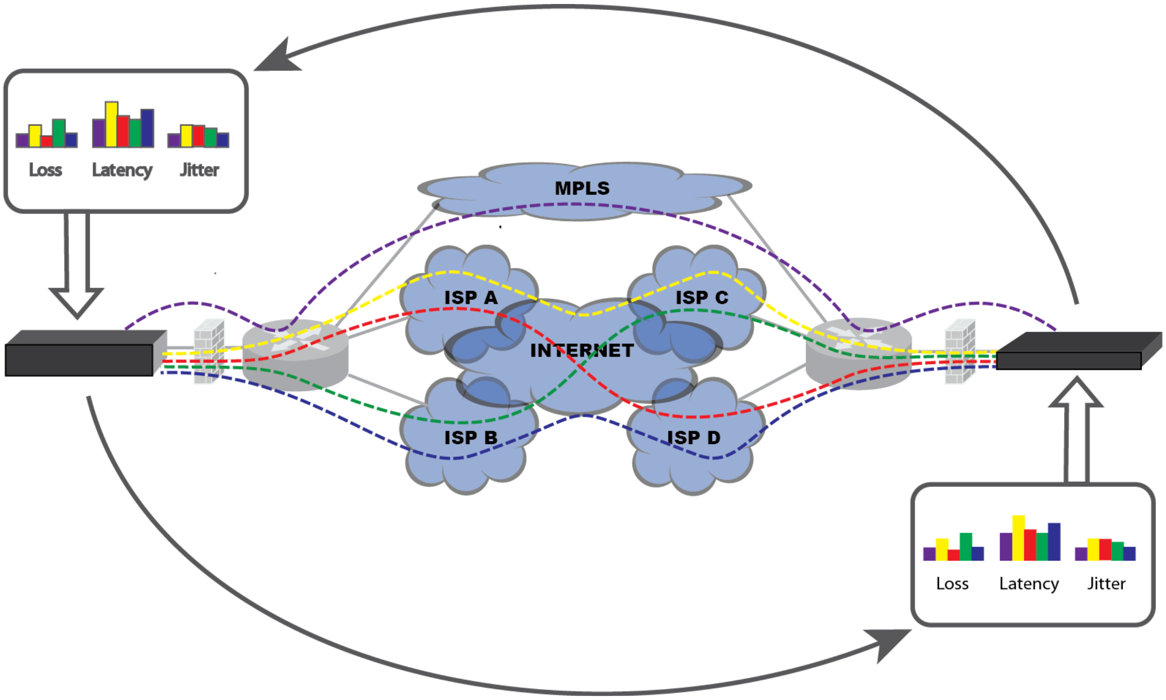

Figure 5.5: WAN Virtualization

As shown in Figure 5.5, because it continuously measures loss, latency, jitter and bandwidth utilization across all of the various paths between any 2 locations, in less than a second WAN Virtualization can switch traffic away from a path that is exhibiting an unacceptable level of performance. This capability, combined with the availability of parallel systems as depicted in Figure 5.4, 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 all of the time. This combination of capabilities also underscores the ability of WAN Virtualization to deliver performance predictability that equals, and in most cases exceeds, that of a single MPLS network.

Because of the high availability and performance predictability of WAN virtualization, 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.

Cloud Bridging

A hybrid cloud computing solution relies on a WAN to provide the connection between the enterprise locations, including the enterprise data center(s), and the public cloud data center(s) providing the IaaS or other cloud service. Ideally, the resulting hybrid cloud would appear to both users and administrators as a single cloud data center, with the physical location of application resources as transparent as possible.

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.***

In a hybrid cloud that leverages a WAN for access to the public portion of the cloud, transparency of application location has a number of implications:

- **VLAN Extension**
  The VLANs within which VMs are migrated must be extended over the WAN between 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 allows these application services to be provisioned from the private enterprise data center and eliminates the need for 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 user location within the enterprise network and the server location within the hybrid cloud. This means that public cloud data center extensions need to offer the same WAN optimization and application acceleration capabilities as are described below. In addition, WAN optimization controllers serving the bridged connection between the enterprise private cloud data center and the public cloud data center can accelerate VM migration, system backups, and other bulk data transfers between these data centers.

The Evolving Branch Office

As recently as 2005 the typical branch office was IT heavy. By that is meant that the branch office circa 2005 contained a lot of network equipment (e.g., switches and branch office routers) as well as a multitude of servers, applications, and storage. The devices in the typical branch office, however, were not aware of application level traffic detail. In addition, the network optimization of application traffic, if it existed at all, was very primitive (e.g., QoS, TOS) and it was focused on a couple of obvious applications. It was also unusual in this time frame for IT organizations to have visibility into the application traffic that was either originated in, or was destined to branch office users.

The current branch office infrastructure is less IT centric than it was in 2005. For example, the majority of IT organizations have removed at least some servers, applications and storage out of the branch office and placed these resources in centralized data centers. This centralization of IT resources is driven both by the desire to reduce computing and IT management costs as well as the desire to get more control over those resources and to implement better security.

In order to improve the performance of applications delivered to branch office employees, many IT organizations have deployed into their branch offices the WAN optimization controllers (WOCs) that are described below. In addition, in order to improve security, many IT organizations have deployed firewalls, intrusion detection systems (IDS) and intrusion protection systems (IPS) into their branch offices.

One of the design challenges facing IT organizations is how to strike a balance between having branch offices be IT heavy and consolidating all IT functionality out of branch offices and into centralized data centers.

To help IT organizations strike that balance, many vendors have developed a branch office box (BOB). The advantage of the BOB is that it can potentially consolidate several functions (i.e., WOC, firewall, IDS, IPS, print server, file server, domain controller, etc.) into a single physical device, while also allowing all of these capabilities to be managed via a single system management interface.

The Webtorials Respondents were asked to indicate the type of device that they would most likely build their branch office infrastructure around. Their responses are shown in Table 5.6.

Type of Device	Percentage of Respondents
A router that supports VMs	30.7%
An appliance, such as a WOC, that supports VMs	25.7%
A virtualized server	27.7%
Other	15.8%

Table 5.6: BOB Form Factor

In the *other* category, the most common response was that based on the situation, all three types of devices are provided.

Local Access to the Internet

The traditional approach to providing Internet access to branch office employees was to carry the 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 the Internet traffic and 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 WAN, which adds to the cost of the WAN. Another disadvantage of this approach is that it adds additional delay to the Internet traffic.

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 5.7.

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%

Table 5.7: Routing of Internet Traffic

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

> ***IT organizations will make an increased use of distributed access to the Internet from their branch offices.***

Network and Application Optimization

The Webtorials Respondents were asked to indicate how important it is to their organization to get better at seventeen different network and application optimization tasks over the next year. They were given the following five-point scale:

1. Not at all important
2. Slightly important
3. Moderately important
4. Very important
5. Extremely important

Table 5.8 shows the ten optimization tasks that are the most important for IT organizations to improve on in the next year. Included in Table 5.8 are the tasks and the percentage of The Webtorials Survey Respondents who indicated that the task was either very or extremely important for their organization to get better at over the next year.

Optimization Tasks	Importance: Very or Extremely
Relating the performance of applications to their impact on the business	70%
Ensuring acceptable performance for VoIP traffic	68%
Improving the performance of applications used by mobile workers	60%
Ensuring acceptable performance for video or telepresence traffic	57%
Ensuring acceptable performance for the applications that you acquire from a Software-as-a-Service (SaaS) provider	56%
Optimizing the performance of TCP	54%
Controlling the cost of the WAN by reducing the amount of traffic that transits the WAN	50%
Optimizing the Web tier of a multi-tiered application for peak utilization	50%
Optimizing the performance of specific applications such as SharePoint	49%
Optimizing the performance of protocols other than TCP; e.g., HTTP and MAPI	49%

Table 5.8: The Importance of Improving Optimization Tasks

The data in Table 5.8 shows that even though acquiring applications from a SaaS provider is a relatively recent phenomenon, more than half of The Webtorials Respondents stated that ensuring acceptable performance for the applications that they acquire from a SaaS provider is either very or extremely important. More detail on the optimization challenges facing IT organizations can be found in the report ***Application Delivery: A Reality Check***.

WAN Optimization Controllers (WOCs)

Goals of a WOC

The goal of a WOC is to improve the performance of applications delivered across the WAN from the data center either to the branch office or directly to the end user, typically over a network such as MPLS. A model is helpful to illustrate how the performance of a WAN can impact the performance of an application and also serves to illustrate how a WOC can improve application performance. The following model (Figure 5.6) is a variation of the application response time model created by Sevcik and Wetzel[4]. Like all mathematical models, the following is only an approximation, and as a result it is not intended to provide results that are accurate to the millisecond level.

