The 2014
Application & Service Delivery Handbook

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

Platinum Sponsors:

Gold Sponsors:

Produced by:

Table of Contents

Executive Summary 1
- Introduction 1
- Second Generation Application and Service Delivery Challenges 2
- Network and Application Optimization 5
- NFV Optimization 6
- Emerging WAN Optimization Services 7
- Management 9
- Application Performance Management 9
- DevOps 10
- Security 11

Second Generation Application and Service Delivery Challenges 13
- First Generation Application & Service Delivery Challenges 13
- Second Generation Application and Service Delivery Challenges 14
- Emerging Challenges 21

Network and Application Optimization 24
- Key Optimization Tasks 24
- Traditional Optimization Appliances 26
- NFV Optimization 31
- Emerging WAN Optimization Services 32

Management & Security 37
- Management 37
- DevOps 44
- Security 45

Conclusions 51

Executive Summary

Introduction

Throughout the ***2014 Application and Service Delivery Handbook*** (***The Handbook***), the phrase ***ensuring acceptable application and service delivery*** will refer to ensuring that the applications and services that an enterprise uses:

- Can be effectively managed;
- Exhibit acceptable performance;
- Incorporate appropriate levels of security;
- Are cost effective.

There is a growing relationship between the requirements listed above. For example, in order to implement an appropriate level of security, an IT organization may implement encryption. However, the fact that the information flow is encrypted may preclude the IT organization from implementing the optimization techniques that are required to ensure acceptable performance.

Starting around 2007, IT organizations began to implement the first generation of application delivery solutions in order to respond to the first generation of application delivery challenges, such as supporting chatty protocols. The first generation of application delivery solutions were typically deployed on-premise. Representative solutions included appliance-based WAN Optimization Controllers (WOCs), management solutions that focused narrowly on the network and myriad security appliances such as firewalls.

A second generation of challenges is described below. To respond to these new challenges a second generation of application delivery solutions is emerging. In many cases these solutions aren't appliance based, but are software based. In a growing number of instances they are provided as part of a managed service or acquired from a public cloud provider. The management component of this new generation of application delivery solutions is less likely to be focused narrowly just on the network and more likely to integrate network and application management.

The goal of the ***The Handbook*** is to help IT organizations ensure acceptable application and service delivery when faced with both the first generation, as well as the emerging second generation of application and service delivery challenges. To help to achieve this goal, in early 2014 a survey was given to the subscribers of Webtorials. Throughout this document, the IT professionals who responded to the surveys will be referred to as The Survey Respondents.

Second Generation Application and Service Delivery Challenges

There are a number of fairly well understood challenges that have over the years complicated the task of ensuring acceptable application and service delivery. Those challenges are described in detail in the document entitled Traditional Application & Service Delivery Challenges. In addition, there are a number of second-generation challenges that are beginning to complicate the task of ensuring acceptable application and service delivery. Those challenges include:

- Mobility and BYOD
- Virtualization
- Cloud Computing

Mobility and BYOD

Previous research has identified a number of key characteristics that are associated with mobility and BYOD including:

- The vast majority of employees require mobile access for at least part of their typical day.
- The BYOD movement has resulted in a loss of control and policy enforcement.

The Survey Respondents were asked how important it is for their IT organization over the next year to get better at improving the performance of applications used by mobile workers. They were also asked how important it is for their IT organization over the next year to get better at managing and monitoring the performance of applications used by mobile workers. Their responses are shown in **Table 1**.

Table 1: Importance of Getting Better Delivering Mobile Applications

	Improving the Performance	Managing and Monitoring
Extremely Important	22%	22%
Very Important	33%	33%
Moderately Important	29%	26%
Slightly Important	11%	15%
Not at all Important	6%	5%

Virtualization

Server & Desktop Virtualization

The vast majority of organizations have made at least some deployment of server virtualization and the deployment of server virtualization will increase over the next several years. Many of the same management tasks that must be performed in the traditional server environment need to be both extended into the virtualized environment and also integrated with the existing workflow and management processes. One example of this is that IT organizations must be able to automatically discover both the physical and the virtual environment and have an integrated view of both environments. This view of the virtual and physical server resources must stay current as VMs move from one host to another. The view must also be able to indicate the resources that are impacted in the case of fault or performance issues. Feedback from The Survey Respondents indicates that almost two

thirds of the IT organizations consider it to be either very or extremely important over the next year for them to get better performing management tasks such as troubleshooting on a per-VM basis.

Over the last couple of years desktop virtualization has begun to gain traction in the market. The growing adoption of desktop virtualization is reflected in the fact that well over half of The Survey Respondents indicated that getting better at optimizing the performance of virtualized desktops is either extremely or very important to their IT organization. That is a significant increase over the responses to the same question in 2013 and the responses in 2013 were a significant increase over the responses to that question in 2012.

Software Defined Data Center (SDDC)

As noted, IT organizations are making increasing use of varying forms of virtualization. SDDC is an emerging concept that is being advocated by a number of vendors. The two primary characteristics of a SDDC are virtualization and automation. In particular, in a SDDC, all of the infrastructure is virtualized and delivered as a service and the control of this datacenter is entirely automated by software. The document entitled The Promise and the Reality of a Software Defined Data Center contains a detailed discussion of SDDCs.

Cloud Computing

Both private and public cloud computing create significant challenges relative to ensuring acceptable application delivery. For example, in most instances the SLAs that are associated with public cloud computing services such as Amazon's Simple Storage System are weak and as such, it is reasonable to say that these services are delivered on a best effort basis.

In order to understand some of the concerns that IT organizations have with cloud computing, The Survey respondents were asked to indicate how important it was over the next year for their organization to get better a managing end-to-end in a private cloud environment. **Table 2** shows how The Survey Respondents answered this question in 2014 and also shows how a corresponding set of survey respondents answered this question in 2013.

Table 2: Importance of Getting Better at Managing Private Cloud: 2014 vs. 2013

	Managing Private Cloud 2014	Managing Private Cloud 2013
Extremely Important	21%	12%
Very Important	39%	30%
Moderately Important	28%	32%
Slightly Important	6%	14%
Not at all Important	6%	12%

Software Defined Networking

According to the Open Networking Foundation (ONF), “Software-Defined Networking (SDN) is an emerging architecture that is dynamic, manageable, cost-effective, and adaptable, making it ideal for the high-bandwidth, dynamic nature of today's applications. This architecture decouples the network control and forwarding functions enabling the network control to become directly programmable and the underlying infrastructure to be abstracted for applications and network services. The OpenFlow™ protocol is a foundational element for building SDN solutions.”

A detailed discussion of Software Defined Networking (SDN) can be found in The 2013 Guide to Network Virtualization and SDN.

Network Function Virtualization

NFV is being driven primarily by telecommunications service providers who feel that they can greatly simplify their operations and reduce expense if all network functions were available as virtual appliances that can be easily provisioned and integrated regardless of the vendor who provided the appliance or the hypervisor(s) on which it runs. In order to bring this vision to fruition, an Industry Specifications Group for Network Functions Virtualization (NFV ISG) was formed under the auspices of the European Telecommunications Standards Institute (ETSI). In October 2013, ETSI published a set of high level reference documents that are openly available on the ETSI website. One of those documents discussed a framework for conducting a NFV Proof of Concept (POC). ETSI currently has eighteen POCs underway.

Until recently, the conventional wisdom in the IT industry was that SDN and NFV were separate topics and didn’t need to be formally coordinated. That conventional wisdom changed in March 2014 when the ONF and ETSI announced the signing of a Memorandum of Understanding (MOU) detailing their intention to work together.

A detailed description of Network Function Virtualization (NFV) can be found at An NFV Reality Check.

Network and Application Optimization

Key Optimization Tasks

The Survey Respondents were asked about the importance of a range of optimization tasks. Their feedback indicates that:

- 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.

- 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.

WAN Optimization Controllers (WOCs)

When WOCs were first introduced in the mid-2000s, they were hardware-based appliances. 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.

In addition, 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
- Client software
- Integrated Hardware/Software Appliances

Feedback from The Survey Respondents indicates that 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.

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. ADCs also provide a wide range of other sophisticated functionality.

