The 2014
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

The ***2014 Application and Service Delivery Handbook (The Handbook)*** will be published both in its entirety and in a serial fashion. The first chapter was already published and it describes how the ***2014 Application and Service Delivery Handbook*** differs from previous editions in this series. That chapter also describes how a variety of traditional and emerging factors are complicating the task of ensuring acceptable application and service delivery.

This is the second of the serial publications and contains the second chapter of ***The Handbook***. This chapter of ***The Handbook*** focuses on describing the technologies, products and services that are available to improve the performance of applications and services. The third chapter of ***The Handbook*** will focus on describing the technologies, products and services that are available to improve the management and security of applications and services.

Each of the first three chapters contains recent market research that identifies how IT organizations are approaching application and service delivery. The fourth and final publication will include an executive summary of ***The Handbook*** as well as a copy of the complete document.

Background

The phrase ***network and application optimization*** refers to an extensive set of techniques that organizations have deployed in an attempt to optimize the performance of networked applications and services while also controlling WAN bandwidth expenses. Over the last several years, the primary role that these techniques have played has been to:

- Reduce the amount of data sent over the WAN;
- Ensure that a WAN link is never idle if there is data to send;
- Dynamically load balance traffic over multiple WAN links;
- Reduce the number of round trips (a.k.a., transport layer or application turns) necessary for a given transaction;
- Overcome the packet delivery issues that are common in shared networks that are typically over-subscribed;
- Mitigate the inefficiencies of protocols and applications;
- Offload computationally intensive tasks from client systems and servers;
- Use multiple LAN links from origin to destination where appropriate;
- Direct traffic to the most appropriate server based on a variety of metrics.

In addition, as described in Chapter 1 of ***The Handbook***, the IT industry is adopting a software-based approach to providing most, if not all types of IT functionality; i.e., software defined data centers, virtual servers, virtual desktops, software defined networking, network function virtualization. This movement to software-based solutions creates some performance challenges. For example, 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.

As mentioned in the previous chapter of ***The Handbook***, a key component of the traditional approach to ensuring acceptable application delivery has been to implement hardware-based appliances on premise. This includes appliances such as WAN optimization controllers (WOCs)

and Application Delivery Controller (ADCs). However, as described in this chapter of ***The Handbook***:

- Most, if not all, of the optimization-related functionality that used to only be available as a hardware-based appliance is now also available as a software-based appliance;

- There is new functionality that is being deployed on premise to enable organizations to:
  - More effectively leverage the Internet;
  - Respond to the performance challenges associated with Network Function Virtualization (NFV).

- There are a set of emerging WAN technologies and services that optimize traffic performance and reduce the cost of the WAN.

Key Optimization Tasks

The previous chapter of ***The Handbook*** discussed a survey that was given in early 2014 to the subscribers of Webtorials. As previously noted, within ***The Handbook*** the respondents to that survey will be referred to as The Survey Respondents. **Table 1** shows how The Survey Respondents answered the survey question about the optimization tasks that their IT organizations are most interested in getting better at over the next year.

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	1.2%	4.3%	11.2%	45.3%	37.9%
Ensuring acceptable performance for VoIP traffic	3.1%	5.7%	15.1%	42.8%	33.3%
Optimizing the performance of TCP	3.7%	8.1%	33.5%	33.5%	21.1%
Improving the performance of applications used by mobile workers	5.6%	10.6%	28.8%	33.1%	21.9%
Optimizing the performance of protocols other than TCP; e.g., HTTP and MAPI	4.4%	13.8%	33.8%	31.3%	16.9%
Optimizing the transfer of storage between different data centers	7.3%	11.3%	23.2%	36.4%	21.9%
Optimizing the performance of servers by offloading SSL and/or TCP processing	9.7%	15.5%	32.3%	30.3%	12.3%
Optimizing the performance of virtual desktops	5.8%	9.7%	26.0%	37.7%	20.8%

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	9.9%	14.8%	25.3%	32.7%	17.3%
Ensuring acceptable performance of traditional video traffic	4.3%	10.5%	27.2%	38.3%	19.8%
Optimizing the performance of applications that you acquire from a SaaS provider such as Salesforce.com	12.3%	16.9%	31.8%	23.4%	15.6%
Ensuring acceptable performance for telepresence traffic	11.6%	11.6%	20.0%	37.4%	19.4%
Optimizing the performance of chatty protocols such as CIFS	10.5%	20.3%	32.7%	26.1%	10.5%
Optimizing the performance of the computing services that you acquire from a third party such as Amazon	14.7%	14.0%	36.4%	23.1%	11.9%

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; followed closely by the need to ensure acceptable performance for VoIP traffic. While these were also the two most important optimization tasks in 2013, their importance has increased notably in the last year.

Some traditional challenges, such as optimizing the performance of TCP, remain very important while other traditional challenges, such as optimizing the performance of chatty protocols, have become notably less important.

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.

Another challenge that is increasing in importance is optimizing the transfer of storage between different data centers.

The 2013 edition of ***The Handbook***[1] contains an extensive discussion of some of the key optimization challenges, such as the challenge associated with moving storage between data centers.

[1] http://www.webtorials.com/news/2013/10/2013-application-service-delivery-handbook-1.html

Traditional Optimization Appliances

WAN Optimization Controllers (WOCs)

The 2013 edition of ***The Handbook*** also contains an extensive discussion of the functionality provided by a WOC.

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, in the current environment, it is also possible for IT organizations to acquire WOC functionality from a managed service provider (MSP). IT organizations also have a third option because some providers offer network and application optimization as part of a WAN service.

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

As was mentioned, there is a 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.

As discussed below, ADCs are also available in a variety of form factors. Given the breadth of ways in which WOC and ADC functionality can be consumed, The Survey Respondents were asked to indicate their organization's current usage as well as their planned usage over the next year of a range of ways that this functionality can be consumed. Their responses are shown in **Table 2**.

