The 2015
Application & Service Delivery Handbook

Part 2: Network and Application Optimization

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

Platinum Sponsors:

Gold Sponsors:

Produced by:

Webtorials

Table of Contents

Executive Summary 1

Key Optimization Tasks 1

Traditional Optimization Appliances 3
- WAN Optimization Controllers (WOCs) 3
- Application Delivery Controllers (ADCs) 5
- The Role of SDN 7
- NFV Optimization 8

The Next Generation WAN 10
- Background 10
- Factors Driving Change in the WAN 11
- Software Defined WANs 12

Executive Summary

The ***2015 Application and Service Delivery Handbook (The Handbook)*** will be published both in its entirety and in a serial fashion. This is the second of the serial publications. The first publication described how the application and service delivery environment is changing and the challenges and opportunities that the changing environment creates. This publication will focus on describing the technologies, products and services that are available to improve the performance of applications and services and the third publication will focus on the products and services that are available to improve the management and security of applications and services. The fourth and final publication will include an executive summary as well as a copy of the complete document.

The goals of the 2015 Application and Service Delivery Handbook are to help IT organizations to understand the emerging application and service delivery environment and to effectively respond to that environment.

Key Optimization Tasks

The previous chapter of The Handbook discussed two surveys that were given in early 2015 to the subscribers of Webtorials. As previously noted, within The Handbook the respondents to those surveys will be referred to as The Survey Respondents. The Survey Respondents were given a number of optimization-related tasks and were asked to indicate how important it was for their IT organization to get better at each task at over the next year. Their responses are shown in **Table 1.**

Table 1: The Importance of Key Optimization Tasks

	Not at All	Slightly	Moderately	Very	Extremely
Optimizing the performance of a key set of applications that are critical to the success of the business	0.0%	3.1%	15.5%	47.4%	34.0%
Ensuring acceptable performance for VoIP traffic	3.2%	14.7%	29.5%	26.3%	26.3%
Optimizing the performance of TCP	4.2%	12.6%	26.3%	40.0%	16.8%
Improving the performance of applications used by mobile workers	0.0%	18.3%	22.6%	37.6%	21.5%
Optimizing the performance of protocols other than TCP; e.g., HTTP and MAPI	5.2%	19.6%	25.8%	37.1%	12.4%
Optimizing the transfer of storage associated with business continuity and disaster recovery between different data centers	3.2%	8.5%	23.4%	42.6%	22.3%
Optimizing the performance of servers by offloading SSL and/or TCP processing	7.5%	17.2%	31.2%	32.3%	11.8%
Optimizing the performance of virtual desktops	11.2%	15.7%	40.4%	23.6%	9.0%

Table 1: The Importance of Key Optimization Tasks					
	Not at All	**Slightly**	**Moderately**	**Very**	**Extremely**
Controlling the cost of the WAN by reducing the amount of traffic by techniques such as compression	5.4%	9.7%	35.5%	34.4%	15.1%
Ensuring acceptable performance of traditional video traffic	3.3%	17.4%	27.2%	28.3%	23.9%
Optimizing the performance of applications and services acquired from public cloud providers	12.6%	20.7%	26.4%	28.7%	11.5%
Optimizing the transfer of virtual machines between data centers	14.6%	15.7%	33.7%	25.8%	10.1%
Optimizing the performance of chatty protocols such as CIFS	14.1%	17.6%	35.3%	22.4%	10.6%

Some of the conclusions that can be drawn from the data in **Table 1** are:

Optimizing the performance of a key set of applications that are critical to the business is the most important optimization task facing IT organizations.

Slightly less important than optimizing the performance of business critical applications is optimizing the transfer of storage associated with business continuity and disaster recovery between different data centers.

A relatively new challenge, ensuring the performance of applications used by mobile workers, is now one of the most important optimization tasks facing IT organizations.

Some traditional optimization challenges, such as optimizing the performance of TCP, VoIP and video traffic remain important.

Some traditional optimization challenges, such as optimizing the performance of chatty protocols, have become notably less important.

Optimizing the performance of applications and services acquired from public cloud providers, while not one of the most important optimization tasks, is growing in importance.

Traditional Optimization Appliances

For the last decade, the two primary optimization appliances have been WAN Optimization Controllers and Application Delivery Controllers.

WAN Optimization Controllers (WOCs)

An extensive discussion of the varying forms of optimization provided by a WAN Optimization Controller (WOC) can be found in the document ***Key WOC Functionality***.

When WOCs were first introduced in the mid-2000s, they were hardware-based appliances that IT organizations typically acquired and implemented on a do-it-yourself (DIY) basis. While that is still an option, as was mentioned in the previous chapter of The Handbook there is a broad movement underway within the IT industry to adopt a software-based approach to implementing virtually all types of IT functionality. Hence, while it is still possible to acquire a hardware-based WOC, software based WOCs are now available in a number of form factors, including:

- ***Standalone Hardware/Software Appliances***
  These are typically server-based hardware platforms that are based on industry standard CPUs with an integrated operating system and WOC software.

- ***Client software***
  WOC software can also be provided as client software for a PC, tablet or Smartphone to provide optimized connectivity for mobile and/or SOHO workers.

- ***Integrated Hardware/Software Appliances***
  This form factor corresponds to a hardware appliance that is integrated within a device such as a LAN switch or WAN router via a line card or other form of sub-module.