ADCs and Security

The majority of serious security attacks are to an organization's data center because that's where most of their applications and most of their data reside. Given that the most common deployment of ADCs has them placed in front of application servers in a data center, ADCs 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, a DNS application firewall and it should also support SSL offload.

IPv6 and ADCs

Some of the IPv6 functionality that ADCs can support include[1]:

- 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)
- 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

Virtual ADCs

There is a wide array of options for implementing virtual ADCs. These options include:

- General Purpose VM Support
- Network Appliance O/S Partitioning
- Network Appliance with OEM Hypervisor
- Network Appliance with Custom Hypervisor

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

While performance bottlenecks are not unique to a virtualized environment such as an NFV environment, as IT organizations adopt a virtualized environment the performance bottlenecks multiply. 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 a smooth transition to support 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.

[1] IPv6 Deployment Starts at Network Edge

Emerging WAN Optimization Services

Emerging Environment

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. While the traditional IT environment 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. 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, 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.

Cloud-Based, Private WAN Optimization Solutions

In a cloud-based, private WAN optimization solution 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. 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

Throughout ***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 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. 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.

Hybrid WAN Optimization Solutions

Throughout ***The Handbook*** the phrase *hybrid WAN* refers to a network that is comprised of two or more WAN services such as MPLS and the Internet.

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. One way that a hybrid WAN can eliminate the disadvantages of backhauling traffic is shown in **Figure 1**.

Figure 1: A Basic Hybrid WAN

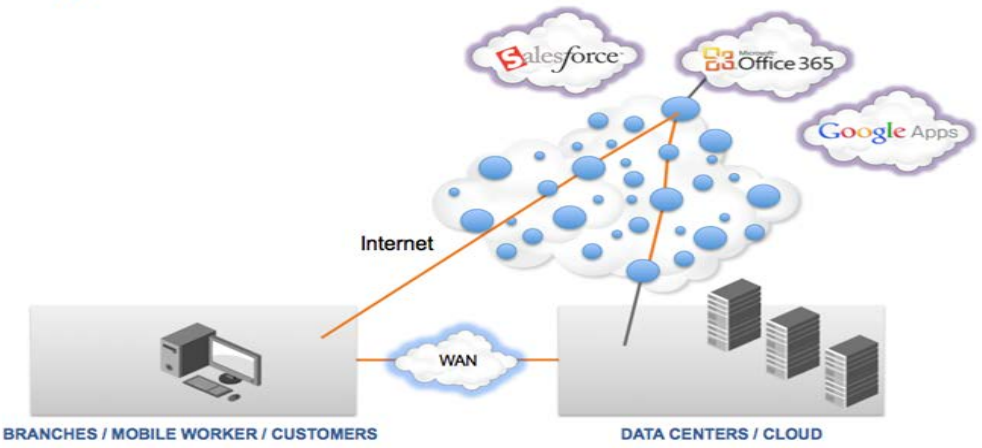

In order for an IT organization to feel comfortable implementing the network shown in **Figure 1**, the organization must find a way to implement the security and control that it has when it backhauls Internet traffic. One way this can be done is to replace the basic Internet connection shown in Figure 1 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

The preceding discussion of a basic hybrid WAN included the use of traditional Policy Based Routing (PBR) to determine which traffic transited which WAN link. One of the concerns with PBR is the static nature of the PBR forwarding policies. A relatively new class of device has emerged to address the shortcomings of PBR: A WAN path controller (WPC). A WPC works 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.

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. This approach also improves performance because an OITS optimizes traffic 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.

Management

Key Management Tasks

The Survey Respondents were asked about the importance of a range of management tasks. Their feedback indicates that:

- Two tasks are in a virtual tie for the most important management task to get better at over the next year: 1) Rapidly identifying the root cause of degraded application performance; 2) Identifying the components of the IT infrastructure that support the company's critical business applications.

- The second most important set of management tasks include: 1) Obtain performance indicator metrics and granular data that can be used to detect and eliminate impending problems; 2) Monitor the end user's experience and behavior; 3) Effectively manage SLAs for one or more business critical applications; 4) Manage the use of VoIP.

Impact of SDN

SDN creates some new management challenges. For example, one of the primary benefits of SDN is the ability to support multiple virtual networks that run on top of the physical network. Effective operations management, however, requires tools that give operators clear visibility into the relationships between the virtual and physical networks and their component devices. In addition, the SDN controller is in the data path for new flows entering the network. During periods when many new flows are being created, the controller can potentially become a performance bottleneck adding significant latency for flow initiation. Performance management systems need visibility not only into application performance but also controller performance in processing flows.

The document entitled *Mock RFI for Enterprise SDN Solutions* contains a number of questions that IT organizations should ask SDN vendors relative to SDN management.

Application Performance Management

A Framework

Since any component of a complex service such as Customer Relationship Management (CRM) can cause service degradation or a service outage, in order to effectively perform application performance management IT organizations need a single unified view of all of the components that support a service. This includes the highly visible service components such as servers, storage, switches and routers, in both their traditional stand-alone format as well as in their emerging converged format. It also includes the somewhat less visible network services such as DNS and DHCP, which are significant contributors to application degradation. Multiple organizational units within an IT organization have traditionally provided all of these service components. On an increasing basis, however, one or more network service providers and one or more cloud computing service providers will provide some or all of these service components. As a result, in order to achieve effective service delivery management, management data must be gathered from the enterprise, one or more Network Service Providers (NSPs) and one or more Cloud Computing Service Providers (CCSPs). In addition,

in order to help relate the IT function with the business functions, IT organizations need to be able to understand the key performance indicators (KPIs) for critical business processes such as CRM and relate these business level KPIs to the performance of the IT services that support the business processes.

Enterprise IT organizations can choose among several types of tools for monitoring and managing application performance over a private enterprise network. These include: application agents, monitoring of real and synthetic transactions, network flow and packet capture, analytics, and dashboard portals for the visualization of results.

At a high level, there are two basic classes of tools. The first class of tool monitors global parameters such as user response time or transaction completion time and provides alerts when thresholds are exceeded. These tools include agents on end user systems and monitoring appliances in the data center. The second class of tool supports triage by monitoring one or more of the components that make up the end-to-end path of the application. These tools include devices that capture application traffic at the flow and packet levels, agents on database, application, and web servers, as well as agents on various network elements.

Application Aware Network Performance Management

In response to the fact that enterprise networks and the applications that transit these networks are becoming increasingly entwined, there has been a movement to bring together two management disciplines: Application Performance Management and Network Performance Management. The result of bringing together those two disciplines is a new discipline that is often referred to as Application Aware Network Performance Management (AANPM). An AANPM solution integrates data that has historically been associated with application performance management with data that has historically been associated with network performance management. The result is a system that provides cross-platform visibility that enables IT organizations to monitor, troubleshoot and analyze both network and application systems.

DevOps

The phrase *DevOps* is a result of bringing to together two phrases: *Development* and *Operations*. That's appropriate because the point of adopting DevOps is to establish tight collaboration between a number of the phases of the application development lifecycle, including application development, testing, implementation and ongoing operations. DevOps is not a technology, but an approach. Some of the key characteristics of the approach are that the applications development team writes primarily small incremental pieces of code that are tested on an architecture that reflects the production architecture. Ideally, the network on which the software is tested will reflect not just the architecture but also the same characteristics (i.e., delay, packet loss) as the production network.

Implementing DevOps provides many advantages. For example, DevOps can provide business value by enabling companies to experience sustained innovation[2]. Examples of companies that claim to have experienced sustained innovation as a result of implementing DevOps include Twitter, Netflix and Facebook. Implementing DevOps has other advantages. According to a recent Information Week Report, eighty two percent of the IT organizations that implemented DevOps saw at least some improvement in infrastructure stability and eighty three percent saw at least some improvement in the speed of application development.

[2] Use DevOps to Turn IT into a Strategic Weapon

Security

The Changing Security Landscape

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, can rent botnets, 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.

IT security systems and policies have evolved and developed around the traditional application delivery architecture in which branch offices are connected using a private WAN service to application servers in a central corporate data centers. In this architecture, the central corporate data center is a natural location to implement IT security systems and policies that provide layered defenses as well a single, cost efficient location for a variety of IT security functions. With the adoption of public cloud computing, applications and services are moving out of the central corporate data center and there is no longer a convenient single location for security policies and systems.