Table 2: Current and Planned Usage of Optimization Functionality					
	Will Begin to Use	**Use now, but will have a major decrease**	**Use now, and no major changes**	**Use now, with major increase**	**Will not use**
Optimization functionality as part of a managed service, such as a managed and optimized WAN service	15.1%	6.5%	29.5%	28.1%	20.9%
A purpose built, stand-alone WAN optimization controller (WOC)	9.2%	6.9%	36.2%	13.8%	33.8%
WAN optimization functionality that is integrated into another device, such as a router	12.6%	6.3%	42.0%	23.8%	15.4%
Software-based WAN optimization functionality running on a virtual machine or a server	16.3%	3.0%	36.3%	22.2%	22.2%
A purpose built, stand-alone ADC	11.6%	5.0%	37.2%	14.9%	31.4%
Software-based ADC functionality running on a virtual machine or a server	14.7%	4.7%	38.0%	16.3%	26.4%

Some of the conclusions that can be drawn from Table 2 are:

Only a tiny minority of IT organizations plan on having a major reduction in their use of any of the ways that they currently consume optimization functionality.

While there is interest in expanding the use of hardware-based optimization solutions, the primary interest is in expanding the use of software-based optimization solutions.

Of the varying ways to consume optimization functionality, the two ways with the largest percentages of respondents indicating "will not use" are hardware-based WOCs and hardware-based ADCs.

The 2013 edition of ***The Handbook*** detailed a number of factors that are driving the adoption of virtualized WOCs (vWOCs) and discussed some of the technical attributes of vWOCs that IT organizations should consider when evaluating this class of product. In addition to technical considerations, IT organizations also need to realize that there are some significant differences in terms of how vendors of vWOCs 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 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 with WAN optimization by implementing vWOCs that have relatively small capacity and are priced accordingly. The IT organization can upgrade to a higher-capacity vWOC when needed and only pay the difference between the price of the vWOC that it already has installed and the price of the vWOC that it wants to install.

In some cases, vendors of virtual ADCs offer the same type of pay as you go and pay as you grow pricing options.

Application Delivery Controllers (ADCs)

Background

ADCs 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 maximizes 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.

The 2013 edition of ***The Handbook*** contains a lengthy discussion of the type of functionality provided by an ADC. That document also discussed high availability options and trends in the evolution of ADCs.

ADCs and Security

The security landscape has changed dramatically in the last few years. In the very recent 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.

Security is both a first and a second-generation application and service delivery challenge and it will remain a significant challenge for the foreseeable future. Rapid 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 regulations pose a spate of new

challenges for IT security systems and policies in much the same manner that they present challenges to the IT infrastructure.

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*[2]. 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. Is should also support SSL offload and high speed SSL decryption with SSL intercept.

IPv6 and ADCs

Background

Gartner estimates that 17% of the global Internet users and 28% of new Internet connections will use IPv6 by 2015.[3] This is creating an imperative for enterprises to develop an IPv6 strategy and migration plan. A key component of that strategy and migration plan is ensuring that devices such as ADCs that you are implementing today, fully support IPv6.

The 2013 edition of ***The Handbook*** describes a number of standards and technologies that help with IPv6 migration. These include:

- **Tunneling** – Transporting IPv6 traffic in IPv4 areas and vice versa.
- **Network Address Translation (NAT)** – Translating between IPv4 and IPv6 addresses, including DNS support.
- **Carrier Grade NAT (CGN)** – Contains more features than NAT and is based on IETF reference [draft-nishitani-cgn-05].
- **Dual Stack** – Both IPv4 and IPv6 packets are processed by devices simultaneously.

There are a variety of approaches that can be used to implement IPv6. One approach is that IPv6 to IPv4 services can be purchased via the ISP. Another approach is that IPv6 can be implemented on the data center perimeter firewalls and translated to the existing IPv4 infrastructure. A third approach is that Application Delivery Controllers can translate between IPv6 and IPv4 for application servers.

IT organizations may choose to utilize all three approaches in stages. For example, an IT organization may choose to start by relying on their ISP for IPv6 presence and then implementing IPv6 on their data center's perimeter firewalls. Once the data center perimeter firewall supports IPv6, attention can now turn to Application Delivery Controllers (ADCs) that provide load balancing, SSL offloading, WAN optimization, etc.

[2] http://quoteinvestigator.com/2013/02/10/where-money-is/
[3] http://www.verisigninc.com/assets/preparing-for-ipv6.pdf

ADCs can have the following IPv6 capabilities[4]:

- Ability to provide IPv6/IPv4 Dual Stack for Virtual IPs (VIP)
- Server Load Balancing with port translation (SLB-PT/SLB-64) to IPv4 servers (and the ability to transparently load balance a mix of IPv4 and IPv6 servers)
- 6rd
- NAT64 and DNS64 (to provide IPv6 name resolution services for IPv4-only servers)
- Dual-stack Lite (DS-lite)
- SNMP IPv4 and IPv6 support for monitoring, reporting and configuration
- Ability to provide utilization and usage statistics separated by IPv4 and IPv6

Using the ADC to implement IPv6 migration gives an IT organization the ability to insert Dual Stack IPv6/IPv4 or IPv6 only servers transparently into production. This is a critical first step to providing a low risk application server IPv6 migration path, which in turn is needed to gain access to a larger IP address pool for new and expanded applications. Just using the ISP or data center perimeter firewall for IPv6 does not provide the scalability nor the routing nor security benefits of IPv6.

Virtual ADCs

The 2013 edition of The Handbook analyzed how network appliances in general, and ADCs in particular were evolving. One of the conclusions that was drawn was that:

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 have 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

[4] http://www.a10networks.com/news/industry-coverage-backups/20120213-Network_World-Clear_Choice_Test.pdf

Microsoft's System Center. Support for loosely couple systems (e.g. VMWare's VMotioin 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 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.