Software based WOCs are often referred to a being a virtual WOC (vWOC). The phrase virtual WOC refers to optimizing the operating system and the WOC software to run efficiently in a VM on a virtualized server. One of the factors that are driving the deployment of vWOCs is the previously discussed growing interest that IT organizations have in using Infrastructure-as-a-Service (IaaS) solutions. IaaS providers often don't want to install custom hardware such as WOCs for their customers. IT organizations, however, can bypass this reluctance by implementing a vWOC at the IaaS provider's site.

As is true with a number of virtual appliances, one advantage of a vWOC is that some vendors of vWOCs provide a version of their product that is completely free and is obtained on a self-service basis. The relative ease of transferring a vWOC also has a number of advantages. For example, one of the challenges associated with migrating a VM between physical servers is replicating the VM's networking environment in its new location. However, unlike a hardware-based WOC, a vWOC can be easily migrated along with the VM. This makes it easier for the IT organization to replicate the VM's networking environment in its new location.

Many IT organizations choose to implement a proof-of-concept (POC) trial prior to acquiring WOCs. The purpose of these trials is to enable the IT organization to quantify the performance improvements provided by the WOCs and to understand related issues such as the manageability and transparency of the WOCs. While it is possible to conduct a POC using a hardware-based WOC, it is easier to do so

with a vWOC. This follows in part because a vWOC can be downloaded in a matter of minutes, whereas it typically takes a few days to ship a hardware-based WOC. Whether it is for a POC or to implement a production WOC, the difference between the amount of time it takes to download a vWOC and the time it takes to ship a hardware-based appliance is particularly acute if the WOC is being deployed in a part of the world where it can take weeks if not months to get a hardware-based product through customs.

When considering vWOCs, IT organizations need to realize that there are some significant technical differences in the solutions that are currently available in the marketplace. These differences include the highest speed LAN and WAN links that can be supported as well as which hypervisors are supported; e.g., hypervisors from the leading vendors such as VMware, Citrix and Microsoft as well as proprietary hypervisors from a cloud computing provider such as Amazon. Another key consideration is the ability of the vWOC to fully leverage the multi-core processors being developed by vendors such as Intel and AMD in order to continually scale performance.

In addition to technical considerations, IT organizations also need to realize that there are some significant differences in terms of how vendors of virtual appliances structure the pricing of their products. One option provided by some vendors is typically referred to as *pay as you go*. This pricing option allows IT organizations to avoid the capital costs that are associated with a perpetual license and to acquire and pay for a vWOC or a virtual Application Delivery Controller (vADC) on an annual basis. Another option provided by some vendors is typically referred to as *pay as you grow*. This pricing option provides investment protection because it enables an IT organization to get stared by implementing vWOCs or vADCs that have relatively small capacity and are priced accordingly. The IT organization can upgrade to a higher-capacity vWOC or vADC when needed and only pay the difference between the price of the virtual appliance that it already has installed and the price of the virtual appliance that it wants to install.

In addition to acquiring a hardware-based or software-based WOC and implementing it on a DIY basis, IT organizations also have an additional way to acquire WOC functionality. They can acquire it from a service provider who offers network and application optimization as part of a WAN service.

IT organizations have a variety of options for how they acquire WOC functionality.

Application Delivery Controllers (ADCs)

Background

The original purpose of an ADC was to provide load balancing across local servers or among geographically dispersed data centers based on Layer 4 through Layer 7 intelligence. By providing this functionality, an ADC maximized the efficiency and availability of servers through intelligent allocation of application requests to the most appropriate server. ADCs, however, have assumed, and will most likely continue to assume, a wider range of more sophisticated roles that enhance server efficiency and provide asymmetrical functionality to accelerate the delivery of applications from the data center to individual remote users. 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 and hence can reduce the number of servers that are required for a given level of business activity.

The primary role of an ADC is to improve the utilization of compute resources.

A discussion of the traditional type of functionality provided by an ADC can be found in the document entitled *Primary Functionality Provided by an Application Delivery Controller*.

ADCs and Security

The security landscape has changed dramatically in the last few years. In the not so distant past, the typical security hacker worked alone, relied on un-sophisticated techniques such as dumpster diving, and was typically motivated by the desire to read about their hack in the trade press. In the current environment, sophisticated cyber criminals have access to malware networks and R&D labs and can use these resources to launch attacks whose goal is often to make money for the attacker. In addition, national governments and politically active hackers (hacktivists) are engaging in cyber warfare for a variety of politically motivated reasons.

The sophistication of computer attacks has increased dramatically in the last few years.

In addition to the growing sophistication of hackers, changes in IT, such as those created by the adoption of cloud computing, social networking and the new generation of mobile devices, combined with the ongoing evolution of regulatory requirements pose a spate of new security challenges. As a result, security currently is and likely will remain a significant application and service delivery challenge for the foreseeable future.