In addition, IT security systems and policies have traditionally distinguished between people who were using IT services for work versus those who were using it for personal use. The use of an employer provided laptop was subject to the employer's IT security policies and systems. In this environment, the use that employees made of personal laptops was generally outside of the corporate IT security policy. With the arrival of smartphones and tablet computers, the ownership, operating systems and security capabilities of the end user devices have changed radically. IT security policies and standards that were developed for PCs are no longer effective nor optimal with these devices. Most corporations have embraced the BYOD movement and end users are less willing to accept strict corporate security policies on devices they own. Additionally, strict separation of work and personal usage for security on an employee owned device is impractical.

Current State of DDoS Attacks

A number of recent reports have highlighted the evolving threats that are associated with DDoS attacks on data centers. One of the key findings of those reports is that DDoS attacks are growing in a variety of ways, including:

- Frequency: A 50% increase in DDoS attacks on a year-over-year basis[3];
- Size: One attack out of three is over 20 Gbps, 60 Gbps attacks are common and 100 Gbps attacks are not uncommon[4];
- Persistence: The average duration of a DDoS attack is 17 hours[5].

[3] Akamai's State of the Internet 2014
[4] Neustar Annual DDoS Attacks and Impact Report
[5] Prolexic Global Attack Report Q1 2014

Web Application Firewall Services

Roughly twenty years ago IT organizations began to implement the first generation of network firewalls. These devices were placed at the perimeter of the organization with the hope that they would prevent malicious activities from causing harm to the organization.

Whereas network firewalls are focused on parameters such as IP address and port numbers, a more recent class of firewall, referred to as a Web application firewall, analyzes messages at layer 7 of the OSI model. Web application firewalls are typically deployed as a hardware appliance and they sit behind the network firewall and in front of the Web servers. In some cases, Web application firewall functionality is provided by an Application Delivery Controller (ADC).

Web application firewalls look for violations in the organization's established security policy. For example, the firewall may look for abnormal behavior or signs of a known attack. It may also be configured to block specified content, such as certain websites or attempts to exploit known security vulnerabilities. Because of their ability to perform deep packet inspection at layer 7 of the OSI model, a Web application firewall provides a level of security that cannot be provided by a network firewall.

In order to be effective, a Cloud Based Security Solution (CBSS) that provides Web application firewall functionality needs to be deployed as broadly as possible, preferably in thousands of locations. When responding to an attack, the service must also be able to:

- Block or redirect requests based on characteristics such as the originating geographic location and whether or not the originating IP addresses are on either a whitelist or a blacklist.
- Direct traffic away from specific servers or regions under attack.
- Issue slow responses to the machines conducting the attack. The goal of this technique, known as tarpits[6], is to shut down the attacking machines while minimizing the impact on legitimate users.
- Direct the attack traffic back to the requesting machine at the DNS or HTTP level.

Impact of SDN

As was the case with management, SDN poses both security challenges and security opportunities. The primary security challenge is to ensure that an attacker cannot compromise the central SDN controller and hence have access to all of the subtending network elements. In addition to securing the controller itself, all communication between the controller and other devices including switches, network services platforms and management systems must be secured.

As noted, SDN also presents opportunities to improve security by implementing security related applications that leverage the control information that has been centralized in the SDN controller. One such application that has been announced[7] is a network service that provides DDoS protection. Another such example is an application[8] that was designed to combat the security challenges that are associated with BYOD.

The document entitled Mock RFI for Enterprise SDN Solutions contains a number of questions that IT organizations should ask SDN vendors relative to security.

[6] Wikipedia Tarpit (networking)
[7] Radware: DefenseFlow
[8] HP Network Protector SDN Application

Second Generation Application and Service Delivery Challenges

First Generation Application & Service Delivery Challenges

There are a number of fairly well understood challenges that have over the years complicated the task of ensuring acceptable application and service delivery. Those challenges are listed below and are described in detail in the document entitled Traditional Application & Service Delivery Challenges.

- Limited Focus on Application Development
- Network Latency
- Availability
- Bandwidth Constraints
- Packet Loss
- Characteristics of TCP
- Chatty Protocols and Applications
- Myriad Application Types
- Webification of Applications
- Expanding Scope of Business Critical Applications
- Server Consolidation
- Data Center Consolidation
- Server Overload
- Distributed Employees
- Distributed Applications
- Complexity
- Increased Regulations
- Security Vulnerabilities

Second Generation Application and Service Delivery Challenges

There are a number of second generation challenges that complicate the task of ensuring acceptable application and service delivery. Some of these challenges result from the adoption of application architectures such as SOA. These application architectures tend to be more susceptible to performance problems due to WAN impairments than do traditional application architectures. In addition, the introduction of technologies such as AJAX creates significant security vulnerabilities[9]. Many of the second generation application and service delivery challenges, such as the ones mentioned in the preceding paragraph, are described in the 2012 Application and Service Delivery Handbook.

The 2014 Application and Service Delivery Handbook will focus on three key second generation challenges:

- Mobility and BYOD
- Virtualization
- Cloud Computing

One of the facts of life in IT is that when IT organizations implement new technologies or new ways of implementing technology, they tend to not completely retire the traditional approaches. For example, as IT organizations make increasing use of Infrastructure-as-a-Service (IaaS) and Software-as-a-Service (SaaS), they still continue to provide infrastructure services and host applications on-premise.

IT managers face the application delivery challenges associated with both the legacy environment and the emerging environment.

Mobility and BYOD

The 2013 edition of *The Handbook* quantified how often employees of a company access business related data and applications by using a mobile device both within a company facility as well as when they are at an external site. The data indicated that:

The vast majority of employees require mobile access for at least part of their typical day.

The 2013 edition of ***The Handbook*** also identified that the majority of IT organizations support a wide range of access devices including company and employee owned laptops and PCs as well as smartphone and tables from a wide range of vendors.

In the majority of instances, this new generation of employee-provided mobile devices doesn't run the Windows O/S and the existing security and management services for PCs must be extended for mobile devices or, alternatively, additional products and/or services added to perform these functions. Similar to PCs, smartphone and tablet computers are subject to malware and network intrusion attacks. On PCs, there are mature, robust products for malware protection (e.g. anti-virus software) and network intrusion protection (e.g., personal firewall), but these protections are just now emerging for smartphones and tablet computers. Similarly, inventorying and updating installed software on smartphone and tablet computers are emerging capabilities and a critical area for Mobile Device Management solutions.

[9] Ajax Security Issues

The BYOD movement has resulted in a loss of control and policy enforcement.

In addition, this new generation of mobile devices were architected and designed primarily for consumer use which is an environment in which the IT security risk is lower than it is in a corporate environment.

Adopting BYOD increases a company's vulnerability to security breaches.

Another key concern relative to supporting mobile workers is how the applications that these workers access have changed. At one time, mobile workers tended to primarily access either recreational applications or applications that are not delay sensitive; e.g., email. However, in the current environment mobile workers also need to access a wide range of business critical applications, many of which are delay sensitive. One of the challenges associated with supporting mobile workers' access to delay sensitive, business critical applications is that because of the way that TCP functions, even the small amount of packet loss that is often associated with wireless networks results in a dramatic reduction in throughput.

In order to quantify the concern amongst IT organizations about ensuring acceptable application and service delivery to mobile workers, The Survey Respondents were asked two questions. They were asked how important it is for their IT organization over the next year to get better at improving the performance of applications used by mobile workers. They were also asked how important it is for their IT organization over the next year to get better at managing and monitoring the performance of applications used by mobile workers. Their responses are shown in **Table 3**.

Table 3: Importance of Getting Better Delivering Mobile Applications

	Improving the Performance	Managing and Monitoring
Extremely Important	22%	22%
Very Important	33%	33%
Moderately Important	29%	26%
Slightly Important	11%	15%
Not at all Important	6%	5%

Getting better at managing and optimizing the delivery of mobile application is either very or extremely important to the majority of IT organizations.