NFV Optimization

As was previously mentioned, 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[5]"

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: Representative Bottlenecks in a Physical Environment

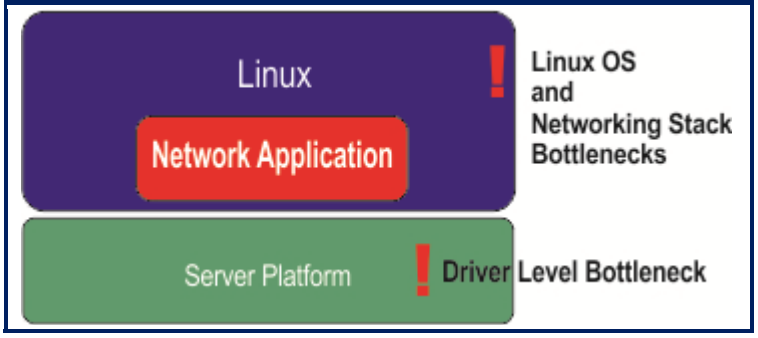

[5] http://docbox.etsi.org/ISG/NFV/Open/Latest_Drafts/NFV-PER001v009%20-%20NFV%20Performance%20&%20Portability%20Best%20Practises.pdf

Figure 2: Performance Bottlenecks in a Virtualized Environment

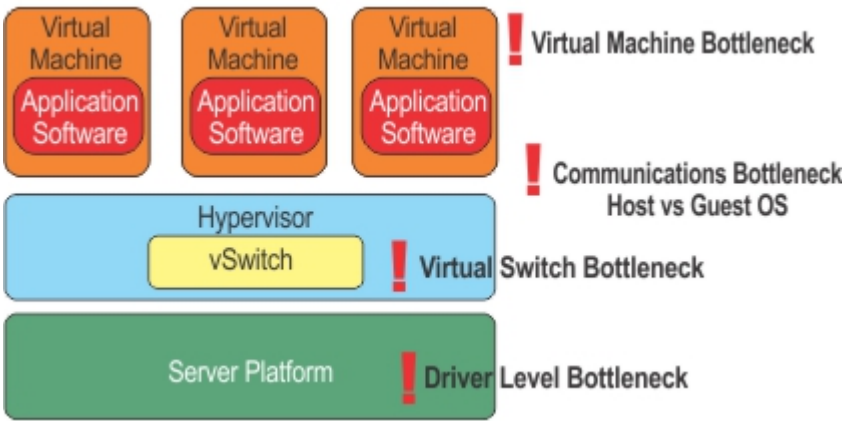

Unfortunately, as shown in **Figure 2**, as IT organizations adopt a virtualized environment the performance bottlenecks multiply. **Figure 2** demonstrates some, but not all of the bottlenecks that can occur in a virtualized environment. For example, while not explicitly shown in **Figure 2**, VM to VM communications can also result in bottlenecks.

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. As shown in **Figure 3**, 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.

Figure 3: Breadth of Environments

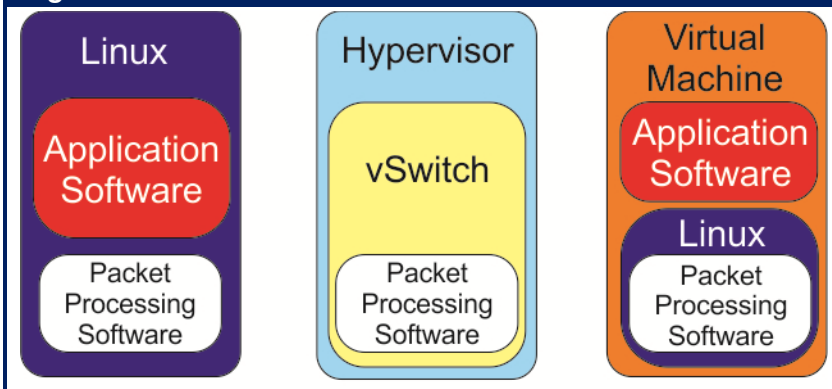

The evaluation criteria listed above are intended to ensure that the packet processing software can be easily and universally implemented on any version of Linux or on any hypervisor, without requiring changes to existing environments.

The types of performance improvements that are possible are significant. For example, it is possible to leverage packet processing software to accelerate the performance of a virtual

switch, such as Open vSwitch, by a factor of 10 or more. Some examples of high performance Virtual Network Functions (VNFs) designed with effective packet processing software include:

- An accelerated TCP/UDP stack that enables the building of products such as stateful firewalls, DPI engines, cloud servers and web servers that support millions of concurrent sessions and also support session setup rates above one million sessions per second.

- A high performance IPsec stack that can sustain more than 190 Gbps of encrypted traffic on a single server.

- High performance and capacity for encapsulation protocols such as GRE, GTP, PPP, L2TP. An example of this is a vBRAS server that can handle 256,000 PPPoE tunnels with 70 Gbps throughput.

Emerging WAN Optimization Services

Background

In the traditional IT environment, the end users reside in a corporate office and the applications and data that the users need to access are housed in a corporate data center. All of the resources in the corporate data center (i.e., the servers, storage or networks) are under the control of the IT organization. In the vast majority of cases the connectivity between the corporate offices and the corporate data center is provided primarily by a WAN service such as MPLS. For all of the reasons highlighted in the document *Traditional Application and Service Delivery Challenges*[6], ensuring acceptable application and service delivery in a traditional environment such as this is challenging.

While the IT environment that was described in the preceding paragraph is still somewhat common, a different IT environment is becoming increasingly common. One of the key characteristics of this new environment is that the users are mobile and use a wide array of access devices. One type of mobile user resides in a corporate facility and as described below, increasingly the Internet is used to provide WAN connectivity out of a corporate facility. A second type of mobile user accesses corporate applications and data from an external facility where network connectivity is either WiFi or cellular access to the Internet. Another key characteristic of this emerging IT environment is that users are increasingly accessing applications and data that are provided by cloud service providers.

The traditional optimization appliances (e.g., WOCs and ADCs) provide significant value in an environment where the users as well as the applications and data the users are accessing are in a fixed location and under the control of the IT organization. However, as described in this section of ***The Handbook***, a new set of optimized WAN services is emerging which is highly complementary to the traditional approach to optimization. This emerging set of solutions is focused on environments in which one or both of the end points is either not in a fixed location or not under the control of the IT organization. An example of this is a mobile user accessing applications and data from a cloud provider.