The role that the ADC plays in providing security was exemplified by the famous criminal Willie Sutton. Sutton was once asked why he robbed banks, and his response was simple, eloquent, and humorous: *Because that's where the money is*. In the case of IT security, the majority of the attacks are to a data center because that's where most of the applications and most of the data resides. Given that the most common deployment of ADCs has them placed in front of application servers in a data center, they are in a strategic position to thwart attacks. In order to be effective thwarting security attacks, ADCs should have an ICSA-certified web application firewall and a DNS application firewall. It should provide protection against DDoS attacks and also support SSL offload and high speed SSL decryption with SSL intercept. Additionally, authentication should be supported to traditional authentication stores such as Microsoft Active Directory. The ADC should support RADIUS, as well as the emerging class of identity

providers (IdPs) for maximum flexibility and authentication options. The ADC should also enable authentication support for passwords, one time passwords, and certificate options.

The Requirement for Programmability

As described in the preceding chapter of The Handbook, one of the ways that the application and service delivery model is changing is that there is an increasingly large adoption of cloud computing. One of the reasons for the ongoing success of cloud computing is that it provides for the dynamic allocation of IT resources. This has the effect of maximizing the utilization and hence minimizing the cost of those resources. As was also described in the preceding chapter of The Handbook, one of the key characteristics of a cloud computing solution is automation.

As cloud computing solutions evolve they tend to be inclusive of a growing range of services and the capability to manage those services. In order to support the scale and automation that is associated with cloud computing while simultaneously interoperating with an ever increasing set of products and services, an ADC needs to support open and standards-based programmability. The APIs that the ADC supports must ensure interoperability with the broadest possible range of automation, orchestration and analytics tools, such as that which is enabled by the use of RESTful APIs.

Virtual ADCs

Network appliances such as ADCs are evolving along two paths. One path is comprised of general-purpose hardware, a general-purpose hypervisor and a specialized O/S. The other path is comprised of specialized network hardware, specialized network hypervisors and a specialized O/S. This two-path evolution of network appliances has resulted in a wide array of options for deploying ADC technology. These options include:

- **General Purpose VM Support**
  A specialized network O/S along with ADC software that has been modified to run efficiently in a general purpose virtualization environment including VMWare's vSphere, Citrix's XenServer and Microsoft's Hyper-V.

- **Network Appliance O/S Partitioning**
  This involves the implementation of a lightweight hypervisor in a specialized network O/S by partitioning critical memory and I/O ports for each ADC instance, while also maintaining some memory and I/O ports in common.

- **Network Appliance with OEM Hypervisor**
  A general-purpose virtualization solution is adapted to run on a network appliance and provides the ability to run multiple ADCs on a single device. Since the hypervisor is based on an OEM product, other applications can be run on the device as it can participate in an enterprise virtualization framework such as VMWare's vCenter, Citrix's Xencenter or Microsoft's System Center. Support for loosely coupled systems (e.g. VMWare's VMotion and Citrix's XenMotion) is common.

- **Network Appliance with Custom Hypervisor**
  General-purpose hypervisors are designed for application servers and not optimized for network service applications. To overcome these limitations, custom hypervisors optimized for network O/S have been added to network appliances. Depending on the implementation, these specialized network hypervisors may or may not support loosely coupled systems.

Each of these approaches has advantages and disadvantages that effect overall scalability and flexibility. General purpose VM support has the most flexibility, but when compared to network appliance hardware, general purpose VM support gives the lowest level of performance and reliability. Network appliances with custom hypervisors can provide the greatest performance levels, but provide the least flexibility with limited co-resident applications and virtualization framework support.

The Role of SDN

In a traditional data center implementing L4 – L7 services such as WOCs and ADCs is cumbersome and time consuming as it requires acquiring the requisite network appliances and cabling them together in the correct order. Since each appliance has its own unique interface, configuring these appliances is an error-prone task. In addition, IT organizations have two alternatives relative to sizing these appliances. They can either size the appliances for the peak application load or they can resize the appliances on a regular basis to account for shifts in the traffic load. The first alternative results in stranded capacity and the second alternative results in an increase in the amount of manual labor that is required. In addition, because setting up a service tier is cumbersome, time consuming and error prone, service tiers are often built to support multiple applications. While this approach reduces the amount of manual labor that is required, it results in traffic needlessly passing through network appliances and consuming bandwidth and CPU cycles.

SDN holds the promise of overcoming the challenges of implementing L4 – L7 services by implementing two closely related techniques: service insertion and service chaining. The phrase *service insertion* refers to the ability to dynamically steer traffic flows to a physical or virtual server that provides L4 – L7 services. The phrase *service chaining* refers to the ability to dynamically steer traffic flows through a sequence of physical or virtual servers that provide a set of L4 – L7 services.

There are some key considerations that need to be taken into account relative to running L4 – L7 network services in a virtualized environment. These considerations include that IT organizations need to:

- Optimize provisioning
    - The virtualized L4 – L7 network services must integrate with orchestration platforms such as OpenStack so that they auto-deploy, auto-configure and auto-scale.

- Optimize placement
    - Service chaining should be implemented in such a way as to minimize traffic repeatedly going into and out of a device; a.k.a., hair-pinning.
    - Services should be placed in such a way that no individual fault zone (like a rack or a POD) becomes a single point of failure.

NFV Optimization

In order to obtain the potential cost and agility benefits of a software-based approach to providing IT functionality, it must be possible to achieve the same or greater performance in a software-based environment as is possible in a traditional hardware-based environment. However, that isn't possible without an enabling software architecture because of the bottlenecks that are associated with the hypervisors, virtual switches and virtual machines that are the foundation of the emerging software-based approach to IT. In response to the performance bottlenecks that are associated with NFV, ETSI has authored a document entitled "NFV Performance & Portability Best Practices".