Virtualization

Server Virtualization

The vast majority of organizations have made at least some deployment of server virtualization and the deployment of server virtualization will increase over the next several years. Many of the same management tasks that must be performed in the traditional server environment need to be both extended into the virtualized environment and also integrated with the existing workflow and management processes. One example of the need to extend functionality from the physical server environment into the virtual server environment is that IT organizations must be able to automatically discover both the physical and the virtual environment and have an integrated view of both environments. This view of the virtual and physical server resources must stay current as VMs move

from one host to another. The view must also be able to indicate the resources that are impacted in the case of fault or performance issues.

To quantify the impact that managing on a per-VM basis is having on IT organizations, The Survey Respondents were asked how important it is for their IT organization over the next year to get better at performing traditional management tasks such as troubleshooting and performance management on a per-VM basis. Their responses are shown in **Figure 2**.

Figure 2: Managing on a per-VM Basis

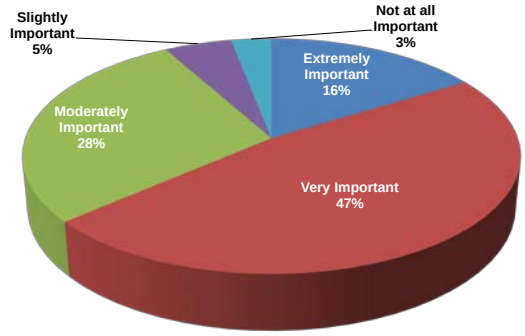

Almost two thirds of the IT organizations consider it to be either very or extremely important over the next year for them to get better performing management tasks such as troubleshooting on a per-VM basis.

Some of the other challenges created by server virtualization include:

- ***Multiple Hypervisors***
  It is becoming increasingly common to find IT organizations using multiple hypervisors, each with their own management system and with varying degrees of integration with other management systems. This creates islands of management within a data center.

- ***Layer 2 Network Support for VM Migration***
  When VMs are migrated, the network has to accommodate the constraints imposed by the VM migration utility. Typically the source and destination servers have to be on the same VLAN.

- ***Manual Network Reconfiguration to Support VM Migration***
  If the source and destination servers are not on the same VLAN, manual reconfiguration is required to adjust parameters such as QoS settings, ACLs, and firewall settings.

- ***Storage Support for Virtual Servers and VM Migration***
  The data storage location, including the boot device used by the VM, must be accessible by both the source and destination physical servers at all times. If the servers are at two distinct locations and the data is replicated at the second site, then the two data sets must be identical.

Desktop Virtualization

The 2013 edition of ***The Handbook*** discusses both the primary types of desktop virtualization and some of the key enabling protocols. As that edition of ***The Handbook*** demonstrates, desktop virtualization can provide significant benefits. However:

From a networking perspective, the primary challenge in implementing desktop virtualization is achieving adequate performance and an acceptable user experience for client-to-server connections over a WAN.

To quantify the concern that IT organizations have relative to supporting desktop virtualization, The Survey Respondents were asked how important it is for their IT organization over the next year to get better at optimizing the performance of virtualized desktops. Their responses are shown in **Figure 3**.

Well over half of The Survey Respondents indicated that getting better at optimizing the performance of virtualized desktops is either extremely or very important to their IT organization. That is a significant increase over the responses to the same question in 2013 and the responses in 2013 were a significant increase over the responses to that question in 2012.

Figure 3: Optimizing the Performance of Virtualized Desktops

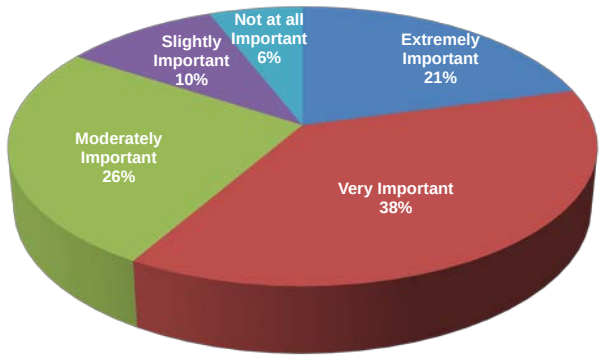

Getting better at optimizing the performance of virtualized desktops is becoming increasingly more important.

Software Defined Data Center (SDDC)

As noted, IT organizations are making increasing use of varying forms of virtualization. SDDC is an emerging concept that is being advocated by a number of vendors. The two primary characteristics of a SDDC are virtualization and automation. In particular, in a SDDC, the entire infrastructure is virtualized and delivered as a service and the control of this datacenter is entirely automated by software. Some vendors also advocate that the software run on commodity hardware. The document entitled The Promise and the Reality of a Software Defined Data Center[10] contains a detailed discussion of SDDCs. As described in that document, while it's true that few if any IT organizations have currently implemented an SDDC, it's also true that the steps that the majority of IT organizations have already taken to implement virtualization and automation are key steps on the path to implementing an SDDC.

[10] Ashton, Metler & Associates and click on *Journey to a new IT Operational Model*

Cloud Computing

The 2013 edition of ***The Handbook*** contains a detailed description of cloud computing, including the primary classes of solutions; the most important characteristics of cloud-based solutions; the drivers and inhibitors; the current and planned adoption; and the decision process relative to adopting public cloud solutions.

Over the last few years IT organizations have made a significant adoption of cloud computing in large part because:

The goal of cloud computing is to enable IT organizations to achieve a dramatic improvement in the cost effective, elastic provisioning of IT services.

In most instances the SLAs that are associated with public cloud computing services such as Salesforce.com or Amazon's Simple Storage System are weak and, as such, it is reasonable to say that these services are delivered on a best effort basis. For example, the SLA[11] that Amazon offers for its Amazon Web Services (AWS) states that, "AWS will use commercially reasonable efforts to make Amazon EC2 available with an Annual Uptime Percentage of at least 99.95% during the Service Year." As part of the Amazon definition of Annual Uptime Percentage, Amazon excludes any outage of 5 minutes or less. The Amazon SLA also states that if their service doesn't meet the Annual Uptime Percentage commitment, the customer will receive 10% off its bill for the most recent month that the customer included in the SLA claim that it filed.

A key attribute of the vast majority of the SLAs that are associated with public cloud computing services is that they don't contain a goal for the end-to-end performance[12] of the service. The reason for the lack of performance guarantees stems from the way that most public cloud computing services are delivered. As shown in **Figure 4**, one approach to providing public cloud computing services is based on the service being delivered to the customer directly from an independent software vendor's (ISV's) data center via the Internet. Another approach is for an ISV to leverage an IaaS provider such as Amazon to host their application on the Internet. Both of these approaches rely on the Internet and it is not possible to provide end-to-end quality of service (QoS) over the Internet. As a result, neither of these two approaches lends itself to providing an SLA that includes a meaningful commitment to critical network performance metrics such as delay, jitter and packet loss.

[11] Amazon EC2 SLA

[12] In this context, *performance* refers to metrics such as delay or response time.

The fact that cloud computing service providers (CCSPs) don't provide an end-to-end performance SLA for applications delivered over the Internet will not change in the foreseeable future. However, as will be described in the optimization section of ***The Handbook***, there are things that can be done to improve the performance of applications delivered over the Internet.

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

Figure 4: Distribution Models for Cloud-Based Solutions

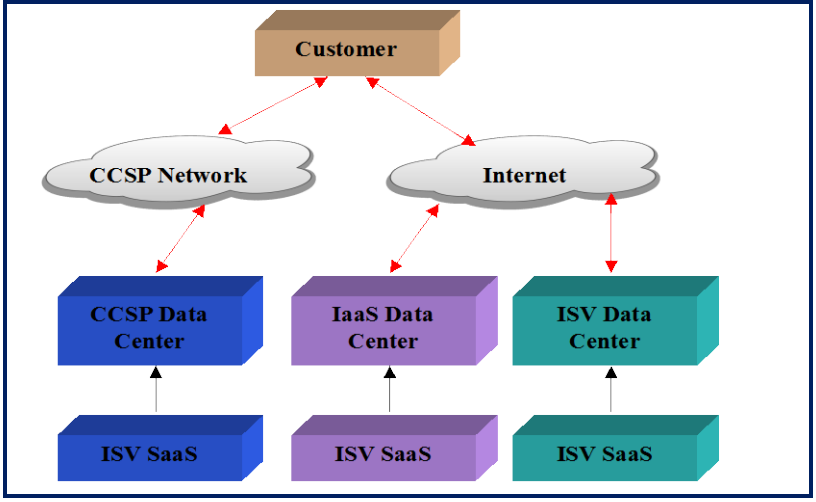

Many of the application delivery challenges associated with server virtualization also apply to the use of private cloud computing. In contrast, many of the application delivery challenges associated with the use of public cloud computing stem from the fact that IT organizations often have less visibility and control over the resources that comprise the cloud-based applications and services. This makes it difficult to manage, secure and optimize those resources.