[6] http://www.ashtonmetzler.com/Traditional%20App%20Delivery%20Challenges%20V2.0.pdf

There is no doubt that ensuring acceptable application and service delivery is even more challenging in the emerging environment than it is in the traditional environment. There is also no doubt that:

In order to continue to show business value, IT organizations must be able to ensure acceptable application and service delivery independent of the type of IT environment.

Cloud-Based, Private WAN Optimization Services

In a cloud-based, private WAN optimization service a variety of types of users (e.g., mobile users, branch office users) access WAN optimization functionality at the service provider's points of presence (POPs) and the POPs are inter-connected by a private WAN. Throughout this chapter of ***The Handbook***, the phrase *private WAN* will refer to WAN services other than the Internet. This includes WAN services such as private lines, MPLS, Frame Relay and ATM.

Ideally a solution of this type supports a wide variety of access services. In addition, the solution must have enough POPs so that there is a POP in close proximity to the users and to the applications and data the users want to access so as to not introduce unacceptable levels of delay.

The Optimization of Internet Traffic

WOCs were designed to address application performance issues at both the client and server endpoints. These solutions make the assumption that performance characteristics within the WAN are not capable of being 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 private 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. Throughout this section of the handbook, the class of WAN optimization service that has a focus on optimizing Internet traffic will be referred to as an Optimizing Internet Traffic Service (OITS).

An OITS leverages service provider resources that are distributed throughout the Internet. The way this works is that as shown in **Figure 4**, all client requests to the application's origin server in the data center are redirected via DNS to a server in a nearby point of presence (PoP) that is part of the OITS. This edge server then optimizes the traffic flow to the OITS server closest to the data center's origin server.

Figure 4: The Structure of an OITS

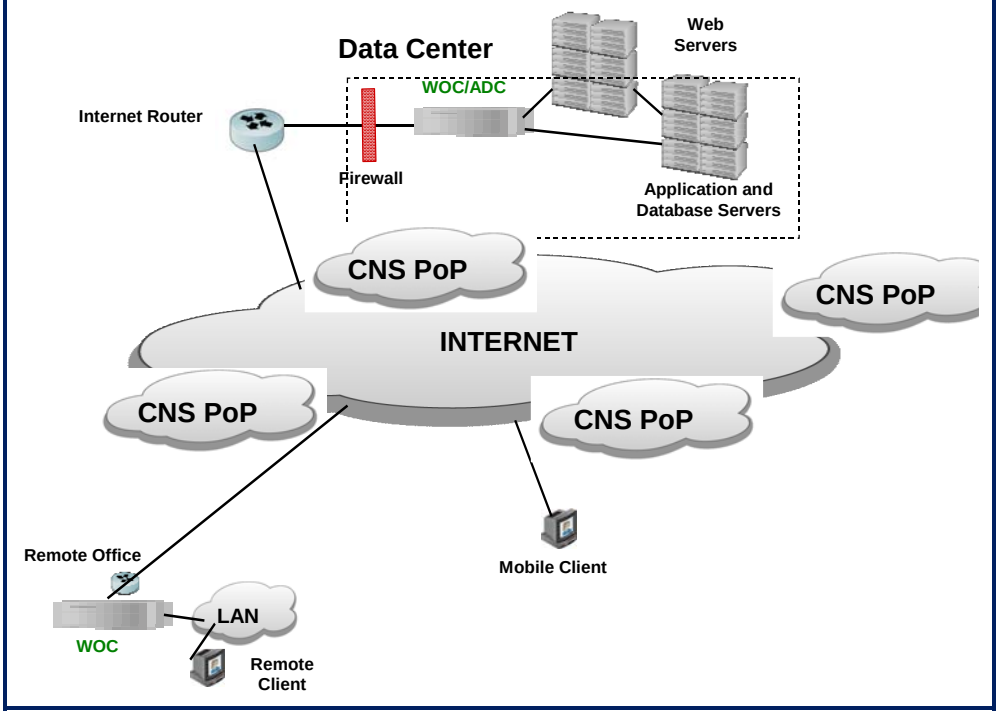

The servers at the OITS provider's PoPs perform a variety of optimization functions. Intelligence within the OITS servers can also be leveraged to provide extensive network monitoring, configuration control and SLA monitoring of a subscriber's application and can also be leveraged to provide security functionality. The management and security functionality that can be provided by an OITS will be discussed in more detail in the next chapter of the handbook.

Some of the optimization functionality provided by an OITS is similar to the functionality provided by a WOC. This includes optimizing the performance of protocols such as TCP and HTTP. Some of the unique optimization functionality that can be provided by an OITS includes choosing the optimal path through the Internet, offloading data out of data-centers to caches in OITS servers close to the users, and increasing availability by leveraging dynamic route optimization technology.

Hybrid WAN Optimization Solutions

Throughout this chapter of ***The Handbook***, the phrase *hybrid WAN* will refer to a network that is comprised of two or more WAN services such as MPLS and the Internet. As explained in the 2013 edition of ***The Handbook,*** having connections to multiple WAN services can enable IT organizations to add inexpensive WAN bandwidth and can dramatically improve the reliability and availability of the WAN.

A Basic Hybrid WAN

The traditional approach to providing Internet access to branch office employees has been to backhaul that Internet traffic on the organization's enterprise network (e.g., their MPLS network)

to a central site where the traffic was handed off to the Internet. The advantage of this approach is that it enables IT organizations to exert more control over their Internet traffic and it simplifies management in part because it centralizes the complexity of implementing and managing security policy. One disadvantage of this approach is that it results in extra traffic transiting the enterprise's WAN, which adds to the cost of the WAN. Another disadvantage of this approach is that it usually adds additional delay to the Internet traffic.

The 2012 and 2013 editions of ***The Handbook*** reported on the results of a survey in which the survey respondents were asked to indicate how they currently route their Internet traffic and how that is likely to change over the next year. The survey responses indicated that:

Although the vast majority of IT organizations currently have a centralized approach to Internet access, IT organizations are continually adopting a more decentralized approach.