Performance bottlenecks are not unique to virtualized environments. For example, some of the bottlenecks that occur in a physical environment are shown in **Figure 1.**

Figure 1: Bottlenecks in a Physical Environment

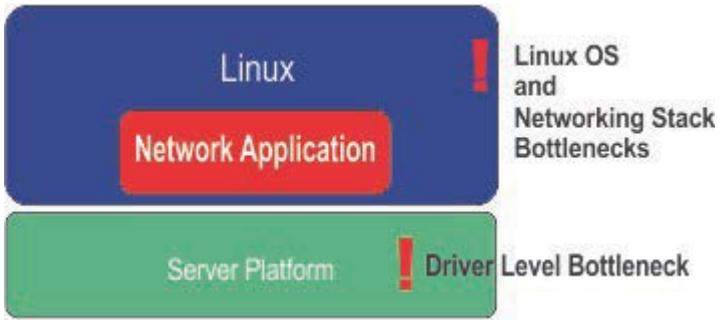

Figure 2: Bottlenecks in a Virtual Environment

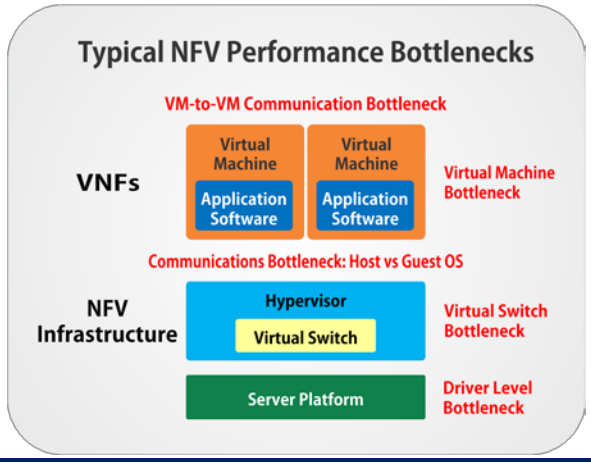

Unfortunately, as shown in **Figure 2**, as IT organizations adopt a virtualized environment the performance bottlenecks increase.

Acquiring solutions that have effective packet processing software that can bypass bottlenecks is one of the primary ways to avoid experiencing unacceptable performance in a virtualized environment. When evaluating the enabling packet processing software, IT organizations should check for the

following criteria in order to ensure a cost effective value proposition, and smooth transition to future requirements:

- Equal performance in both physical and virtual environments;
- Transparency: No change should be required to the operating system, the hypervisor, the virtual switch or to the management tools;
- Availability: The solution must work across multi-vendor processors, NICs and hardware platforms

The Next Generation WAN

Because the WAN introduces a range of demanding challenges relative to ensuring acceptable application and service delivery, this section will describe the current state of the WAN and how the WAN might evolve.

Background

WAN Evolution

The modern WAN got its start in 1969 with the deployment of the 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 around 1984 saw the deployment of four distinct generations of enterprise WAN technologies. 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.

WAN Services

As discussed in The 2014 State of the WAN Report, network organizations currently make relatively little use of WAN services other than MPLS and the Internet and the use they do make of those other services is decreasing somewhat rapidly. That report also identified the concerns that network organizations have with those two services. Those concerns are shown in Table 2 in descending order of importance.

Table 2: Concerns with WAN Services

Concerns with MPLS	Concerns with the Internet
Cost	Security
Uptime	Uptime
Latency	Latency
Lead time to implement new circuits	Cost
Security	Packet loss
Lead time to increase capacity on existing circuits	Lead time to increase capacity on existing circuits
Packet loss	Lead time to implement new circuits
Jitter	Jitter

Traditional WAN Design

The traditional approach to designing a branch office WAN is for each branch office to have either a T1 link or a set of bonded T1 links that provide access to a service provider's MPLS network and to have one or more higher speed links at each data center. In this design, it is common to have all or some of a company's Internet traffic be backhauled to a data center before being handed off to the Internet. One of the limitations of this design is that since the Internet traffic transits the MPLS link, this adds both cost and delay.

One alternative to the traditional approach to designing a branch office WAN is to supplement the T1 access link(s) in a branch office with direct Internet access and to also leverage technology such as Policy Based Routing (PBR). 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.

One advantage of this alternative design is that it enables network administrators to take Internet traffic off the relatively expensive MPLS link and put it on the relatively inexpensive Internet link. One disadvantage of this approach is that configuring PBR is complex, time consuming and error prone. Another limitation of this approach it that it creates a static allocation of traffic to multiple links which means that it isn't possible to reallocate the traffic when the quality of one of the links degrades.

Factors Driving Change in the WAN

As described in the preceding chapter of The Handbook, the application delivery model is changing. Two of the key changes are the growth in the number of mobile workers and the increasing use of public cloud services. In addition, as discussed in the preceding sub-section, it is important for network organizations that they get better over the next year at tasks such as supporting VoIP and video traffic. Supporting mobile workers, providing access to public cloud services and supporting real time applications are examples of requirements that are difficult and/or expensive to satisfy with a traditional WAN.