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

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).

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

Figure 5.6: Application Response Time Model

In order to improve the performance of applications that are delivered over the WAN, WOCs implement a variety of techniques. These techniques are summarized in Table 5.9 and are explained in detail in ***The 2010 Application Delivery Handbook***.

WAN Characteristics	WAN Optimization Techniques
Insufficient Bandwidth	Data Reduction: • Data Compression • Differencing (a.k.a., de-duplication) • Caching
High Latency	Protocol Acceleration: • TCP • HTTP • CIFS • ICA • RDP Application Acceleration • SharePoint • SAP • Oracle Mitigate Round-trip Time • Request Prediction • Response Spoofing
Packet Loss	Congestion Control Forward Error Correction (FEC) Packet Reordering
Network Contention	Quality of Service (QoS)

Table 5.9: Techniques to Improve Application Performance

Virtually all WAN Optimization Controllers (WOCs) on the market support the functions listed above. However, as described below, there are some significant differences in terms of how the functionality is implemented and how well it performs. In addition, some WOC vendors provide functionality not included in the above list. A recent report[5] provides insight into the primary WOC vendors and their products.

Enabling Virtual Desktops

As was previously mentioned, one of the factors that will drive more traffic over the WAN is the implementation of virtual desktops. As explained in the report entitled ***Virtualization: Benefits, Challenges and Solutions***, the two fundamental forms of desktop virtualization are client side (a.k.a., streamed desktops) and server side; a.k.a., hosted desktops.

> ***The ICA and RDP protocols employed by many hosted desktop virtualization solutions are examples of protocols that can be difficult to optimize.***

One of the reasons that these protocols can be difficult to optimize is that they only send small request-reply packets. Byte-level caching best optimizes this form of communications. Unfortunately, not all WOCs support byte-level caching. Implementers of desktop virtualization need to understand the functionality provided by the various WOCs and to evaluate that functionality in the context of the types of desktop virtualization that they want to deploy.

As shown in Table 5.10, techniques such as byte level compression, caching, protocol (e.g., ICA, RDP) optimization, and QoS can provide benefits for hosted desktops. Before implementing them, however, an IT organization must determine which acceleration techniques are compatible with the relevant display protocols. For example, in order to be able to compress ICA traffic, a WOC must be able to decrypt the ICA workload, apply the optimization technique, and then re-encrypt the data stream.

	Streamed Desktops	**Hosted Desktops**
Block Level Compression	**X**	
Byte Level Compression	**X**	**X**
Caching	**X**	**X**
Staging	**X**	
Protocol Optimization (e.g., TCP, IP, UDP)	**X**	**X**
Protocol Optimization (e.g., ICA, RDP)		**X**
Protocol Optimization (e.g., CIFS, HTTP, MAPI)	**X**	
QoS	**X**	**X**

Table 5.10: Applicability of Common WAN Optimization Techniques

[5] http://searchenterprisewan.techtarget.com/generic/0,295582,sid200_gci1381156,00.html

The support of streamed desktops also creates some significant WAN performance problems that may require the deployment of a WOC. For example, 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. Hence, in order to effectively support application streaming, IT organizations need to be able to optimize the performance of protocols such as CIFS, MAPI, HTTP, and TCP. In addition, IT organizations need to implement other techniques that reduce the bandwidth requirements of application streaming. For example, by using a WOC it is possible to cache the virtual application code at the client's site. Caching greatly reduces the volume of traffic for client-side virtualized applications and it also allows applications to be run locally in the event of network outages. Staging is a technique that is similar to caching but is based on pre-positioning and storing streamed applications at the branch office on the WOC or on a branch server. With staging, the application is already locally available at the branch when users arrive for work and begin to access their virtualized applications.

Whether it is done by the WOC itself, or in conjunction with the WOC, supporting desktop virtualization will require that IT organizations are able to apply the right mix of optimization technologies for each situation. For example, protocols such as ICA and RDP already incorporate a number of compression techniques. As a result, any compression performed by a WAN optimization appliance must adaptively orchestrate with the hosted virtualization infrastructure to prevent compressing the traffic twice - a condition that can actually increase the size of the compressed payload.

Mobile Workers

In the last few years there has been an explosive growth in the number of mobile workers.

In order to enable this growing population of mobile workers to access enterprise applications as easily as do workers in branch offices, the communications between the mobile worker and the data center (whether it is owned by the enterprise or a third party provider such as a cloud computing service provider) has to be optimized. The importance of optimizing this traffic is reflected in the data in Table 5.8. As shown in that table, sixty percent of The Webtorials Respondents stated that improving the performance of applications used by mobile workers is either very or extremely important.

One way to optimize this communications is to deploy client software that provides WOC functionality onto the user's mobile device. In many cases the mobile worker will use some form of wireless access. Since wireless access tends to exhibit more packet loss than does wired access:

> ***The WOC software that gets deployed to support mobile workers needs functionality such as forward error correction that can overcome the impact of packet loss.***

In addition, as workers move in and out of a branch office, it will be necessary for a seamless handoff between the mobile client and the branch office WOC.

Until recently, the typical device that mobile workers used to access enterprise applications was a laptop. While that is still the most common scenario, today many mobile workers use their smartphones to access enterprise applications. Therefore, over the next few years it is reasonable to expect that many IT organizations will support the use of smartphones as an access device by implementing server-side application virtualization for those devices. This means that in a manner somewhat similar to remote workers, mobile workers will access corporate applications by running protocols such as ICA, RDP and PCoIP over a WAN.

Many IT organizations, however, resist putting any more software on the user's device. In addition, many users resent having multiple clients (e.g., WOC, SSL VPN, IPSec VPN, wireless/cellular access) that are not integrated on their access device. On a going forward basis, IT organizations should look to implement WOC software that is integrated with the other clients used by mobile workers.

Application Delivery Controllers (ADCs)

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 an 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 5.6, 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 deployment of an SLB enables an IT organization to get a *linear benefit* out of its servers. That means that if an IT organization that has implemented an SLB doubles the number of servers supported by that SLB that it should be able to roughly double the number of transactions that it supports. The traffic at most Web sites, however, is not growing at a linear rate, but at an exponential rate. To exemplify the type of problem this creates, assume that the traffic at a hypothetical company's (Acme) Web site doubles every year. If Acme's IT organization has deployed a linear solution, such as an SLB, after three years it will have to deploy eight times as many servers as it originally had in order to support the increased traffic. However, if Acme's IT organization were to deploy an effective ADC then after three years it would still have to increase the number of servers it supports, but only by a factor of two or three – not a factor of eight. The phrase *effective ADC* refers to the ability of an ADC to have all features turned on and still support the peak traffic load.

Key ADC Functionality

Below is a listing of the type of functionality that is provided by an ADC. This functionality is explained in detail in ***The 2010 Application Delivery Handbook***.

- **TCP Offload**
  TCP offload functionality is designed to deal with the complexity associated with the fact that each object on a Web page requires its own short-lived TCP connection. Processing all of these connections can consume an inordinate about of the server's CPU resources. Acting as a proxy, the ADC terminates the client-side TCP sessions and multiplexes numerous short-lived network sessions initiated as client-side object requests into a single longer-lived session between the ADC and the Web servers.

- **SLB and Global SLB**
  As noted, an ADC sits in front of a server farm and receives service requests from clients and delivers the requests for service to the most appropriate servers. As such, an ADC functions as a traditional SLB. In addition, an ADC can function as a global server load balancer (GSLB). In this role the ADC balances the load across geographically dispersed data centers by sending a service request to the data center that is exhibiting the best performance metrics.

- **SSL Offload**
  The ADC terminates the SSL session by assuming the role of an SSL Proxy for the servers. SSL offload can provide a significant increase in the performance of secure intranet or Internet Web sites. SSL offload frees up server resources, allowing existing servers to process more requests for content and handle more transactions.

- **XML Offload**
  XML is a verbose protocol that is CPU-intensive. Hence, another function that can be provided by the ADC is to offload XML processing from the servers by having an ADC serve as an XML gateway.

- **Scripting**
  One of the characteristics of most IT environments is that the environment is comprised of a large and growing number of servers and applications. Another characteristic is that most IT organizations have very limited control as to which users access which applications and which servers.   An ADC gives control to the IT organization through functionality sometimes referred to as scripting, and sometimes referred to as a rules engine. This functionality allows the IT organization to directly classify and modify the traffic of any IP-based application.