One way that a hybrid WAN can eliminate the disadvantages of backhauling traffic is shown in **Figure 5.**

Figure 5: A Basic Hybrid WAN

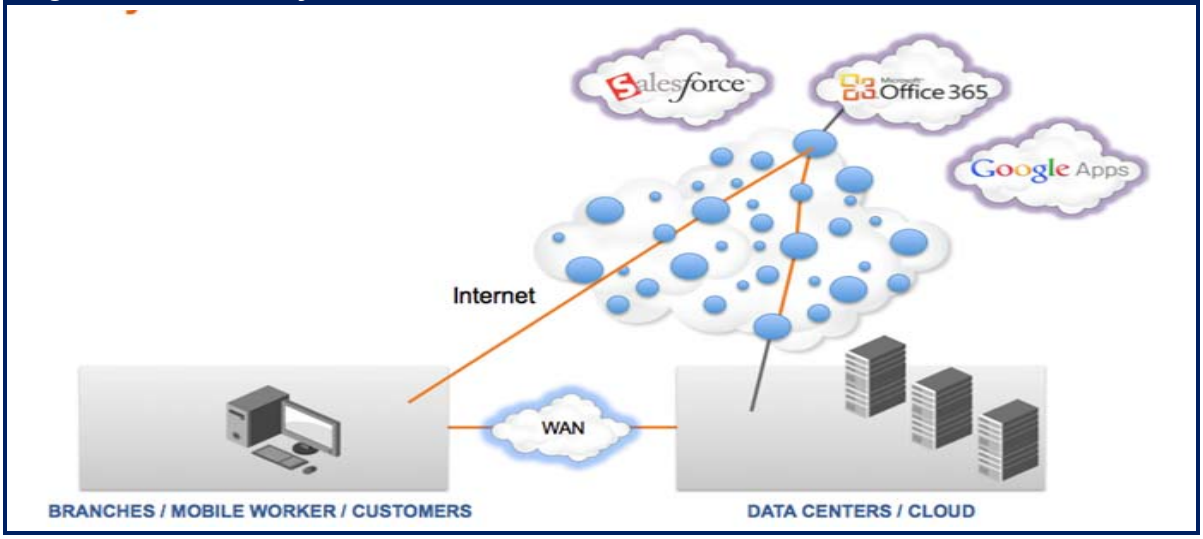

The network in **Figure 5** utilizes Policy Based Routing (PBR) in such a way that traffic destined to a public cloud provider transits the Internet while traffic that is destined for the corporate data center transits the MPLS WAN.

In order for an IT organization to feel comfortable implementing the network shown in **Figure 5**, the organization must find a way to implement the security and control that it has when it backhauls Internet traffic. One way this can be done is to replace the basic Internet connection shown in **Figure 5** with an OITS. The advantage of this is that in addition to providing optimization functionality, the OITS can provide the security functionality that was previously provided in the corporate data center.

The hybrid WAN that is described above is deemed to be a *basic hybrid WAN* service because it doesn't layer any additional intelligence over what is typically contained in the primary components of the service; e.g., a private WAN service such as MPLS; the basic Internet; or an OITS.

Intelligent Hybrid WANs

As documented in *The 2014 State of the WAN Report*[7], 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 that has more intelligence than the basic hybrid WAN described above. Looking just at the use of varying transmission services, 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 and 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 a basic hybrid WAN included the use of PBR to determine which traffic transited which WAN link. One of the concerns about the conventional way of implementing PBR is that the static nature of the PBR forwarding policies which results in the network not being able to respond in real time to changing network conditions. A relatively new class of device has emerged to address the shortcomings of PBR. WAN path controller (WPC) is one phrase that is often used to describe devices that work in conjunction with WAN routers to simplify PBR and to make 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.

One way to construct an intelligent hybrid WAN is to leverage WPC to apportion traffic over two WAN links where one WAN connection is a basic Internet connection and the other connection is MPLS. The added intelligence found in a WPC will improve the performance of the WAN and this WAN design alleviates at least some of the concerns about cost and uptime.

Another option is to still have one WAN connection be MPLS, but instead of using the basic Internet, have the other connection use an OITS. Because of the security functionality provided in an OITS, this approach should alleviate the previously mentioned security concerns. In addition, this approach results in improved performance due to the fact that as previously discussed, one of the ways that an OITS optimizes traffic is by offloading data out of data-centers to caches in OITS servers close to the users. It is possible, however, to further leverage the intelligence of an OTIS. For example, instead of offloading data out of data-centers to caches in OITS servers, it is possible to offload data out of data-centers to caches in the branch office and hence eliminate the round-trip delay on the access links.

[7] http://www.webtorials.com/news/2014/04/2014-wide-area-networking-state-of-the-market-report.html

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 at *jim@webtorials.com* or Steven Taylor at *taylor@webtorials.com*.

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

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

Professional Opinions Disclaimer

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

6WINDGate™

Packet Processing Software

Increase Data Plane Performance
No Change To Linux Environments
Available Across All Major Platforms
Support Extensive Set Of Protocols

L2-L4 Acceleration
IPsec VPN Gateways
TCP / UDP Termination
Virtual Switching
DPDK
And More ...

6WIND.com

SPEED MATTERS

ARE YOUR USERS WAITING MORE THAN DOING?

A10 helps more than 3,000 Enterprises and Service Providers deliver quality experiences

MAKE YOUR APPS RUN FASTER

Experience the Performance

www.a10networks.com

Enterprises today aspire to grow revenue by expanding globally and acquiring new customers, while also cutting costs and finding ways to become more agile. To realize their goals, every Enterprise has a core set of applications that they rely on to run their business operations.

Current Application Delivery Landscape

The user requirements for accessing and business applications is changing dramatically, and Enterprises must support more applications across a broader user base including customers, suppliers, partners, and employees. In order to leverage their applications to achieve their business goals, Enterprises must optimize the delivery of their applications to support fast, reliable, and secure access to ensure all users, both inside and outside of their organization, have the best possible experience.