The Survey Respondents were given a set of factors and were asked to indicate which three factors will likely have the most impact on their WAN over the next twelve months? Their responses are shown in Figure 3.

Figure 3: Factors Driving Change in the WAN

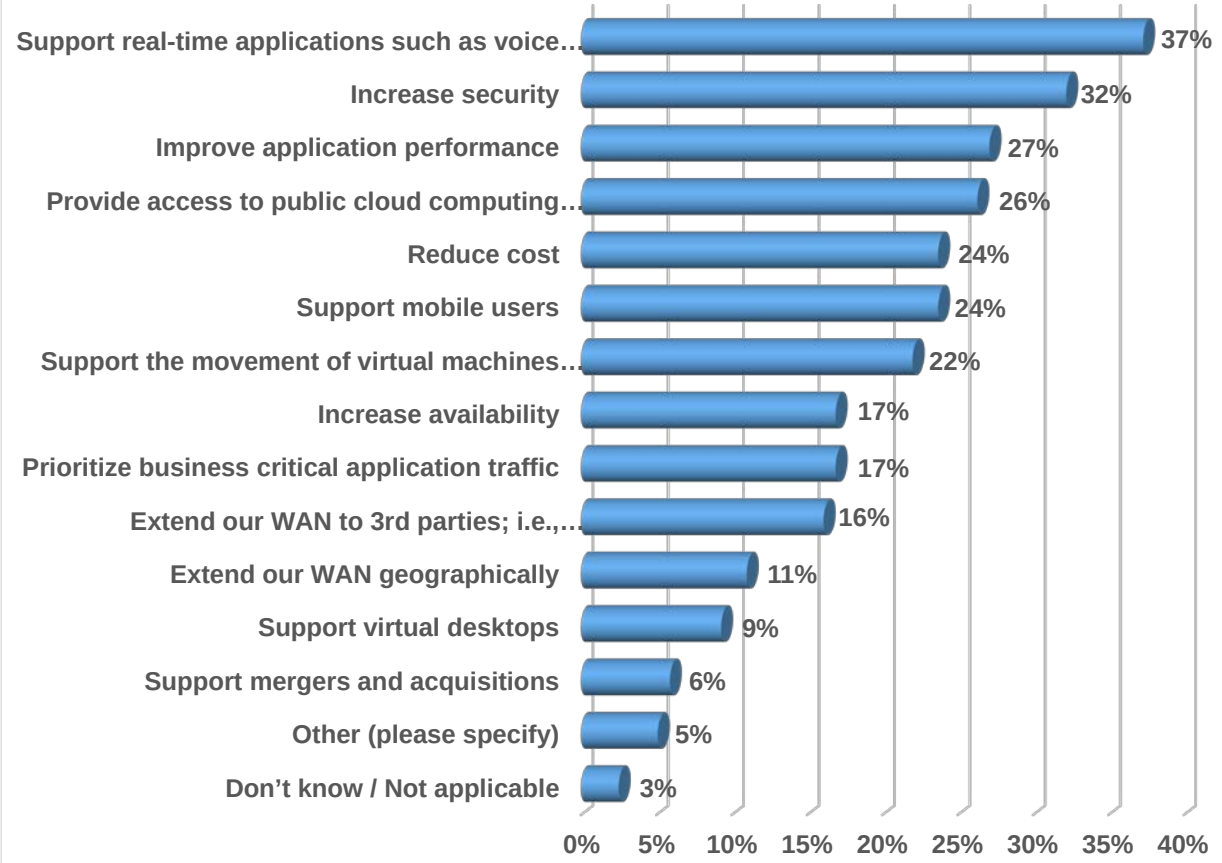

The data in **Figure 3** indicates that there is a wide range of WAN challenges that are important for network organizations to respond to.

Software Defined WANs

Hybrid WAN

As previously mentioned, the two primary concerns that IT organizations have relative to the use of the Internet are security and uptime and the two primary concerns that they have relative to the use of MPLS are cost and uptime. IT organizations can overcome some or all of these concerns by implementing a hybrid WAN; i.e., a WAN based on having two or more disparate WAN links into branch offices. There are many ways to construct such a hybrid WAN. One option is to have two connections to the Internet that are provided by different ISPs and which use diverse access such as DSL, cable or 4G. Another option is to have one WAN connection be an Internet connection and the other be a connection to an MPLS service.

The preceding discussion of the traditional approach to WAN design discussed having multiple WAN links at each branch office and using PBR to determine which traffic transited which WAN link. That discussion mentioned that the conventional way of implementing PBR results in the network not being able to respond in real time to changing network conditions. A relatively new class of functionality has

emerged to address the shortcomings of PBR. WAN Path Control (WPC) is one phrase that is often used to describe functionality that simplifies PBR and makes the selection of the best end-to-end WAN path based on real-time traffic analytics, including the instantaneous end-to-end performance of each available network; the instantaneous load for each end-to-end path; and the characteristics of each application.

Interest in Leveraging SDN in the WAN

The *2015 Guide to SDN and NFV* (The Guide) reported on the results of a survey that was administered in late 2014. The respondents to this survey were asked to indicate the factors that were driving their company's interest in SDN. The two factors that were indicated the most were:

- Better utilize network resources;
- Perform traffic engineering with an end-to-end view of the network.