- **Application Firewalls**
  ADCs may also provide an additional layer of security for Web applications by incorporating application firewall functionality. Application Firewalls are focused on blocking increasingly prevalent application-level attacks. Application firewalls are typically based on Deep Packet Inspection (DPI), coupled with session awareness and behavioral models of normal application interchange.

Virtual Appliances

Section 4 of this report 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 two fundamentally different ways. For example, two or more appliances, such as ADCs, can be combined to appear as a single logical ADC. However, the more 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 among others. As previously discussed, in the branch office a suitably placed virtualized server could potentially host a virtual WOC appliance as well as other virtual appliances. Alternatively, a router or a WOC that supports VMs could also serve as the infrastructure foundation of the branch office.

> ***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 cost of a software-based appliance can be a third less than the cost of a hardware-based appliance[6]. 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[7].

In addition to cost savings, another advantage of a virtual appliance is that it offers the potential to alleviate some of the management burdens in branch offices because most of the provisioning, software updates, configuration, and other management tasks can be automated and centralized at the data center.

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

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

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

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.

A virtualized ADC 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 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 possibility because any actions taken by the application group relative to the ADC will only impact their application.

One approach to monitoring and troubleshooting inter-VM traffic is to deploy a virtual performance management appliance or probe (vProbe).

One of the characteristics of a virtualized server is that each virtual machine only has at its disposal a fraction of the resources (i.e., CPU, memory, storage) of the physical server on which it resides. As a result, in order to be effective, a vProbe must not consume significant resources. The way that a vProbe works is similar to how many IT organizations monitor a physical switch. In particular, the vSwitch has one of its ports provisioned to be in promiscuous mode and hence forwards all inter-VM traffic to the vProbe. As a result, the use of a vProbe gives the IT organization the necessary visibility into the inter-VM traffic.

A virtual firewall appliance can 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 well as between VMs running on the same physical server. The security appliance can potentially provide highly integrated functionality to help secure virtual machines, applications, and traffic. This includes firewall, VPN, anti-malware, IDS/IPS, integrity monitoring (e.g., registry changes), and log inspection functionality.

Virtualized security management makes it is possible to meet critical regulatory compliance requirements for full application segregation and protection within the confines of virtualized physical servers. Through tight integration with the virtual server management system, 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 recently published 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.

Software as a Service (SaaS) Based Solutions

As previously noted, one of the primary forms of public cloud computing is SaaS. As was also noted, the phrase ***cloud networking*** refers to the LAN, WAN and management functionality that must be in place to support cloud computing. As described below, there is a distinct synergy in which one of the approaches that typifies public cloud computing, SaaS, can be used as part of a cloud network to better support cloud computing. In particular, IT organizations can acquire the optimization and security functionality they need to support cloud computing from a SaaS provider. Currently, the most developed form of a SaaS solution that provides optimization and security functionality is an Internet overlay.

An Internet Overlay

As previously described, IT organizations often implement WAN Optimization Controllers (WOCs) and Application Delivery Controllers (ADCs) 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 does not 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 5.7, 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) close to application users, typically within a single network hop. This edge server then optimizes the traffic flow to the server closest to the data center's origin server.

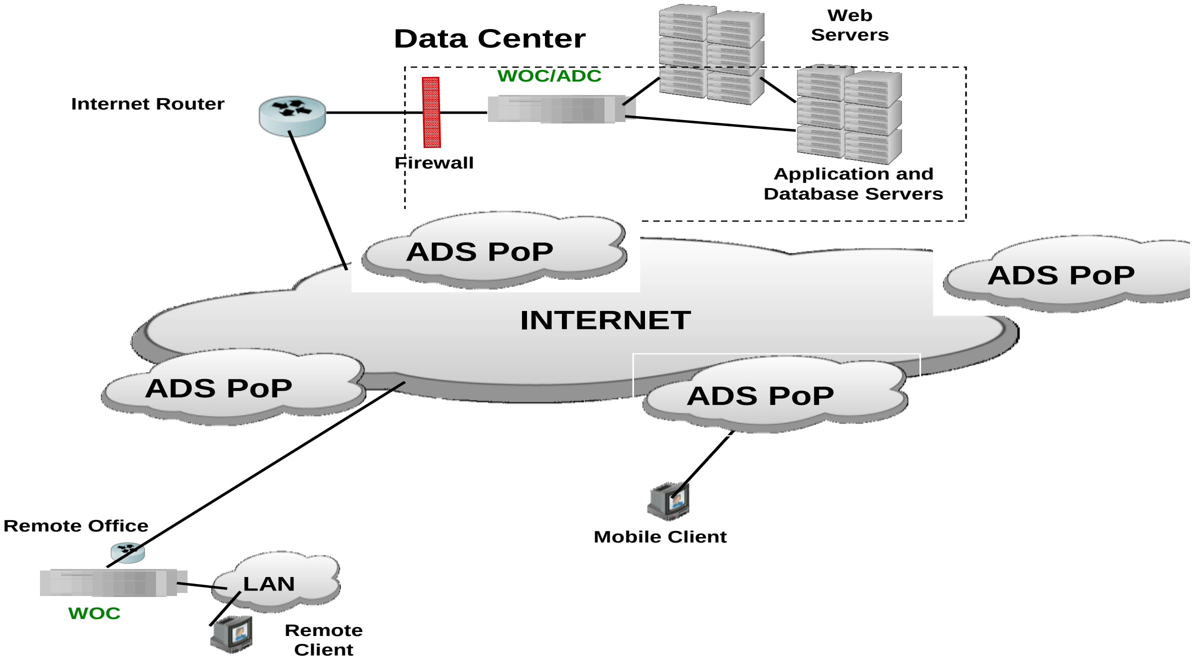

Figure 5.7: An Internet Overlay

An Internet overlay provides a variety of optimization functions that generally complement solutions such as an ADC rather than overlap or compete with them. One such function 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. 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 common functionality associated with an Internet overlay include:

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

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

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.Webtorias.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.

Virtualization: A Key Enabler for Effective Cloud Implementations

Why Use Virtualization?

As Cloud Computing adoption increases, virtualization is a key enabler, driving economies of scale and the ability to scale with hardware appliances or commodity hardware.

Virtualization solutions allow:

- Delivery of elastic, flexible and scalable solutions for changing-traffic volumes
- Enablement of a cost effective on-demand approach to reduce capital expenditure
- Efficiency for Public or Private Clouds

A10 offers a wide range of options, as one solution does not fit all requirements. Beyond the hype of Cloud generalizations is the reality of making the solution work for your unique needs. While Cloud providers take the burden off internal IT organizations, the risks of not considering the hardware used and potential issues of the wrong solution are apparent.

Organizations may no longer require owning the hardware in Cloud implementations, but they will still use similar devices to handle traffic. The advent of hypervisor solutions, or "virtual appliances" are serious options that offer an alternative to fixed hardware appliances, but each solution has its own pros and cons that must be considered, both from the feature and performance angles.

This is the reason A10 Networks' AX Series offers many solutions, from the flexible SoftAX to high performance hypervisor free AX Virtualization.

AX Series Virtualization Products & Solutions

Based on A10's award-winning AX Series Application Delivery Controllers (ADC) and Advanced Core Operating System (ACOS) architecture, enterprises and service providers have the flexibility to choose the following scale-as-you-grow virtualization options.

SoftAX

- SoftADC: AX virtual machine (VM) atop a hypervisor on commodity hardware
- Rapidly scale with commodity hardware
- Reduce hardware costs and upload to compatible cloud providers
- Flexible solution leveraging an existing Cloud provider or internal virtualized infrastructure

AX Virtual Chassis System (aVCS)

- Cost effective alternative to fixed ADC pairs and fixed chassis systems
- Massively increase performance to hundreds of Gbps and multiple millions of L4 connections per second
- Cluster multiple AX devices to operate as a unified single device
- Scale multiple AX devices with shared capacity, High Availability (HA) and single IP management
- Reduce cost and simplify management while adding devices as you grow

AX Virtualization

- High performance multi-tenancy without hypervisor cost and hypervisor performance hit
- Application Delivery Partitions (ADPs) divide the AX platform resources for individual applications
- Enables quality multi-tenancy with granular resource allocation
- Reduce the number of appliances to host multiple applications

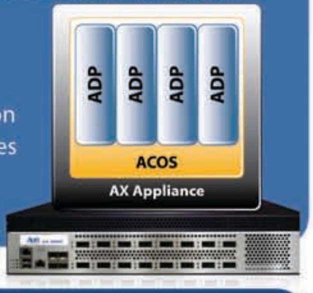

AX-V Appliance

- The first dedicated hardware platform designed specifically for hypervisor based ADCs
- Multiple SoftADCs: AX virtual machines (VMs) on dedicated AX Series hardware
- SoftAX flexibility with AX hardware performance and reliability
- Guaranteed performance, certifications, support and optimized hardware

The AX Series virtualization products and features are in addition to existing integration with leading third party virtualization vendors, such as VMware and associated solutions for vSphere acceleration, vCenter dynamic provisioning and VMotion with Global Server Load Balancing (GSLB).

vmware PARTNER
TECHNOLOGY ALLIANCE

Virtualization at Work: Subaru Canada and A10 Case Study

Subaru Canada had been using the Foundry ServerIron 4G-SSL to provide server load balancing for its website (www.subaru.ca). However, when it came time to renew the support contract with Brocade Communications Systems, Inc., which had acquired Foundry Networks in 2008, Subaru Canada decided to evaluate some of the newer technologies available.