The Survey Respondents were asked to indicate how important it was over the next year for their organization to get better a managing end-to-end in a private cloud environment. **Table 4** shows how The Survey Respondents answered this question in 2014 and also shows how a corresponding set of survey respondents answered this question in 2013.

Table 4: Importance of Getting Better at Managing Private Cloud: 2014 vs. 2013

	Managing Private Cloud - 2014	Managing Private Cloud - 2013
Extremely Important	21%	12%
Very Important	39%	30%
Moderately Important	28%	32%
Slightly Important	6%	14%
Not at all Important	6%	12%

The Survey Respondents were also asked to indicate how important it was over the next year for their organization to get better a managing end-to-end in a public cloud environment. **Table 5** shows how The Survey Respondents answered this question in 2014 and also shows how a corresponding set of survey respondents answered this question in 2013.

Table 5: Importance of Getting Better at Managing Public Cloud: 2014 vs. 2013		
	Managing Public Cloud 2014	**Managing Public Cloud 2013**
Extremely Important	21%	11%
Very Important	29%	26%
Moderately Important	25%	28%
Slightly Important	16%	14%
Not at all Important	9%	21%

Two key conclusions can be drawn from the data in **Table 4** and **Table 5**:

Managing end-to-end in both a private and public cloud environment has become notably more important to IT organizations over the last year.

Managing end-to-end in a private cloud environment is slightly more important to IT organizations than is managing end-to-end in a public cloud environment.

Emerging Challenges

Software Defined Networking

A detailed discussion of Software Defined Networking (SDN) can be found in The 2013 Guide to Network Virtualization and SDN.

The ONF is the organization that is most associated with the development and standardization of SDN. According to the ONF, "Software-Defined Networking (SDN) is an emerging architecture that is dynamic, manageable, cost-effective, and adaptable, making it ideal for the high-bandwidth, dynamic nature of today's applications. This architecture decouples the network control and forwarding functions enabling the network control to become directly programmable and the underlying infrastructure to be abstracted for applications and network services. The OpenFlow™ protocol is a foundational element for building SDN solutions."

Figure 5 contains a graphical representation of the SDN solution architecture as envisioned by the ONF.

Figure 5: ONF's SDN Architecture

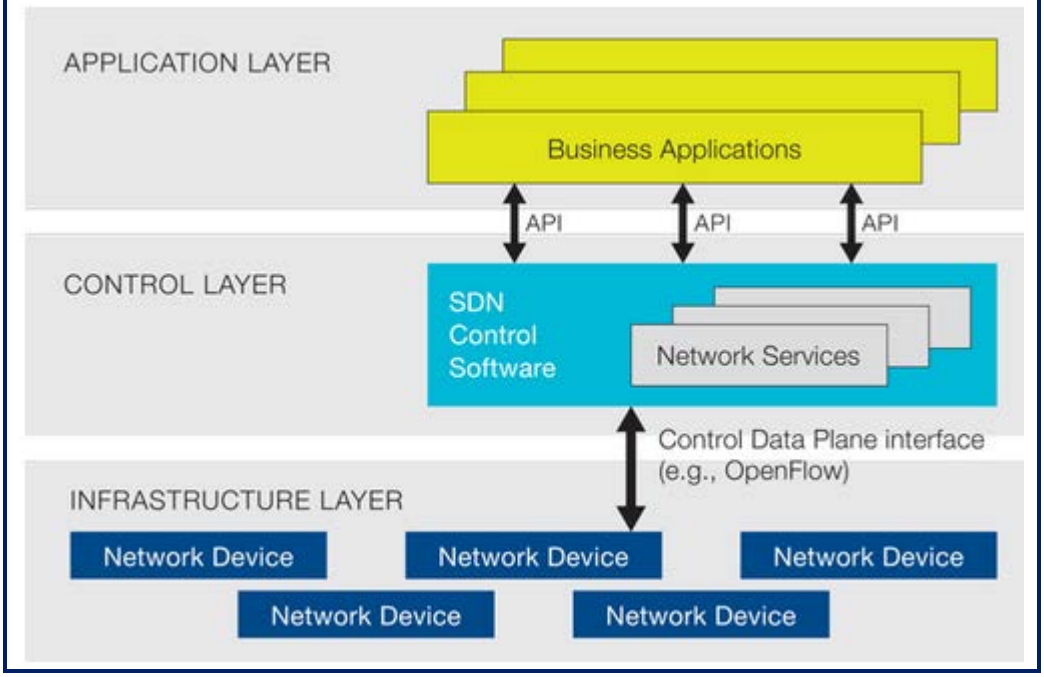

While the use of SDN in data centers receives the majority of attention in the trade press, it is also possible to implement SDN in branch and campus networks as well as in wide area networks (WANs).

The OpenFlow Protocol

OpenFlow is an example of a protocol that runs between the SDN control layer and the SDN infrastructure layer and which can be used to program the forwarding behavior of the switch. With OpenFlow, a single central controller, or cluster of controllers, can program all the physical and virtual switches in the network.

A high level description of what OpenFlow does is that when a packet arrives at an OpenFlow switch, the header fields are compared to flow table entries. One option for when a match is found, is that the packet is forwarded to specified port(s) or dropped depending on the action stored in the flow table. When an OpenFlow Switch receives a packet that doesn't match the flow table entries, it encapsulates the packet and sends it to the controller. The controller then decides how the packet should be handled and notifies the switch to either drop the packet or make a new entry in the flow table to support the new flow.

The most current version of OpenFlow that vendors have begun to support is OpenFlow 1.3. A complete, detailed description of the functionality provided by OpenFlow 1.3 can be found at the ONF Web site[13].

Network Function Virtualization

A detailed description of Network Function Virtualization (NFV) can be found at An NFV Reality Check.

Figure 6: The Virtualization of Network Appliances *Source: NFV ISG*

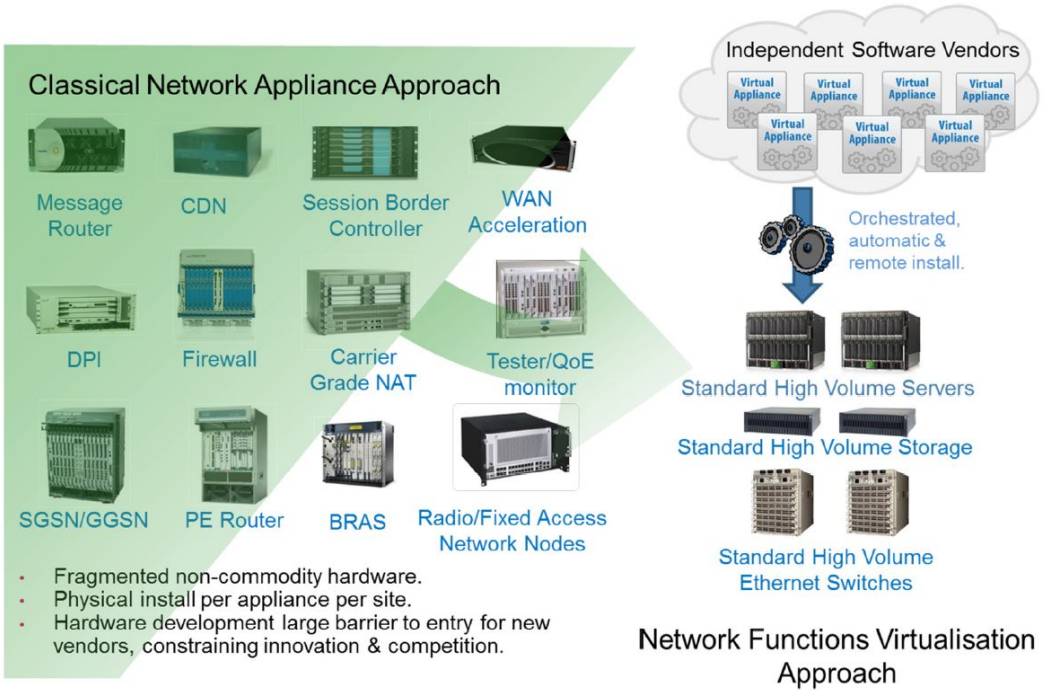

[13] Open Flow Switch Specification

NFV is being driven primarily by telecommunications service providers to meet their specific requirements.