While better utilizing network resources is a benefit of implementing SDN in either the LAN or the WAN, performing traffic engineering with an end-to-end view of the network is primarily a benefit of implementing SDN in the WAN.

The respondents to this survey further demonstrated their interest in implementing SDN in the WAN when they indicated how broadly they expected their campus, WAN and data center networks would be based on SDN three years from now. Their responses (**Table 3**) show that IT organizations believe that three years from now that SDN deployment in data centers will be highly pervasive and that there will also be significant SDN deployment both in the WAN and in campus networks.

Table 3: Anticipated SDN Deployment

	Campus Networks	WAN	Data Center Networks
Exclusively based on SDN	1%	2%	6%
Mostly SDN	10%	6%	20%
Hybrid, with SDN and traditional coexisting about equally	34%	36%	50%
Mostly traditional	29%	31%	10%
Exclusively traditional	13%	13%	4%
Don't know	12%	12%	10%

Definition of a Software Defined WAN

As is the case with any software defined network, a software defined WAN (SD-WAN) centralizes the control function into a SDN controller. The controller abstracts the user's private network services from the underlying IP network and it enables the operations of the user's private network services via centralized policy. The controller also enables the automation of management tasks such as configuration and provisioning.

Leveraging the underlying WAN platforms, which may include physical or virtual routers, the controller sets up virtual overlays that are both transport and technology agnostic. Under the direction of the controller, the WAN platforms implement functionality such as quality of service, path selection, optimization and security, often using dynamic multi-pathing over multiple WAN links.

Over half of The Survey Respondents (**Figure 4**) indicated that they either just learned about a SD-WAN from the definition that was in the survey instrument or they were just somewhat familiar with the concept. This lack of familiarity isn't surprising given that a SD-WAN is an emerging concept. It does, however, highlight the need for more education on this topic.

Figure 4: Familiarity with SD-WAN

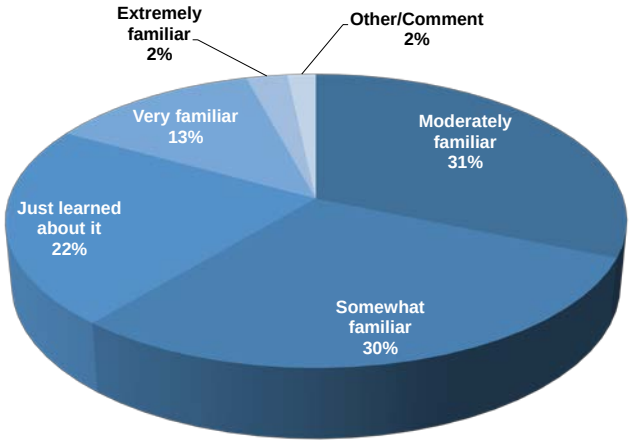

Drivers and Inhibitors of SD-WAN Adoption

The Survey Respondents were given a set of possible outcomes and were asked to indicate which outcomes would drive their company to implement a SD-WAN. Their responses are shown in **Table 4**.

Table 4: Drivers of SDN WAN Adoption

Drivers	Percentage
Increase flexibility	42%
Simplify operations	34%
Deploy new functionality more quickly	28%
Reduce OPEX	24%
Improve application performance	20%
Improve security	19%
Reduce CAPEX	19%
Improve availability	17%
Add bandwidth more quickly	17%
Provide better visibility	13%
Don't know/NA	9%
Other	3%

The perception of The Survey Respondents is that the top three drivers of SD-WAN deployment are:

- Increase flexibility;
- Simplify operations;
- Deploy new functionality more quickly.

There is no question that each of these drivers is important. However, each of these drivers is considered to be a soft savings which means that it can be difficult to show direct tangible benefits. For example, nobody would argue that it isn't a good thing to be able to deploy new network functionality more quickly, but what are the associated business benefits? Does it increase revenue? Reduce the company's bottom line cost? Reduce customer churn?

It is interesting and somewhat surprising that reducing OPEX was fourth on the list. While it can be difficult to build a business case for an investment in the WAN based on soft savings, it is relatively easy to build such a business case if there are hard cost savings. One of the key promises of a SD-WAN is that it will either reduce the amount of money that a company spends with their service providers or reduce how much that spend increases. The potential hard cost savings that result from implementing a SD-WAN is an important topic for vendors and network organizations to explore. Even if these hard savings don't justify a company making an investment in the WAN, the combination of hard and soft savings might.

The Survey Respondents were also given a set of factors and were asked to indicate which factors would inhibit their company from implementing a SD-WAN. Their responses are shown in **Table 5**.

Table 5: Inhibitors to SD-WAN Deployment	
Inhibitors	**Percentage**
The current technologies are unproven and/or immature	42%
It would add complexity	28%
The current products and/or services are unproven and/or immature	23%
We don't see a strong reason to adopt a SD-WAN	17%
It would not improve security and it could make it worse	15%
It could result in degraded application performance	14%
We would be locked into one vendor	13%
Don't know/NA	15%
Our contractual constraints with our WAN service providers limit what we can do	13%
It would increase CAPEX	8%
It would not improve visibility into WAN performance and it could make it worse	6%
Other (please specify)	6%

Some of the top inhibitors to SD-WAN deployment are the unproven and/or immature nature of the current technologies, products and services. Most likely these inhibitors will dissipate over time as the enabling technologies mature and vendors and service providers evolve their products and services. The fact that this survey data indicates that complexity is an inhibitor to SD-WAN deployment is in line with survey data presented in The Guide. That survey data shows that network organizations are concerned with the complexity associated with any implementation of SDN. Hopefully as technologies, services and products mature, vendors and service providers will ensure that complexity is no longer an issue.