Subaru Canada's Director of eBusiness & Information Systems, George Hamin, became impressed with A10 Networks' AX Series New Generation Server Load Balancers while running a proof of concept using the AX 1000.

While Hamin and his team were impressed with the performance of the AX 1000, due to the rapid growth rate of sales at Subaru Canada, they decided they might later appreciate having the additional overhead provided by the AX 2500, with its 10 Gbps throughput capacity, as opposed to the 4 Gbps capacity of the AX 1000. With a list price of $2,500 per Gbps, the AX 2500 was a bargain, costing less than one-third of competing solutions (based on throughput-$-per-Gbps metric). Hamin said it was an easy choice, since the AX appliance cost "just a little more than the cost of renewing support on our 4G-SSL."

Hamin was originally interested in the AX's Application Acceleration features. The AX Series is optimized for SSL and L4-7 acceleration, and web caching further accelerates the user experience by reducing the time required to download each page. This, in turn, reduces the amount of bandwidth needed to serve pages and decreases the total number of requests placed to web servers. Furthermore, the AX Series offers several compression algorithms to reduce the size of each object on the page. Again, this helps reduce the amount of bandwidth being used. Hamin said he was able to leverage the compression and caching features in order to greatly accelerate the delivery of the enterprise's web content.

It was only after Subaru Canada had installed the 64-bit AX 2500 appliances that Hamin and his team learned of the additional AX virtualization feature. They were intrigued by the possibility that this feature might help them reduce the costs associated with supporting both mail and web applications. Virtualization allows customers to sub-divide an AX internally for multi-tenant purposes, whether for multiple organizations, departments, or simply, as in Subaru's case, multiple disparate applications. Each segmented area becomes an Application Delivery Partition (ADP). Within ADPs, various resources and elements are available. Layer 2/3 virtualization on a per-ADP basis was a particularly interesting enhancement to the ADP feature, as this guarantees true network segmentation between Subaru's applications.

SUBARU.

> Subaru Canada, Inc. markets and distributes Subaru vehicles, parts and accessories through a network of over 86 authorized dealers across Canada. This past March was their website's busiest ever, with 306,000 visitors viewing 1.97 million web pages.
>
> "Once a potential buyer test drives one of our vehicles, the rest is easy. I feel the same way about A10's AX Series of appliances - once you try them you'll be sold… While we were originally drawn to the AX's application acceleration features, the recent enhancements to the AX Virtualization Multi-tenancy feature will allow us to consolidate our Microsoft Exchange 2010 environment and our web environment to a single pair of appliances, with high availability. This reduces the amount of Application Delivery Controllers in our network and saves us money in the process."
>
> **George Hamin**
> *Director eBusiness & Information Systems for Subaru Canada, Inc*

The AX Virtualization Multi-tenancy feature will allow Hamin to consolidate his distinct environments as if the ADCs were different platforms (i.e., a Microsoft Exchange Server 2010 environment and a web environment) onto a single pair of AX appliances. The pair of AX appliances will be set up in High Availability (HA) mode to mirror the content on the primary appliance and to act as a failover. This implementation will enable Subaru to reduce the total number of ADCs in the network, saving the company a large amount of money in the process.

"So rather than buying a pair of AX 2500s for HA web, another pair for HA Exchange, and another pair for HA SharePoint, you can virtualize a single pair and just keep throwing applications at it until you hit the limits imposed by your applications' collective peak load conditions, CPU, RAM, or ports," Hamin said.

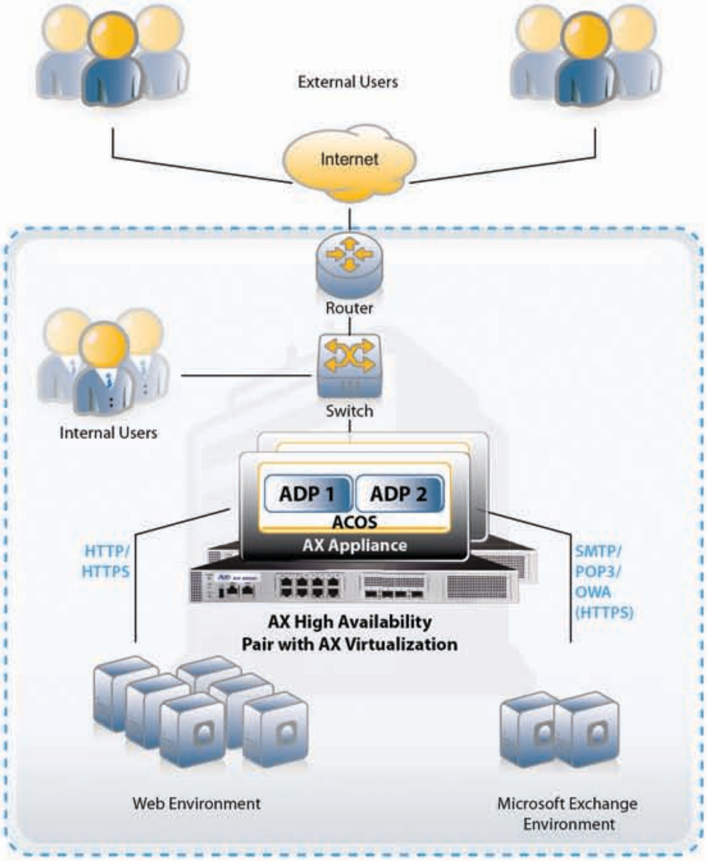

Summary

A10 Networks offers innovative virtualization solutions to enable any Public or Private Cloud deployment. With the widest range of solutions organizations can ensure they receive the right solution for their business and customers.

Please contact A10 for a free consultation of which solution would work best for your organization or to arrange a demonstration or trial at inquire@a10networks.com or www.a10networks.com

About A10 Networks

A10 Networks was founded in 2004 with a mission to provide innovative networking and security solutions. A10 Networks makes high-performance products that help organizations accelerate, optimize and secure their applications. A10 Networks is headquartered in Silicon Valley with offices in the United States, United Kingdom, France, The Netherlands, Germany, Brazil, Japan, China, Korea and Taiwan. For more information, visit: www.a10networks.com

About AX Series

A10 Networks' AX Series is the industry's best price/performance advanced traffic manager – helping enterprises and ISPs maximize application availability through a high-performance and scalable web Application Delivery platform. The AX's Advanced Core Operating System (ACOS) architecture has garnered the company numerous awards and is revolutionary by market standards due to its scalable symmetrical multiprocessing (SSMP), shared memory architecture. AX includes an optimized multi-CPU architecture built from the ground up that leaps the competition in terms of performance, scalability and reliability. For more information, visit: www.a10networks.com/products/axseries

Software WAN Optimization

Accelerate Your Business

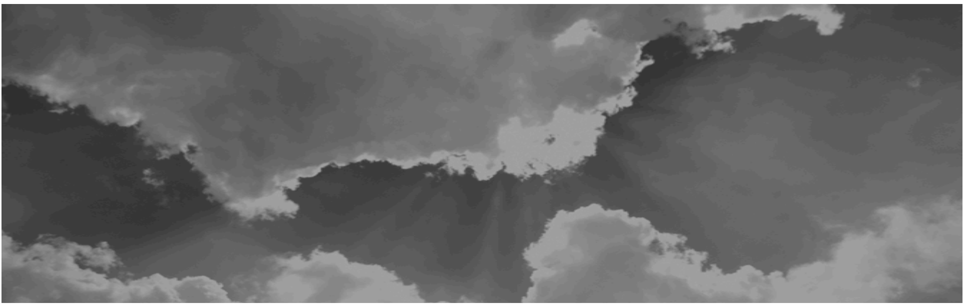

certeon®

Transform computing, storage and networking resources *into* an integrated, agile and scalable cloud infrastructure *with* aCelera WAN Optimization software

certeon
THE Application Performance Company

".. application performance .. one of the top three inhibitors of cloud adoption"

Clouds and Beyond: Positioning for the Next 20 Years of Enterprise IT, Frank Gens, IDC

"Deploying virtual WAN optimization software has been as simple and inexpensive as remotely connecting to the server over the WAN"

"Virtual WAN Optimization software gives much more flexibility, which is imperative"
Ernest Ostro: Director of Information Services, Pathfinder International

Certeon Inc.
4 Van de Graaff Drive
Burlington, MA 01803
781 425 5200
http://www.certeon.com

Cloud Promise and Challenge

Cloud services look like a $100 billion-plus opportunity by mid decade, but is cloud computing worth this level of excitement? **Think, Internet 1997**. Companies were excited about the technology potential and worried about *security, privacy, bandwidth*, standards and more. In spite of these questions, what transformed communication and commerce? The ability to deliver **business value!**

In 2010 and beyond Cloud successes will be measured in **business value**. The units of measure will be the ability to increase business agility, decrease cost through on-demand provisioning and teardown of infrastructure and services, speed development, and improved reliability. It must be utility-based, self-service, secure and most importantly, have levels of application performance that improve productivity. User adoption is the linchpin of any business value equation.