Telecommunications service providers feel that they can greatly simplify their operations and reduce expense if all network functions were available as virtual appliances that can be easily provisioned and integrated regardless of the vendor who provided the appliance or the hypervisor(s) on which it runs. In order to bring this vision to fruition, an Industry Specifications Group for Network Functions Virtualization (NFV ISG) was formed under the auspices of the European Telecommunications Standards Institute (ETSI). Their vision for the transition from hardware appliances of today to a fully virtualized appliance environment is depicted in **Figure 6**.

The approach that the NFV ISG is taking is that the virtualization of network functionality is applicable to any data plane packet processing and control plane function in both fixed and mobile networks. In October 2013, ETSI published a set of high level reference documents that are openly available on the ETSI website[14]. One of those documents discussed a framework for conducting a NFV Proof of Concept (POC). ETSI currently has eighteen POCs underway.

Until recently, the conventional wisdom in the IT industry in general, and on the part of the ONF and ETSI in particular, was that was that SDN and NFV were separate topics and didn't need to be formally coordinated. That conventional wisdom changed in March 2014 when the ONF and ETSI announced the signing of a Memorandum of Understanding (MOU). As part of the announcing the MOU, the ONF and ETSI said that "Together the organizations will explore the application of SDN configuration and control protocols as the base for the network infrastructure supporting NFV, and conversely the possibilities that NFV opens for virtualizing the forwarding plane functions."

Some of the challenges associated with NFV include:

- Carrier-grade scalability and robustness.
- Real-time and dynamic provisioning. The virtual network functions must be automatically deployed and managed in the NFV infrastructure.
- Seamless control and provisioning of physical and virtual networking infrastructures.

[14] ETSI Network Functions Virtualisation

Network and Application Optimization

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 6** 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 6: 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%
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%

Table 6: The Importance of Key Optimization Tasks *(continued)*					
	Not at All	**Slightly**	**Moderately**	**Very**	**Extremely**
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 6** 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*** contains an extensive discussion of some of the key optimization challenges, such as the challenge associated with moving storage between data centers.

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 7**.

Table 7: 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 7** 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. 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.[15] 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.

ADCs can have the following IPv6 capabilities:

- 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)

[15] Preparing for IPv6: Verisign's Perspective

- 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 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.

Performance bottlenecks are not unique to virtualized environments. For example, some of the bottlenecks that occur in a physical environment are shown in **Figure 7**.

Figure 7: Representative Bottlenecks in a Physical Environment

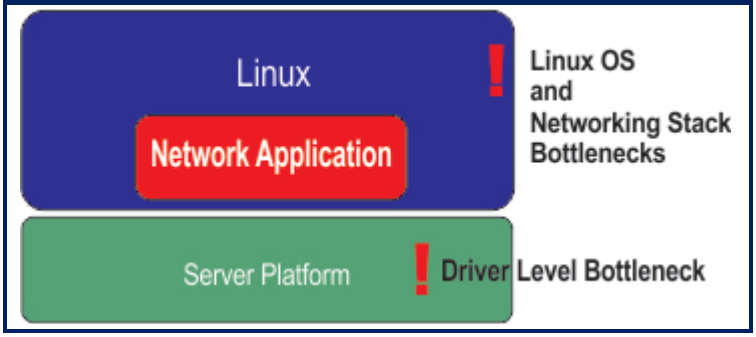

Figure 8: Performance Bottlenecks in a Virtualized Environment

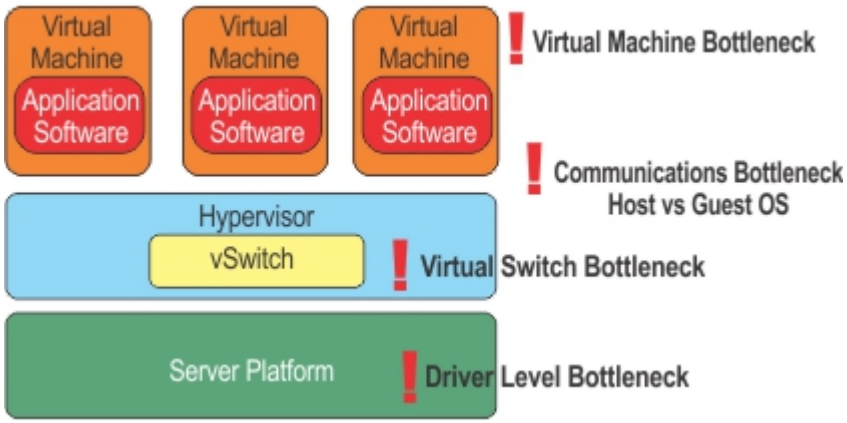

Unfortunately, as shown in **Figure 8**, as IT organizations adopt a virtualized environment the performance bottlenecks multiply. **Figure 8** demonstrates some, but not all of the bottlenecks that can occur in a virtualized environment. For example, while not explicitly shown in **Figure 8**, VM to 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 9**, 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 9: Breadth of Environments

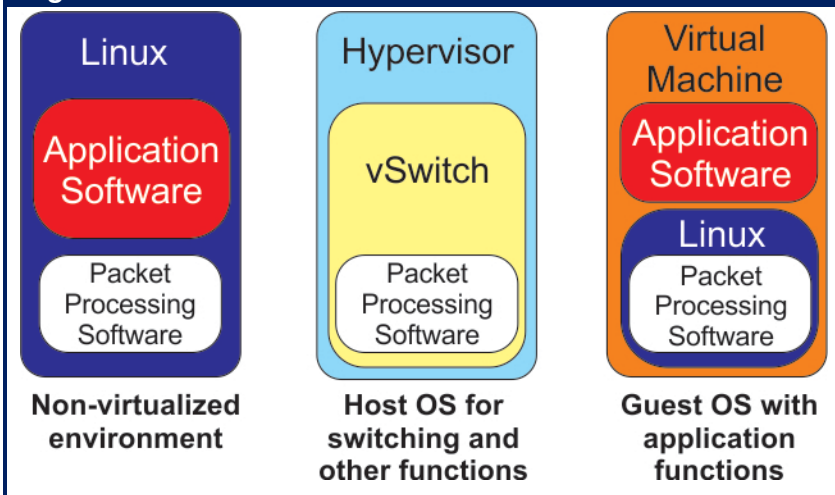

The evaluation criteria listed above are intended to ensure that the packet processing software can be easily and universally implemented on any version of Linux or on any hypervisor, without requiring changes to existing environments.

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, 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.