In the past, Enterprises would resort to optimizing their application delivery using a physical hardware box or a virtual appliance that was deployed within a data center and any offices where users were located. While costly to deploy and manage, this approach did a good job of optimizing application delivery between the data center and branch office locations that were connected via a private network. Today, this approach is no longer effective due to several factors including:

- The complexity of having more users outside the organization's private network
- Applications distributed across multiple data centers and in the cloud
- End-users located all over the world using all sorts of different devices and networks, and
- A growing list of critical business applications such as CRM, collaboration, product lifecycle management, and support portals that users rely on every day.

It's not realistic for IT organizations to establish private network connections between all their users and all the data centers where their applications are hosted, or implement an application delivery box or virtual appliance in every data center, cloud environment, and every location where their end-users are located today.

In order to leverage their applications to achieve their business goals, organizations today cannot only rely exclusively on their private WAN to deliver their applications, but they must also leverage the ubiquity and scale of the Internet in order to embrace the trends of globalization and consumerization within their organizations.

Considering Akamai's Cloud-based Application Delivery Platform

Akamai's Terra Alta solution is a cloud-based Application Delivery Platform that enables Enterprises to leverage the Internet to deliver all their web-based applications in a fast, reliable, secure, and cost-effective way. Terra Alta is a managed service that empowers Enterprises to overcome the challenges related to delivering their applications over the Internet by placing all of the application delivery capabilities within Akamai's cloud-based Intelligent Platform, instead of requiring IT organizations to take on the burden of deploying and managing these critical capabilities on their own in the form of hardware boxes or virtual appliances. With Akamai, application optimizations are distributed globally across our Intelligent Platform, not constrained within the four walls of a few data centers, or restricted only to those users on a private network connection.

Akamai's Intelligent Platform is deployed on over 150,000+ servers which are embedded deeply into thousands of networks worldwide, which means we are very close to nearly all of the world's Internet users and datacenters. This means that users can benefit from fast, reliable, and secure business applications regardless of where they are located in the world! In addition to being a cloud-based platform, Akamai is device agnostic and does not require any application changes, which means it's quick and easy to implement and allows organizations to lower their IT costs and reduce complexity as compared to alternative application delivery optimization solutions. Akamai's unique cloud-based architecture also means that applications can be seamlessly migrated across data centers or cloud providers at will, and the application delivery optimizations will automatically move with the application. Terra Alta empowers Enterprises to embrace their cloud, mobile, and big data initiatives without the fear of increased costs or low application adoption.

Conclusion

By overcoming the new realm of global application delivery challenges, Akamai's cloud-based Application Delivery Platform empowers organizations to meet the demands of globalization and consumerization and instantly enter new markets, acquire new customers, improve customer interactions, do business via lower-cost online channels, enable end-users to get more done in less time, and achieve their goal of increasing revenue and reducing costs.

We make the Internet fast, reliable, and secure.