Leveraging cloud computing and maximizing its value business value requires full featured, secure, scalable, high performance WAN Optimization software that allows applications to perform as expected, and can be part of any on demand architecture, rather than part of a farm of tactical hardware or limited virtual appliance solutions.

Cloud success requires integrating network services that are very far away and often owned by strangers

Business information and resources are increasingly being accessed at global scale distances, from enterprise and cloud sources using Internet, VPN or MPLS connections. At the same time, expectations for application performance are rising.

Enterprises embracing the cost and scalability benefits of cloud computing and service providers delivering consumption and utility-based models, balance the need for security and user expectations for access and application performance. Users don't care if the resource is in a cloud or on the moon, they expect their applications to work quickly and flawlessly.

Bottom line: the success of cloud computing is irreversibly linked to software based WAN Optimization and Application Acceleration technologies as the result of distance induced latency and the need to provide ad-hoc secure and multi-tenant access. aCelera software WAN Optimization's ability to provide secure access, application performance and global scale make it the ideal cornerstone of cloud environments, from Private to Public to Hybrid.

Certeon

Certeon is the leading supplier of 21st century WAN optimization software for agile, elastic, and multi-tenant deployment. Certeon aCelera solves application performance challenges for cloud-based networks as effectively as it does for corporate networks. aCelera software and virtual appliances enable automated, secure and optimized performance for any application, on any device, across any network reducing response time by up to 95% while reducing the bandwidth used from 65 to 95 percent.

aCelera's creates global web of data that will enable businesses to leverage corporate and cloud provider networks to create new services or revenue streams. Certeon aCelera enables cloud service providers to offer on demand WAN Optimization to their catalogs as a one click value-added service.

Enterprise heterogeneous and decentralized needs

Enterprises today are a heterogeneous mix of hardware and virtualization platforms, custom and off the shelf applications, storage technologies, networking equipment and service providers all strung together in a web around the globe.

Decentralization of information sources, delivery workloads and productive users takes this heterogeneous infrastructure and explodes it's management and access problems across the globe. Clouds, company datacenters, branch offices, home offices, coffee shops are all part of the new enterprise.

The effort to make this mix of services and technologies useful, affordable and valuable has service providers of all types rolling out a range of cloud service models (IaaS, SaaS, PaaS, "X"aaS) and an array of deployment models (private, public, and hybrid clouds), that promise provide flexibility, scalability, cost savings that will create competitive advantage. But, even these environments are a heterogeneous mix of virtualization technologies from 3 or 4 vendors.

The combination of a heterogeneous infrastructure and decentralized enterprise with cloud services demands that WAN Optimization solutions be built to support this heterogeneous flexible infrastructure; they must use and be managed with the same building blocks as the environments they support. WAN optimization cannot just be a halo product targeting, or moving to a solution to cloud problems from outside the stack.

aCelera WAN optimization software: built for the cloud, not just moving to the cloud

Solutions "moving to" clouds do not support the dynamic, global and heterogeneous nature of enterprise or **"X"aaS**ervice models. aCelera software and virtual appliances are "built" for the cloud and seamlessly integrate with all of these emerging technologies, delivering resources and services without compromising performance, scalability, or cost reduction.

The future of enterprise business success requires integrating network services that are very far away and often owned by strangers

"… 34% of survey respondents are using 2 virtualization solutions and 36% are using three or more."

"Users should plan for multiple virtualization platforms."

Fall 2010 ESG research study of 463 North American-based IT pros at organizations larger than 500 employees

"Productivity isn't everything, but in the long run it is almost everything." Paul Krugman

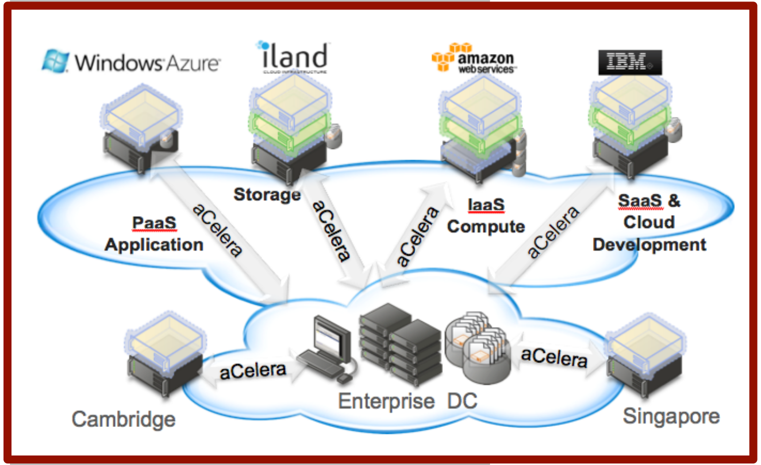

Enterprises and cloud service providers can deploy aCelera in any form factor, using any number of instances, delivering any throughput capacity, aligned with any application SLA requirement while meeting cost savings objectives and footprint limitations. This can be done in seconds on the enterprise premise, hosted, cloud sourced or in any mix of locations.

aCelera leverages enterprises' and service providers' growing heterogeneous virtualized infrastructures, in data centers, branches and on clouds. This allows organizations to turn clear TCO benefits into innovation. Saved acquisition, operations, real estate, power, cooling and maintenance/support costs create this opportunity where solutions not built for virtualized and cloud environments limit innovation.

Virtualization was just the first step

Virtualization is a driving IT strategy and initiatives from SMBs through large enterprises, up to the very large hosting companies, carrier data centers and cloud providers. "Server virtualization provides a foundation for IT automation, dynamic workload mobility, and finally, a bridge to cloud computing."[1]

Virtualization cannot be a single vendor strategy. ISVs creating virtual appliances that support a single hypervisor platform are "moving to the cloud" with products that don't match the requirement to support heterogeneous environments. Single platform virtualization creates castaway technology - islands of virtualization capabilities that are an extension of hardware appliance platform.

aCelera: built for heterogeneous, decentralized work

Certeon's aCelera software is built to provide ALL the performance advantages of any HARDWARE WAN Optimization product along with the flexibility, scalability, manageability and cost-savings of software and virtualization. aCelera supports In-line & out-of-line deployment with software and hardware failover and any level of SSL security.

aCelera can be deployed in any virtualized private, public, and hybrid cloud computing environments and is poised to meet ANY future performance and agency demand imposed by any enterprise's heterogeneous, decentralized and cloud environments.

aCelera software and virtual appliances deliver performance benefits and advantages without the downsides of hardware costs or the friction of limited scope virtualization. aCelera can easily be scaled on any existing hardware platform or migrated to more powerful platforms and processors when business conditions dictate, leveraging all the tools of any virtualization infrastructure.

aCelera software exceeds the scalability and performance of purpose-built hardware appliances. aCelera software is built to support global enterprise scalability requirements and is ready for the Internet scale usage demands of managed services and cloud computing.

aCelera software WAN optimization - 60% better 3 year TCO and 50% better connection scalability

aCelera™

Secure Automated Optimized

- Any form factor
- Any number of instances
- Any throughput capacity
- Any security requirement
- Any routing mode
- Any deployment model: enterprise, hosted, cloud sourced or combination

- Meet cost savings objectives
- Match footprint limitations

Microsoft GOLD CERTIFIED Partner

vmware READY

certeon
THE Application Performance Company

[1] 1, Enterprise Strategy Group

WAN Governance in a Cloud environment

Perform today, take control of tomorrow

Ipanema enables any large enterprise to have full control and optimization of their global networks; private Cloud, public Cloud or both. Moreover, Ipanema is the only system with a central management and reporting platform that scales to the levels required by service providers and large enterprises.