The fact that network organizations don't see a strong reason to adopt a SD-WAN is in line with the previous discussion that network organizations see that the top three drivers of SD-WAN are soft savings and that it can be difficult to make a compelling business case based on soft savings. As previously mentioned, it is relatively easy to make a compelling business case if there are hard savings and vendors need to help network organizations create these business cases.

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.

Jim Metzler has a broad background in the IT industry. This includes being a software engineer, an engineering manager for high-speed data services for a major network service provider, a product manager for network hardware, a network manager at two Fortune 500 companies, and the principal of a consulting organization. In addition, he has created software tools for designing customer networks for a major network service provider and directed and performed market research at a major industry analyst firm. Jim's current interests include cloud networking and application delivery.

For more information and for additional Webtorials® Editorial/Analyst Division products, please contact Jim Metzler or Steven Taylor.

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

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

Professional Opinions Disclaimer

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

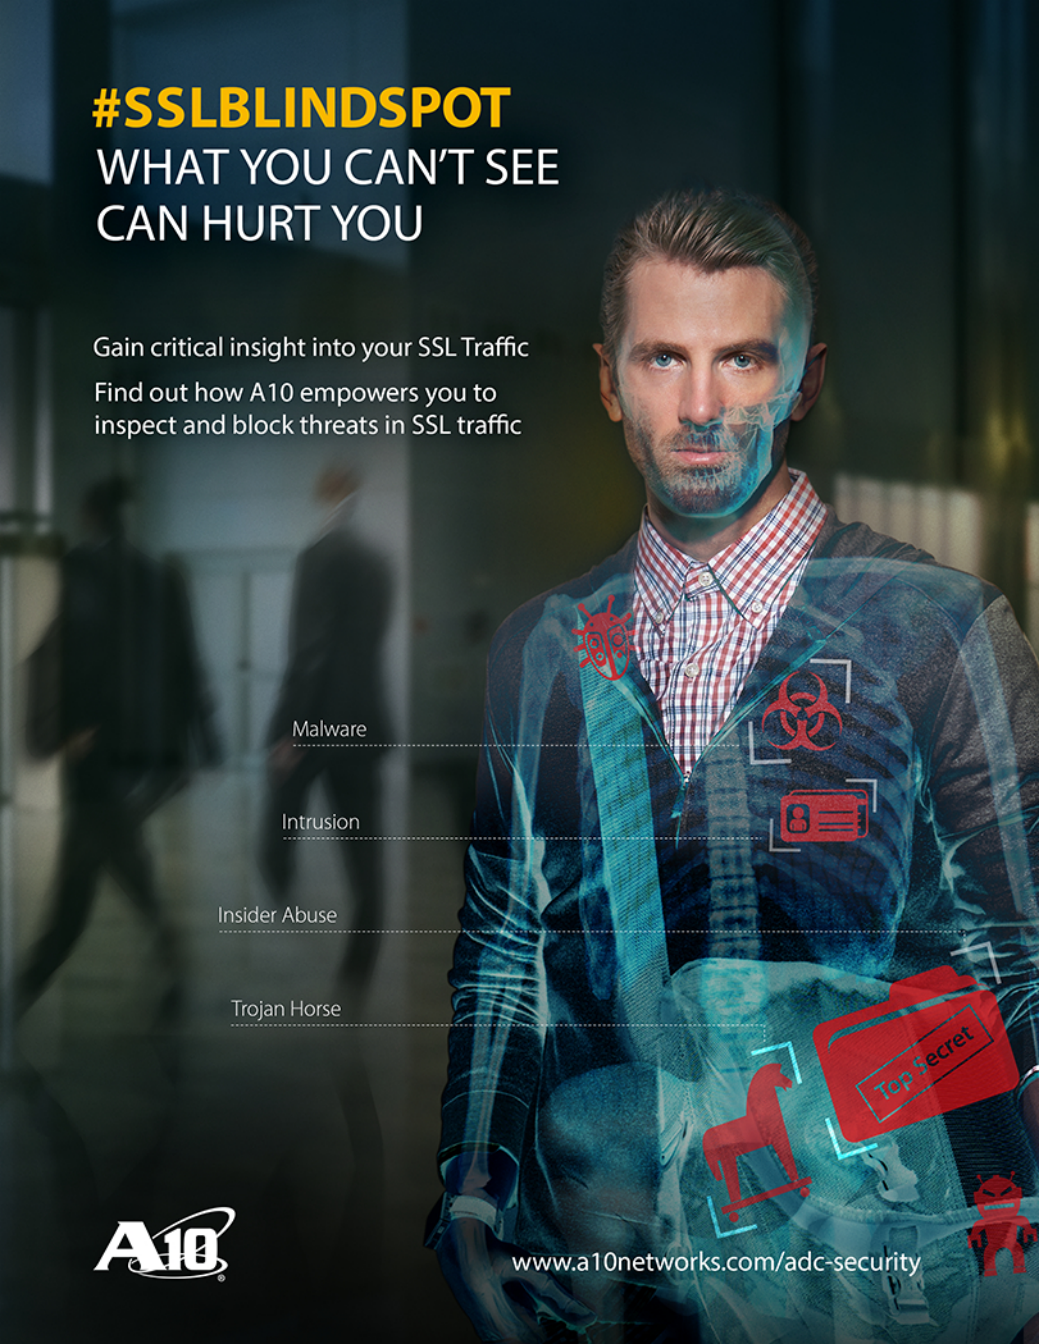
#SSLBLINDSPOT
WHAT YOU CAN'T SEE
CAN HURT YOU
Gain critical insight into your SSL Traffic
Find out how A10 empowers you to
inspect and block threats in SSL traffic
Malware
Intrusion
Insider Abuse
Trojan Horse
Top Secret
A10
www.a10networks.com/adc-security