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 10**, 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 10: The Structure of an OITS

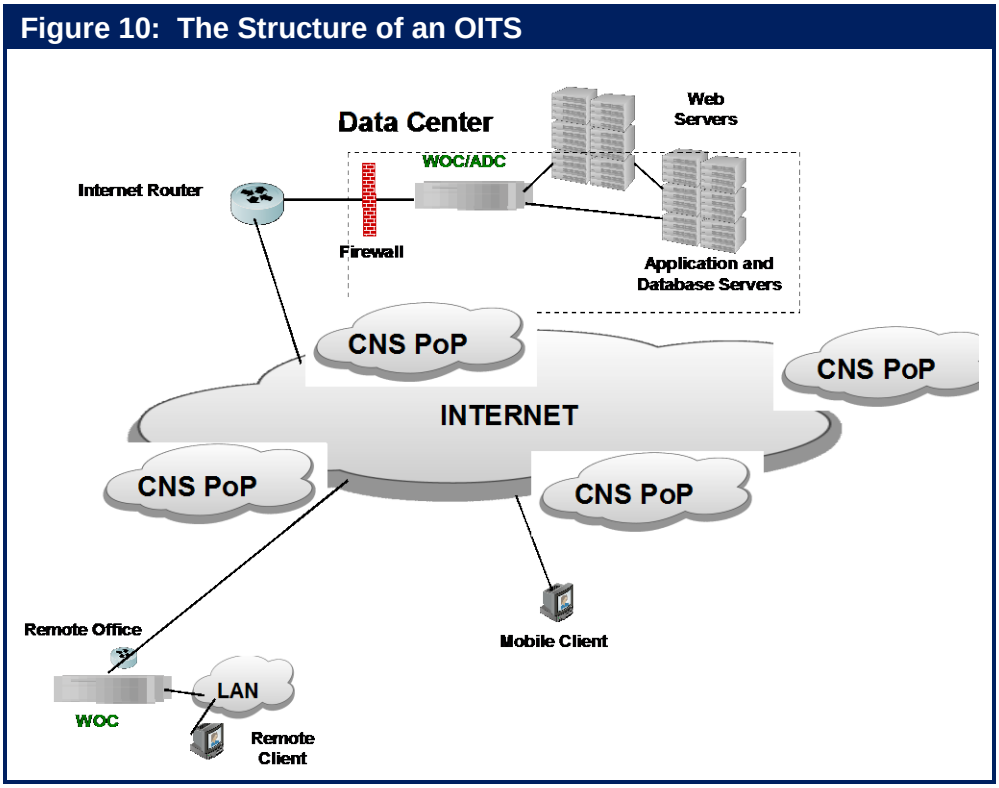

The servers at the OITS provider's PoPs perform a variety of optimization functions. Intelligence within the OITS servers can also be leveraged to provide extensive network monitoring, configuration control and SLA monitoring of a subscriber's application and can also be leveraged to provide security functionality. The management and security functionality that can be provided by an OITS will be discussed in more detail in the next chapter of t ***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 11.**

Figure 11: A Basic Hybrid WAN

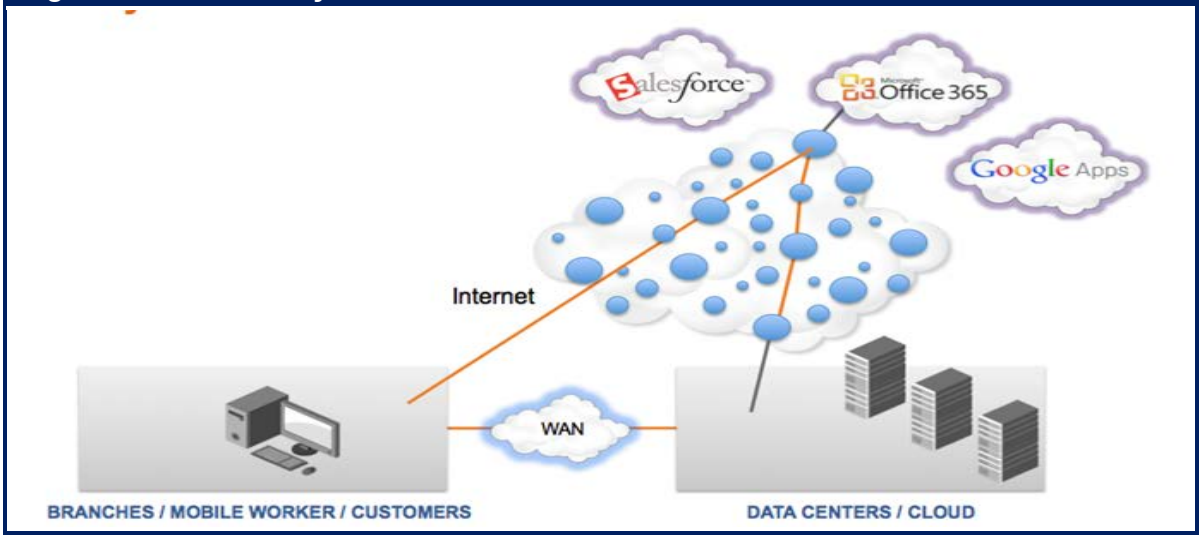

The network in **Figure 11** utilizes Policy Based Routing (PBR) in such a way that traffic destined to a public cloud provider transits the Internet while traffic that is destined for the corporate data center transits the MPLS WAN.

In order for an IT organization to feel comfortable implementing the network shown in **Figure 11**, 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 11** 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, 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.

Management & Security

Management

Market Research

The previous chapters 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 8** shows how The Survey Respondents answered the survey question about the management tasks that their IT organizations are most interested in getting better at over the next year.

Table 8: The Importance of Getting Better at Key Management Tasks

	Not at All	Slightly	Moderately	Very	Extremely
Rapidly identify the root cause of degraded application performance	0.0%	4.8%	11.5%	39.4%	44.2%
Identify the components of the IT infrastructure that support the company's critical business applications	1.2%	4.3%	10.4%	44.8%	39.3%
Obtain performance indicator metrics and granular data that can be used to detect and eliminate impending problems	1.2%	7.8%	16.8%	47.3%	26.9%
Monitor the end user's experience and behavior	0.0%	6.6%	21.0%	42.5%	29.9%
Effectively manage SLAs for one or more business critical applications	1.2%	5.5%	18.8%	41.2%	33.3%
Manage the use of VoIP	5.6%	7.4%	24.7%	30.9%	31.5%
Perform traditional management tasks such as troubleshooting and performance management, on a per VM basis	3.0%	5.5%	27.9%	47.3%	16.4%
Monitor and manage the performance of applications delivered to mobile users	4.9%	14.7%	25.8%	32.5%	22.1%
Manage end-to-end in a public cloud computing environment	9.0%	16.0%	25.0%	29.5%	20.5%

Table 8: The Importance of Getting Better at Key Management Tasks					
Manage end-to-end in a private cloud computing environment	6.3%	6.3%	27.5%	38.8%	21.3%
Effectively monitor and manage an application acquired from a SaaS provider such as Salesforce	13.7%	7.2%	29.4%	33.3%	16.3%
Manage end-to-end in a public cloud computing environment	9.0%	16.0%	25.0%	29.5%	20.5%
Effectively monitor and manage computing services acquired from a IaaS provider such as Rackspace	12.8%	16.2%	32.4%	23.6%	14.9%

Some of the conclusions that can be drawn from the data in **Table 8** include:

Two tasks are in a virtual tie for the most important management task to get better at over the next year: 1) Rapidly identifying the root cause of degraded application performance; 2) Identifying the components of the IT infrastructure that support the company's critical business applications.

The second most important set of management tasks include: 1) Obtain performance indicator metrics and granular data that can be used to detect and eliminate impending problems; 2) Monitor the end user's experience and behavior; 3) Effectively manage SLAs for one or more business critical applications; 4) Manage the use of VoIP.

While managing the use of services acquired from an IaaS provider such as Rackspace is relatively unimportant, it is more important than it was last year.

Forces Driving Change

Previous sections of this handbook described the traditional and emerging service and application delivery challenges. This subsection will identify how some of those challenges are forcing a change in terms of how IT organizations manage applications and services.

Server Virtualization

Until recently, IT management was based on the assumption that the IT organizations performed tasks such as monitoring, baselining and troubleshooting on a server-by-server basis. Now, as highlighted by the data in **Table 8**, IT organizations understand that they must also perform management tasks on a virtual machine (VM)-by-VM basis. Another assumption that underpinned the traditional approach to IT management was that the data center environment was static. For example, it was commonly assumed that an application resided on a given server, or set of servers, for very long periods of time. However, part of the value proposition that is associated with server virtualization is that it is possible to migrate VMs between physical servers, both within the same data center and between disparate data centers.

IT organizations need to adopt an approach to management that is based on the assumption that the components of a service, and the location of those components, can and will change frequently.

Cloud Balancing

IT management has historically been based on the assumption that users of an application accessed that application in one of the enterprise's data centers and that the location of that data center changed very infrequently over time. The adoption of IaaS solutions in general, and the adoption of cloud balancing in particular, demonstrates why IT organizations need to adopt an approach to IT management that is based on gathering management data across myriad data centers, including ones that are owned and operated by a third party. The adoption of cloud balancing is also another example of why IT organizations need to adopt an approach to management that is based on the assumption that the components of a service, and the location of those components, can and will change frequently.