Leading the service providers market for application-centric network services, Ipanema has been proven in large enterprise global networks.

Enterprise infrastructure and WAN, are under constant transformation

Enterprises are on their way to the Cloud…

- They deploy private and hosted datacenters
- They use more and more SaaS applications (Salesforce, Googleapps…)
- Social media (LinkedIn, Twitter) and recreational applications (YouTube, Facebook) are popular
- Employees work not only from branch offices, also from home, hotels, airports…

… and yet to perform today they require:

- Guaranteed application performance
- Total business continuity
- Business process agility
- IT cost savings

WAN Governance aligns the network to IT priorities

WAN Governance is a unique Top-Down approach enabling enterprises to align their global network to IT and business priorities.

It fully controls and optimizes the global network, private Cloud, public Cloud or both. It guarantees that enterprises are always in control of critical applications. It unifies application performance across disparate networks. It dynamically adapts to whatever is happening in the network.

WAN Governance is the answer to all these challenges:

How to get full visibility of your global network:

- Discover which applications use your network resources
- Understand what is the root cause of slow applications
- Communicate clear data about application performance

How to deliver business applications:

- Guarantee voice, tele-presence and data applications over a converged network
- Ensure excellent application performance to your distributed workforce
- Manage social media and recreational applications

How to cost optimize your WAN:

- Reduce your WAN bandwidth requirements now and plan for tomorrow
- Use the Internet as a business network
- Get global control without deploying extra technology everywhere

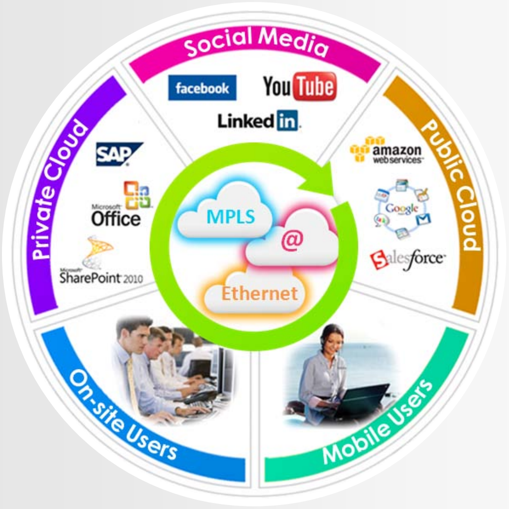

ANS^{TM}, the Autonomic Networking System is the way to deliver WAN Governance

The Ipanema Autonomic Networking System is unique in many aspects:

- Its **central management** based on application performance objectives provides unmatched operation simplicity and automation
- It tightly couples key features in an **All-in-One** approach to ensure the best possible user experience
- Based on a **fully automated** “sense-and-respond” architecture, it adapts to any traffic situation and any network topology
- Its **collaborative agents** deliver full control with physical deployment in only 10-20% of locations
- It **scales** up to 10M users, 100K sites and 10K networks and can match any enterprise and large Service Provider deployment

Key features for an All-in-One system

Application Visibility provides full transparency for application traffic using a true L7 deep packet inspection, topology and performance. Its unique end-to-end metrics (like one-way-delay) easily differentiates network and IT problems. Embedded data consolidation and reporting provides all needed reports from C-level KPIs to technical information for the helpdesk team.

QoS and Control dynamically allocates network resources and combines all type of traffic (voice and tele-presence, Citrix, file transfer, CIFS…) fluidly – based on user behavior, application technical requirements and business criticality. It automatically takes into account complex situations like some-to-many and any-to-any traffic mesh and Cloud-based application delivery over private and hosted datacenters as well as SaaS.

WAN Optimization accelerates application response time and reduces bandwidth requirements by using all up-to-date techniques like byte caching, CIFS acceleration, TCP acceleration, etc.

Dynamic WAN Selection (DWS)automatically selects the best network for each new communication according to their availability, load and performance. Taking full advantage of Autonomic Networking System, DWS delivers many benefits to enterprises including:

- Unify application performance across hybrid networks
- Improve business communication continuity
- Seamlessly integrate Cloud based applications
- Exploit large network capacity at low cost
- Turn back-up lines into business lines

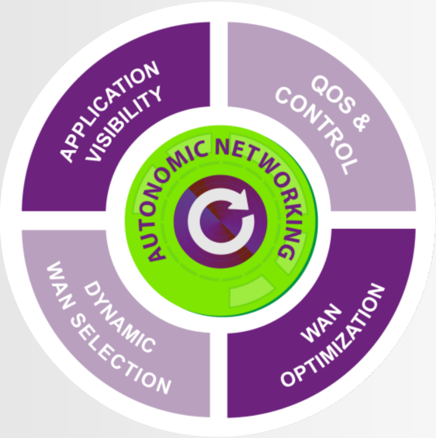

Powered by ANS^{TM}, WAN Governance brings tangible results

Get full visibility over your global network

- Eliminate 90% of network application performance issues
- Reduce problem identification and time-to-repair by 80%
- Ensure performance SLAs for all critical applications for 99,9% of the time

Deliver business applications

- Improve response time by 20x
- Reduce document download times from 5 minutes down to 15 seconds
- 0 business application brownouts during Olympic Games and Tour de France

Cost optimize your WAN

- Delay bandwidth upgrades by 24 months
- ÷3 the cost to transfer a Gbyte of data across the network
- Get full control with only 20% of technology expenses

ENSURE THE BEST NETWORK PERFORMANCE FOR PUBLIC CLOUD COMPUTING

STREAMCORE
MAKE YOUR NETWORK CONSCIOUS

Increasingly enterprises are using cloud computing to improve agility, efficiency and cost-effectiveness of IT operations. However, some enterprises fear the risks of migrating critical, time sensitive business applications to the public cloud because guaranteeing network performance over the Internet is very difficult. By providing visibility and performance control over centralized corporate Internet access links or sites with direct-to-branch Internet connectivity, Streamcore solutions ensure that the network does not negatively impact the performance of public cloud services.

The use of enterprise software-as-a-service (SaaS) applications, such as Webex, GoToMeeting, Salesforce, Google Apps and Microsoft Online Services, is on the rise. Aside from security and regulatory compliance issues, maintaining acceptable service levels is the biggest concern that enterprises have when considering public cloud services. These interactive or real-time SaaS applications are accessed by employees through corporate Internet access links, whether centralized or not, and compete with bandwidth intensive traffic such as recreational Web surfing, emails and software updates. Consequentially, network congestion can severely degrade the performance of SaaS traffic, hindering all the benefits of public cloud services.

WHAT IS NEEDED:

DEEP PACKET INSPECTION + AUTOMATED QOS + ADVANCED VISIBILITY

The adoption of cloud computing services results in the need for both controlled network performance and better WAN traffic visibility, two of Streamcore's core competencies. In order to apply visibility and control for cloud computing traffic, a third key feature is required, the capability to identify these cloud computing services on the network.

DPI engine for cloud traffic

Public cloud computing traffic is always encrypted and exchanged over HTTPS for obvious security reasons, making useless traditional classification processes based on TCP/UDP ports or even on URL for HTTP traffic.

Streamcore has developed a powerful Deep Packet Inspection (DPI) engine focused on business traffic, such as VoIP, videoconferencing and Web business applications, whether encrypted or not. The Streamcore solutions allows automatic identification and classification of encrypted Webex, Salesforce and other public cloud computing traffic in specific classes for monitoring and prioritization.

Automated advanced QoS

Streamcore dynamically applies traffic shaping and prioritization based on the DPI classification process. It eliminates network congestion on corporate Internet access links by prioritizing cloud-based traffic. A single business criticality parameter is required, making provisioning extremely simple.

Another unique Streamcore feature is the ability to automatically manage competition between users of the same application, based on each session's behavior. For example, if different users access a SaaS application, Streamcore's patented QoS engine will analyze the behaviour of each encrypted HTTPS session, and perform appropriate automated prioritization for interactive flows.

Advanced visibility

Streamcore also provides visibility of traffic usage and performance, with application response measurements and quality indicators for voice and video communications. These measurements help IT staffs continually monitor traffic performance and ensure that all cloud-based applications and communications are performing at acceptable levels for end users.

Visibility is provided in true real-time (over the last 10 seconds) and over the long-term for up to two years. Different set of tools are available, either through a Web portal or PDF email report, to share information with stakeholders or within the IT team.