Akamai

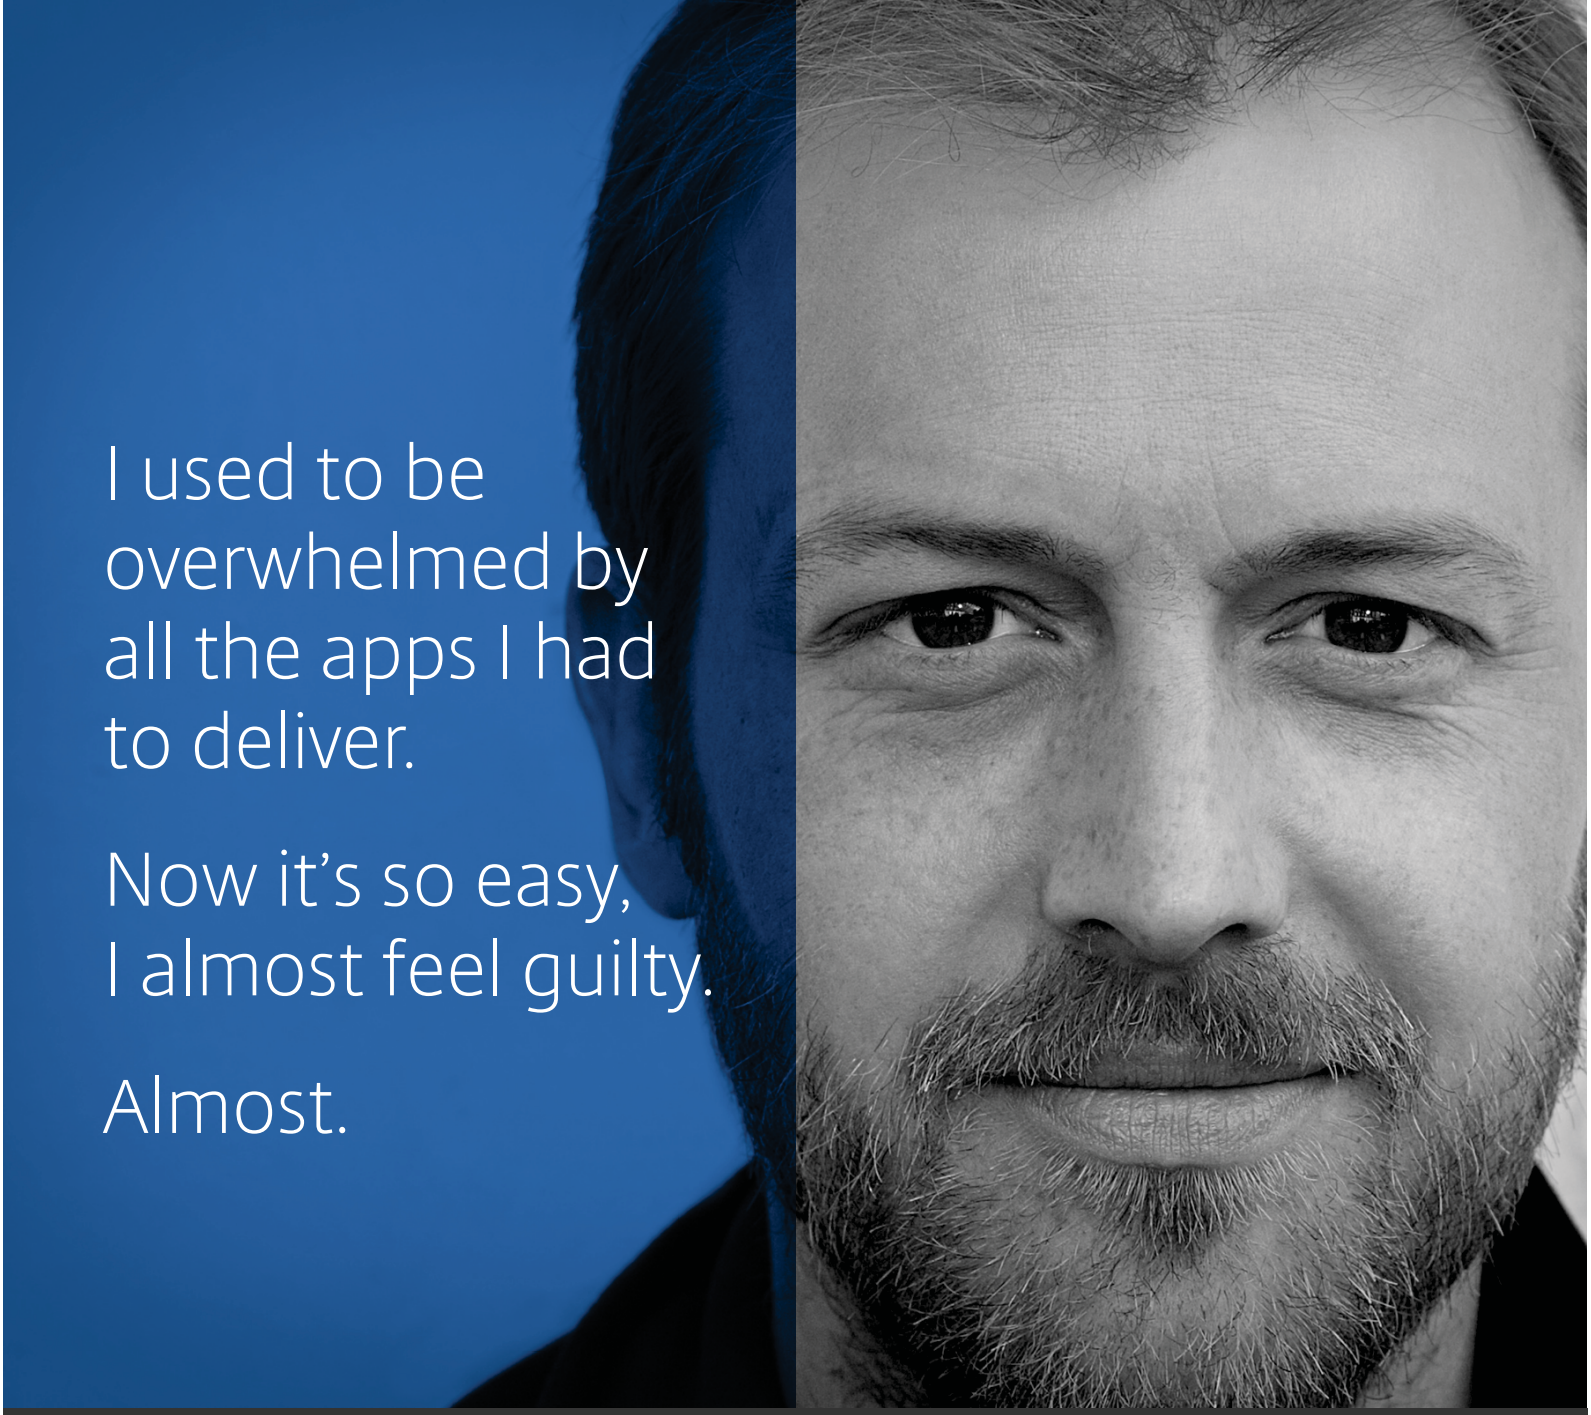

NetScaler with TriScale harnesses the power of software so you can effortlessly customize your app delivery for any business need.

NetScaler with TriScale
SOFTWARE SMART. HARDWARE STRONG.

CITRIX®

www.citrix.com/netscaler

CITRIX

Customer Spotlight

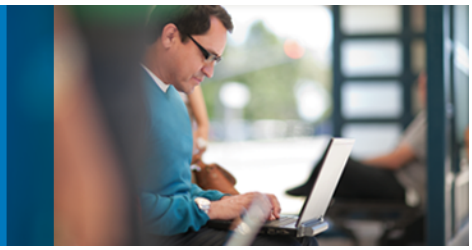

City Index

Industry:

Financial Services

Key Benefits:

- Reduces capital and operating expenses
- Provides an extended feature set on demand
- Ensures uninterrupted availability of trading platforms and applications

Citrix Products:

- Citrix NetScaler SDX
- Citrix NetScaler VPX

NetScaler reduces expenses with data center consolidation

City Index (www.cityindex.co.uk), is a leading global provider of retail trading services, including Spread Betting (UK only), Contracts for Difference (CFDs) and margined foreign exchange (FX). With offices in London (HQ), Warsaw, Tel Aviv, Singapore, Sydney and Shanghai, City Index supports its global clients on various trading platforms running on multiple backend applications. City Index is committed to providing a market-leading client service, transparent prices and innovative technology. All IT operations are centralized from two datacenters in London and one in Florida.

The challenge: need for a cost effective application delivery controller

Marc Morgan-Davies, Infrastructure Manager for City Index explained, our existing load balancing F5 solution that consisted of six BIG-IP 3400s and eight BIG-IP 6400s reached the end of support. There were two legacy Citrix Access Gateways (CAG) that we wanted to replace as well. We wanted a more cost effective and consolidated solution for our two London datacenters with an extended feature set that included global load balancing, DNS responder and rewrite, SSL offload, compression and caching.