SDN Today:

Delivered by Citrix NetScaler and Cisco ACI

Learn more at citrix.com/netscaler/cisco

CITRIX®

cisco™

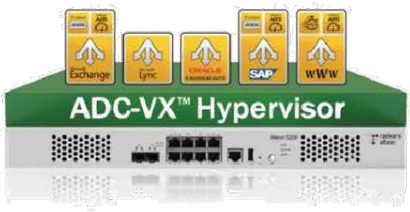

Predictable Application Service Levels, Guaranteed—Only with Alteon NG

Whether it's an online web application, or an internal mission-critical enterprise application such as CRM, ERP, or an organizational portal, end-users expect to receive the same, unchanged quality of experience. The conclusion is clear: today's organizations require **predictable application service levels** and need tools to proactively monitor and manage application service levels.

The Standard ADC: Not Good Enough Anymore

For years, companies have been using application delivery controllers (ADC) to optimally deliver applications. However, the standard/legacy ADC is not enough anymore as it is based on a **best-effort approach**.

In contrast to the legacy ADC, **a next-generation (NG) ADC** can provide full application SLA assurance through reserving resources per application. This allows the addition of new services without performance penalty and the inclusion of real-user monitoring, best-in-class application-level acceleration features and an innovative security offering.

Alteon NG: Complete Application Service Level Assurance

The Alteon® next-generation (NG) ADC solution is the industry's only ADC built from the ground up to ensure application service levels at all times. It innovatively leverages several next-generation services that are not available in any other ADC on the market:

☑ Alteon NG is **architecturally designed to ensure application service levels** by delivering full resource isolation per application, service, or department. Each virtual ADC (vADC) instance is completely isolated from neighboring instances with independent CPU cores, memory, network stack, management control, and operating system. Our unique solution is designed to dynamically scale to add more throughput, services, and vADCs without hardware modification resulting in fast provisioning of additional vADC instances and no service degradation, interruption, or resource overcapacity.

☑ Alteon NG is designed to deliver **secured ADC services**, both through its integrated security modules, such as the web application firewall (WAF), its ADoS and DDoS protection module, and also through its tight integration with Radware's unique **Attack Mitigation System (AMS).** The result is an architecture which enables accurate

detection and mitigation of the most advanced cyber-attacks at the ADC level, and then by leveraging the unique Defense Messaging™ the application delivery service signals attack information to Radware DefensePipe cloud service and/or Radware DefensePro data center attack mitigator, located in the cloud or the network perimeter, respectively to block the attack before it even reaches the datacenter's network.

☑ Alteon's Integrated advanced **Web Application Firewall (WAF)** module, enables risk-free implementation thanks to a unique out-of-path WAF deployment mode along with auto-policy generation capabilities. ADC resources are ensured via full instance isolation and resource reservation, even when WAF policies are updated there's no impact on application availability and performance. Moreover, as attacks are mitigated through DefensePro and/or defense pipe in the perimeter / cloud (thanks to the Defense Messaging™ mechanism), the WAF module can never become a bottleneck for detecting and mitigating attacks. This results in secured web applications with SLA guarantee.

☐ Radware's Application Performance Monitoring (APM) module provides real-time tracking of application service levels by measuring real-user transactions and errors. Embedded in Alteon NG,
Radware's APM is an out-of-the-box solution which doesn't require synthetic transaction scripting or additional installation - reducing deployment time and costs. Radware's APM intuitively tracks SLA by location, user, application and transaction type to expedite root cause analysis. In addition, it provides historical reports based on user-defined SLA that feature granular analysis allowing the measurement of the delay per transaction phase including data center time, network latency and browser rendering time.

☑ Alteon NG integrates **FastView®**, the industry's most advanced **Web Performance Optimization (WPO)** technology – which accelerates application response by up to 40% – for higher conversion rates, revenues, productivity, and customer loyalty. FastView acceleration treatments are optimized according to each user, end-user device and browser - with specific optimization for mobile devices. In addition, FastView automatically optimizes new applications, new application versions and new application modules – reducing manual code optimization while letting you focus on core business competencies.

☑ Alteon NG features a built-in authentication gateway with **Single Sign On (SSO)** capabilities by supporting Radius, Active Directory, LDAP and RSA SecurID – simplifying the user experience without compromising on application security.

Want to see more for yourself? We invite you to visit www.radware.com or contact us at: info@radware.com.