STREAMCORE SOLUTIONS: FOR ANY INTERNET ACCESS ARCHITECTURE

Streamcore provides monitoring and traffic shaping with asymmetrical deployment. Therefore, any type of network architecture for Internet and public cloud computing access is supported.

Centralized Internet Access

Today, most enterprises centralize their gateway and accompanying demilitarized zones (DMZ) toward the public Internet through major data center hubs. This type of architecture is often required by the IT security team, in order to minimize risk and costs, and to ease management of security products. In this case, enterprises can deploy StreamGroomers, Streamcore traffic management appliances, in front of the centralized Internet access link, in order to manage all public cloud computing traffic and guarantee its performance.

If SaaS and cloud computing traffic has to be delivered to remote branch offices from the data centers via the centralized Internet access, additional StreamGroomers can be deployed in front of the data center's private WAN access links. The StreamGroomers can manage cloud computing traffic delivery to remote branch offices over the WAN.

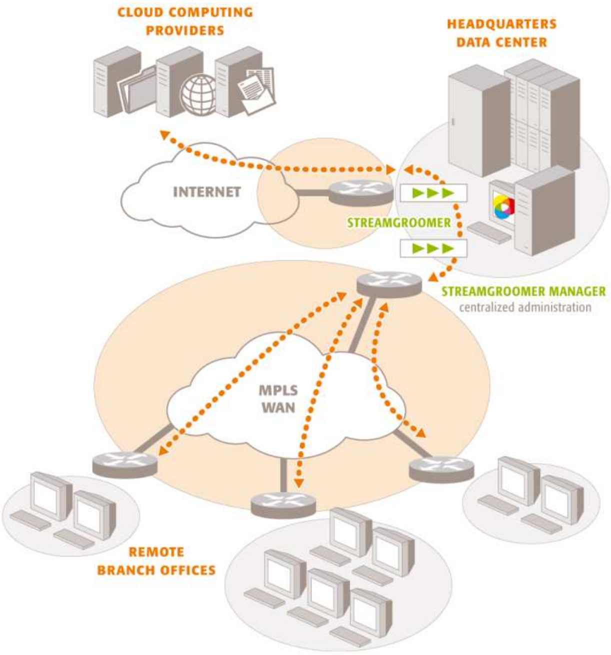

Fig. 1:
Centralized Internet access with branches over a private WAN

Branches with Direct Connections to the Internet

Backhauling Internet traffic to a centralized data center gateway adds latency and load on the private WAN.

Therefore, some enterprises prefer to provide Internet access directly into the VPN core, especially when they begin to rely heavily on public computing resources:

- If the private WAN uses IPSec technology over the Internet, this type of architecture is quite relevant. However, it can be challenging in terms of security because firewalls and security solutions (secure web gateways, antivirus...) must be fully distributed.
- Companies using a MPLS WAN sometimes have the option to migrate their Internet gateways and DMZs to the MPLS provider core. But, carriers may only offer a limited number of Internet gateways hubs around the world.

In such cases of branches with direct connections to the Internet, enterprises can deploy StreamGroomers in each branch office in order to manage and guarantee public cloud computing traffic performance over the branch WAN access link.

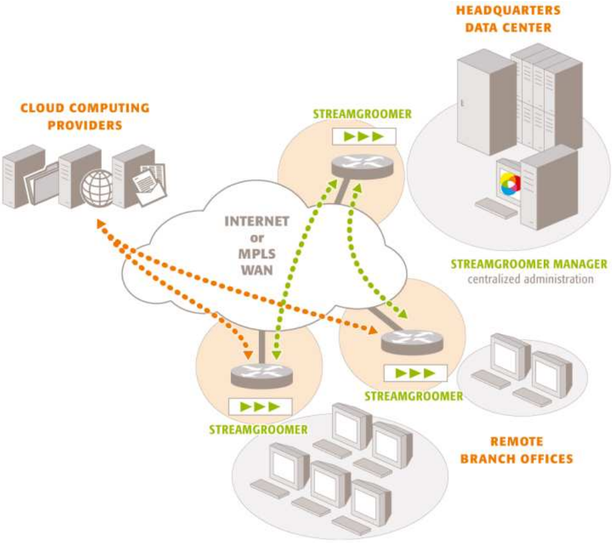

Fig. 2:
Branches with direct connections to the Internet

Hybrid Networks

On rare occasions, enterprises select an architecture in which there are two types of connectivity for each branch: an access link connected to a private MPLS network and another access link connected to a private IPSec network with direct-to-branch Internet access. This hybrid architecture combines the disadvantages of the two previous architectures: the high cost associated with MPLS, the complexity of securing distributed Internet gateways and the additional burden of managing traffic routed between the MPLS and the IPSec networks. However, this hybrid architecture can present advantages as well, such as extreme high availability, for enterprises with the budget and the right network/security team to manage it.

The full benefits of this architecture can be achieved by adding StreamGroomers at the branch: in addition to providing visibility and control for public cloud computing traffic, the Streamcore appliances can offer advanced load balancing per application. Bandwidth intensive applications can be automatically offloaded from the MPLS network to the IPsec network, and the MPLS access links can be dedicated to time-sensitive, real-time and business critical traffic.

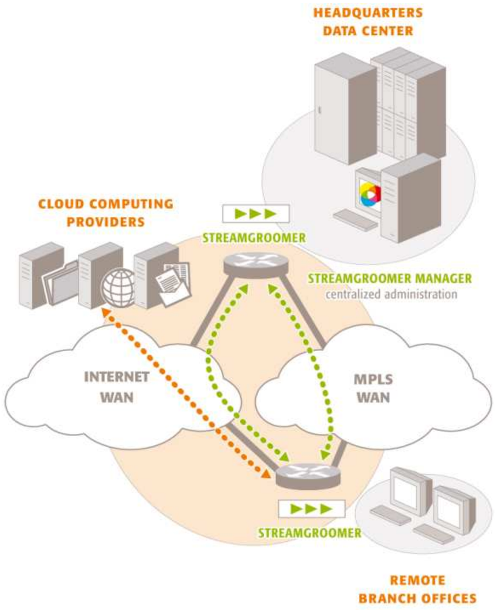

Fig. 3:
Hybrid network with load balancing
per application performed by StreamGroomers

SUMMARY

By providing DPI, automated QoS, and advanced visibility for public cloud computing traffic, Streamcore provides the best solutions to monitor and ensure the best performance for SaaS applications. Streamcore products are suitable for all types of architectures that provide access to public cloud computing applications including centralized Internet access, direct-to-branch Internet access, and even hybrid networks that combine MPLS and IPSec technologies.

For more information, visit www.streamcore.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.