Our choices were F5, Citrix, A10 Networks and Riverbed and we narrowed it down to F5 VIPRION 4400 and Citrix NetScaler 11500 SDX based on features and reputation. Our emphasis for the solution was more on the available feature set rather than raw processing power due to the nature of our platforms. The licensing model employed by Citrix is much simpler and more cost effective in my opinion than the competitors. For instance, if you want to enable a fourth module on the F5 you require another physical blade." Marc Morgan-Davies continued, "NetScaler on the other hand provided all of our required features on a single appliance with a simpler licensing model as well as allowing us to consolidate the existing CAGs onto the new devices further reducing our physical footprint and operating expenses. NetScaler gave us more features at a lower cost so was our chosen solution."

The solution: NetScaler SDX

Citrix NetScaler is an Application Delivery Controller (ADC) that optimizes the security, availability, scalability and performance of web-based applications and is available as a physical or virtual appliance. Citrix NetScaler

SDX is a true service delivery networking platform for enterprises and cloud datacenters. NetScaler SDX provides an advanced virtualized architecture that supports multiple NetScaler instances on a single hardware appliance, while an advanced control plane unifies provisioning, monitoring and management to meet the most demanding multi-tenant requirements**.**

NetScaler VPX is a software-based virtual appliance built for cloud scale. As an easy-to-deploy application delivery solution that runs on multiple virtualization platforms, the simplicity and flexibility of NetScaler VPX make it simple and cost-effective to fully optimize every web application and more effectively integrate networking services with application delivery. Performance capacities can be upgraded in production with the simple addition of a pay-as-you-grow license. NetScaler VPX helps organizations control costs by leveraging processing capacity already in place, including existing virtualized servers and associated resources.

"In October 2012, we installed 2 NetScaler SDXs as HA pair in production in each of our London datacenters. Each SDX box have 2 VPX instances that have discrete security layers In addition, we installed 2 SDXs as HA pair for staging in London with each SDX running 9 VPX instances. We are extremely pleased with NetScaler's ease of configuration and use." said Morgan Davies.

Key benefit: reduces capital and operating expenses

Using NetScaler we were able to prevent appliance sprawl by upgrading and consolidating 14 F5 Big-IP appliances and 2 Citrix Access Gateways to just 6 Citrix NetScaler appliances. This helped reduce support costs, rack space, ongoing power and cooling requirements drastically." Marc Morgan-Davies emphasized.

Key benefit: provides an extended feature set on demand

According to Marc Morgan-Davies, "Citrix NetScaler helped upgrade infrastructure while controlling costs. NetScaler provided the complete ADC feature set we required with the ability to enable features on demand."

Key benefit: ensures uninterrupted availability of trading platforms and applications

NetScaler ensured 24X7 availability of City Index's trading platforms and applications by providing global load balancing and SSL offloading between the London datacenters.

Looking ahead to the future

"NetScaler is a requirement to deploy Citrix XenMobile and MDM that would satisfy our requirement for a secure solution for users to access the corporate resources from any location using any device. Citrix mobility technologies are now very much on our scope for implementation in the near future. We are also looking into NetScaler App Firewall feature as well." Marc Morgan-Davies concluded.

About Citrix

Citrix (NASDAQ:CTXS) is the cloud company that enables mobile workstyles—empowering people to work and collaborate from anywhere, easily and securely. With market-leading solutions for mobility, desktop virtualization, cloud networking, cloud platforms, collaboration and data sharing, Citrix helps organizations achieve the speed and agility necessary to succeed in a mobile and dynamic world. Citrix products are in use at more than 330,000 organizations and by over 100 million users globally. Annual revenue in 2012 was $2.59 billion. Learn more at www.citrix.com.

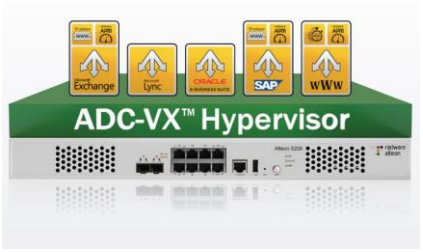

Predictable Application SLA, Guaranteed. Only with Alteon NG.

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 SLAs** and need tools to proactively monitor and manage application SLAs.

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 SLA Assurance

The Alteon® next-generation (NG) ADC solution is the industry's only ADC built from the ground up to ensure application SLAs 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 SLA** 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.
- ☑ Radware's **Application Performance Monitoring (APM)** module provides real-time tracking of application SLAs 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 is part of Radware's unique **Attack Mitigation System (AMS)**, which enables accurate detection and mitigation of the most advanced cyber-attacks. Leveraging a unique Defense Messaging™ mechanism, AMN efficiently mitigates attacks by signaling attack information to Radware DefensePipe cloud service and Radware DefensePro data center attack mitigator, located in the cloud or the network perimeter, respectively.
- ☑ Integrating advanced **Web Application Firewall (WAF)** capabilities, Alteon NG enables risk-free implementation thanks to a unique out-of-path WAF deployment mode along with auto-policy generation capabilities. Moreover, as 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. This results in secured web applications with SLA guarantee.
- ☑ 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.
- ☑ Alteon NG employs Radware's **AppShape™** offering configuration templates for leading business applications (e.g. Microsoft, Oracle, SAP). This helps customers roll out ADC-optimized applications in a simple, fast risk-free manner. In addition, Radware's AppShape++ scripting technology lets customers customize any ADC service per specific application flow/scenario. Using the AppShape++ script library, customers can refine various Layer 4-7 policies including HTTP, HTTPS, TCP, UDP, SSL and more – without application modifications to reduce cost and risk.

Complete Load Balancing/Layer 4-7 Feature Set

Alteon NG delivers a complete set of layer 4-7 services to ensure the availability, performance and security of mission-critical applications in the local and cloud data centers. These extend to traffic redirection, content modification, persistency, redundancy, advanced health monitoring and global server load balancing (GSLB). In addition, Alteon NG integrates advanced modules such as bandwidth management and link load balancing – reducing data center footprint and simplifying deployment. The combination of these advantages – along with an industry unique 5-year longevity guarantee, "pay-as-you-grow" approach in throughput, number of vADCs and services, plus performance leadership in all layer 4-7 metrics – makes Alteon simply your best application delivery choice.

Want to see more for yourself? We invite you to download our Radware ADC solution white paper here or contact us at: info@radware.com.