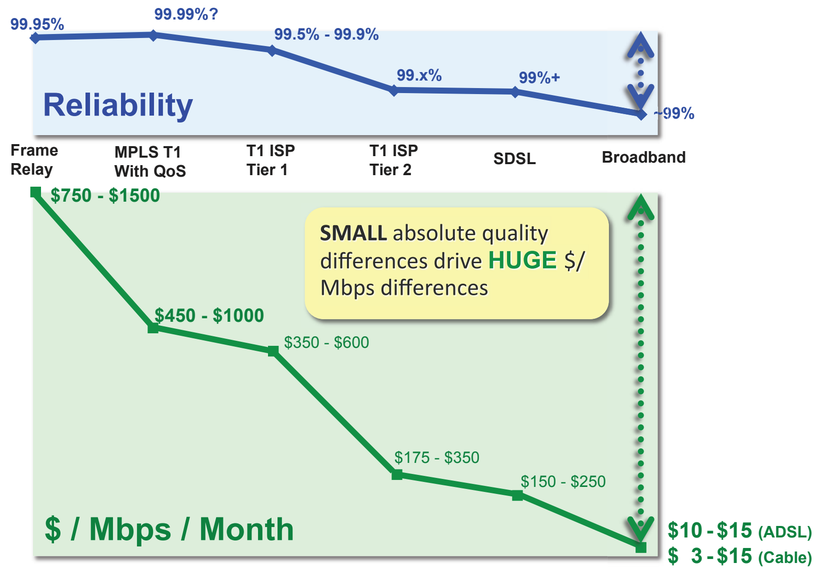

Figure 1: Private / Public WAN Pricing Disparity

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

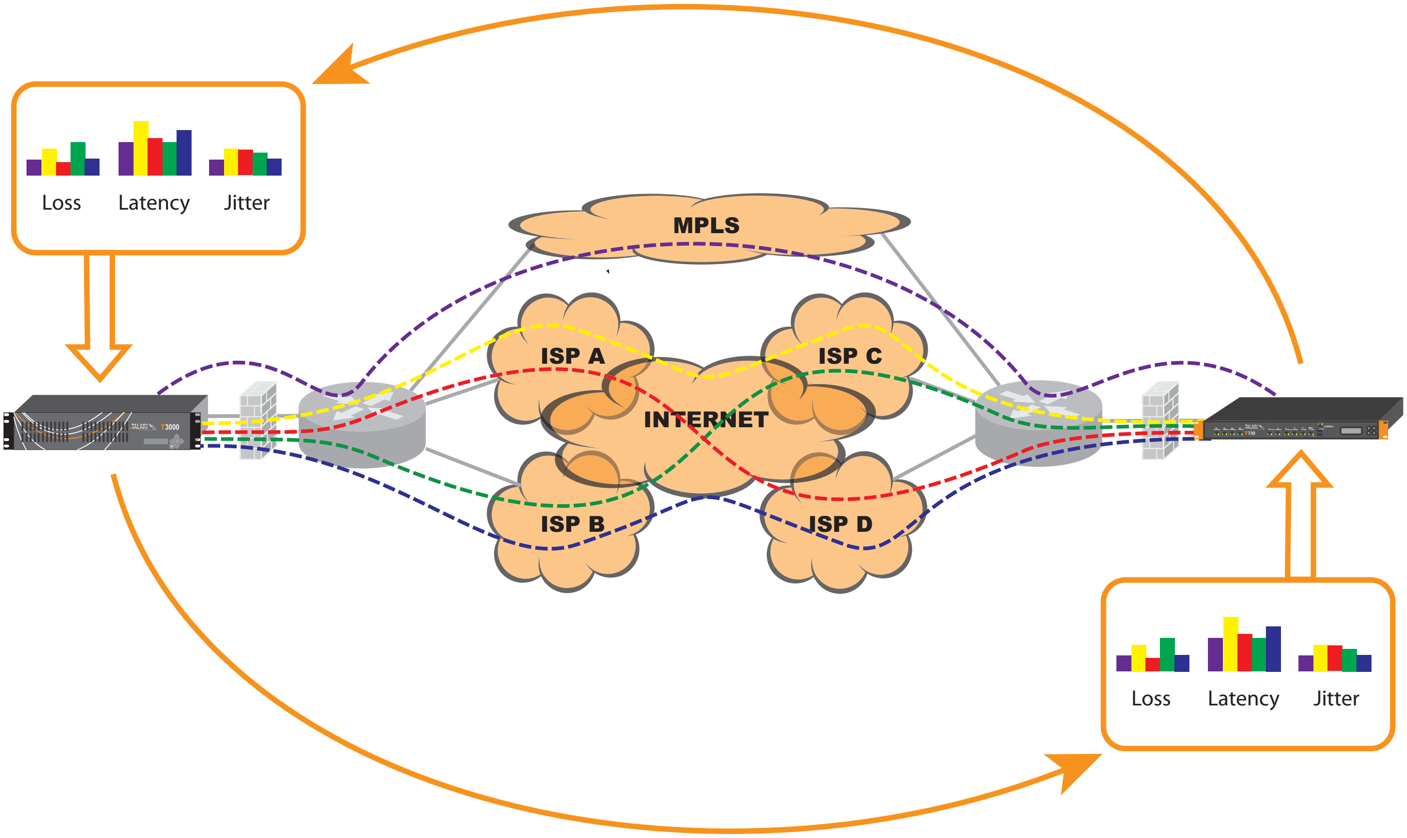

Figure 2: Continuous Measurment and Adaptation to Network Conditions

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**.

VYATTA®
Open Networking

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

Industry Praise

Gartner

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

Tolly

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

DELL

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

ZDNet

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

TechRepublic

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

The Vyatta Network OS

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

The Power of Open Networking

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

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

The New Network Requirements

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

The Vyatta Advantage

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

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

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

Deploying Vyatta in the Cloud: Common Use Cases:

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

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

CLOUD INFRASTRUCTURE

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

SECURE CONNECTIVITY

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

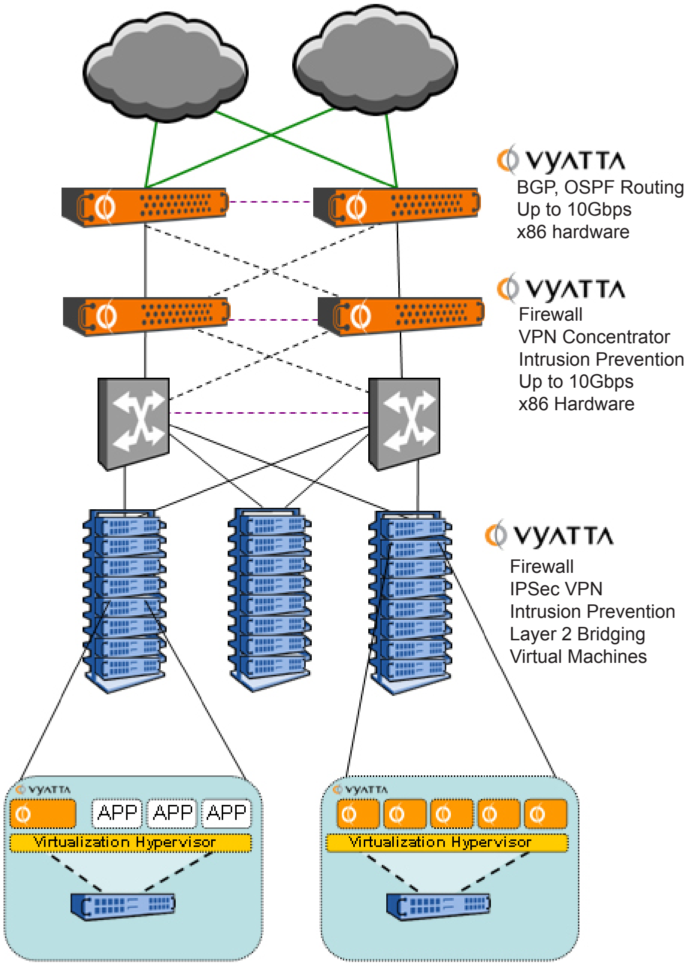

CLOUD ON-BOARDING - SECURE LAYER 2 BRIDGING

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

VIRTUAL FIREWALLING

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

The Vyatta Network OS

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

Vyatta Software Highlights:

Network Connectivity

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

Firewall Protection

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

Content and Threat Protection

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

Secure Connectivity

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

Traffic Management

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

High Availability

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

IPv6 Compatibility

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

Administration & Authentication

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

Monitoring and Reporting

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

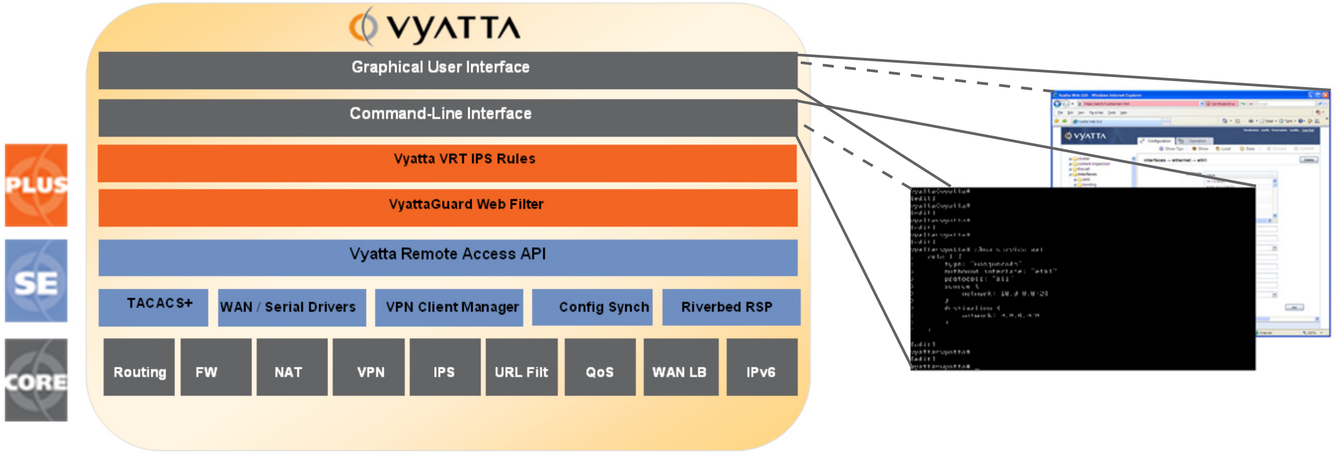

About Vyatta

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