Delay Sensitive Traffic

Voice and video are examples of applications that have high visibility and which are very sensitive to transmission impairments. As was also highlighted in **Table 8**, getting better at managing VoIP is one of the most important management tasks facing IT organizations.

As part of the traditional approach to IT management, it is common practice to use network performance measurements such as delay, jitter and packet loss as a surrogate for the performance of applications and services. A more effective approach is to focus on aspects of the communications that are more closely aligned with ensuring acceptable application and service delivery. This includes looking at the application payload and measuring the quality of the voice and video communications. In the case of Unified Communications (UC), it also means monitoring the signaling between the components of the UC solution.

In addition to having a single set of tools and more of a focus on application payload, IT organizations need to implement management processes that understand the impact that each application is having on the other applications and that can:

- Analyze voice, video, UC and data applications in consort with the network;
- Support multi-vendor environments;
- Support multiple locations.

Converged Infrastructure

One of the characteristics that is frequently associated with cloud computing is the integration of networking, servers and computing in the data center. While a converged data center infrastructure offers a number of benefits, it does create a number of management challenges. In particular, the converged infrastructure requires a management system and management processes that have the same level of integration and cross-domain convergence that the infrastructure has. For example, in order to support the requirement for the dynamic provisioning and re-allocation of resources to support a given IT service, the traditional manual processes for synchronizing the required server, network and storage resources will have to be replaced with integrated, automated processes. In order to enable

this change, the provisioning and change management processes will need to be integrated and will need to feature the automatic configuration of network and storage resources when additional infrastructure services are deployed or when additional physical or virtual servers are brought on line or are moved. In a similar fashion, operations management needs to be consolidated and automated to keep service quality in line with user expectations.

Impact of SDN

SDN management is a combination of good news and bad news. The good news is that SDN has the potential to make network management easier. For example, in theory at least, SDN enables IT organizations to centralize configuration and policy management.

The bad news is that SDN creates some new management challenges. For example, one of the primary benefits of SDN is the ability to support multiple virtual networks that run on top of the physical network. Effective operations management, however, requires tools that give operators clear visibility into the relationships between the virtual and physical networks and their component devices. In particular, when performance or availability problems occur, both root cause analysis and impact analysis require bilateral mapping between the physical and virtual infrastructures. With SDN, the flows between a pair of VMs can be distributed among a number of alternate paths through the network. Mapping a flow to the physical path it takes can be a challenge unless the flow monitoring solution can involve the controller's end-to-end view of the network

With SDN solutions, the controller is in the data path for new flows entering the network. During periods when many new flows are being created, the controller can potentially become a performance bottleneck adding significant latency for flow initiation. Performance management systems need visibility not only into application performance but also controller performance in processing flows. One of the characteristics of SDN is that network functions such as load balancing and firewalls are increasingly implemented in software as network services that can be integrated with virtual networks or SDN flow under programmatic control; a.k.a., service chaining. Implementing these functions in software both increases the delay associated with performing these functions and it also increases the variability of that delay. The result is an increased need for insight into the performance of each component of the overall SDN solution.

The document entitled Mock RFI for Enterprise SDN Solutions contains a number of questions that IT organizations should ask SDN vendors relative to SDN management.

Application Performance Management

Background

Since any component of a complex service such as Customer Relationship Management (CRM) can cause service degradation or a service outage, in order to effectively perform application performance management IT organizations need a single unified view of all of the components that support a service. This includes the highly visible service components such as servers, storage, switches and routers, in both their traditional stand-alone format as well as in their emerging converged format. It also includes the somewhat less visible network services such as DNS and DHCP, which are significant contributors to application degradation. Multiple organizational units within an IT organization have traditionally provided all of these service components. On an increasing basis, however, one or more network service providers and one or more cloud computing service providers will provide some or all of these service components. As a result, in order to achieve effective service

delivery management, management data must be gathered from the enterprise, one or more Network Service Providers (NSPs) and one or more Cloud Computing Service Providers (CCSPs). In addition, in order to help relate the IT function with the business functions, IT organizations need to be able to understand the key performance indicators (KPIs) for critical business processes such as CRM and relate these business level KPIs to the performance of the IT services that support the business processes.

As shown in **Table 8**, being able to monitor the end user's experience and behavior is a very important management task. One of the reasons for that importance is that in spite of all of the effort and resources that have gone into implementing IT management to date:

It is the end user, and not the IT organization who typically is the first to notice when the performance of an application begins to degrade.

An effective approach to application performance management must include the automatic discovery of all the elements in the IT infrastructure that support each service. This functionality provides the basis for an IT organization to being able to create two-way mappings between the services and the supporting infrastructure components. These mappings, combined with event correlation and visualization, can facilitate root cause analysis, significantly reducing mean-time-to-repair. **Table 8** demonstrated how important this functionality is to IT organizations.

If IT organizations can effectively identify which components of the infrastructure support a particular application or service, monitoring can much more easily identify when services are about to begin to degrade due to problems in the infrastructure. As part of this monitoring, predictive technique such as heuristic-based trending of software issues and infrastructure key performance indicators can be employed to identify and alert management of problems before they impact end users. In addition, outages and other incidents that generate alerts can be prioritized based on their potential business impact.

As **Table 8** also demonstrates, getting better at rapidly identifying the causes of application degradation is the most important management task facing IT organizations. Once the components of the infrastructure that support a given application or service have been identified, triage and root cause analysis can be applied at both the application and the infrastructure levels. When applied directly to applications, triage and root cause analysis can identify application issues such as the depletion of threads and pooled resources, memory leaks or internal failures within a Java server or .NET server. At the infrastructure level, root cause analysis can determine the subsystem within the component that is causing the problem.

Application Performance Management in the Private Enterprise Network[16]

Enterprise IT organizations can choose among several types of tools for monitoring and managing application performance over a private enterprise network. These include: application agents, monitoring of real and synthetic transactions, network flow and packet capture, analytics, and dashboard portals for the visualization of results.

At a high level, there are two basic classes of tools. The first class of tool monitors global parameters such as user response time or transaction completion time and provides alerts when thresholds are

[16] This refers to managing the performance of applications that are delivered over WAN services such as Frame Relay, ATM and MPLS.

exceeded. These tools include agents on end user systems and monitoring appliances in the data center. The second class of tool supports triage by monitoring one or more of the components that make up the end-to-end path of the application. These tools include devices that capture application traffic at the flow and packet levels, agents on database, application, and web servers, as well as agents on various network elements.

Each type of individual tool has its strengths and weaknesses. For example, agents can supply the granular visibility that is required for complex troubleshooting but they represent an additional maintenance burden while also adding to the load on the servers and on the network. Monitoring appliances have more limited visibility, but they don't require modification of server configurations and don't add traffic to the network. Taking into consideration these trade-offs, IT organizations need to make tool decisions based on their goals for application performance management, their application and network environment as well as their existing infrastructure and network management vendors.

Independent of the approach that IT organizations take towards application performance management, a critical component of application performance management is end-to-end visibility.

End-to-end visibility refers to the ability of the IT organization to examine every component of IT that impacts communications once users hit ENTER or click the mouse button until they receive a response back from the application.

End-to-end visibility is one of the cornerstones of assuring acceptable application performance. This functionality is important because it:

- Provides the information that allows IT organizations to notice application performance degradation before the end user does.
- Identifies the symptoms of the degradation and as a result enables the IT organization to reduce the amount of time it takes to identify and remove the causes of the degraded application performance.
- Facilitates making intelligent decisions and getting buy-in from other impacted groups.
- Allows the IT organization to measure the performance of a critical application before, during and after a change is made.

The value of providing end-to-end visibility is maximized if two criteria are met. One criterion is that all members of the IT organization use the same tool or set of tools. The second criterion is that the tool(s) are detailed and accurate enough to identify the sources of application degradation.

Application Performance Management in Public and Hybrid Clouds

There are a number of possible ways that an IT organization can adjust their application performance management strategies in order to accommodate accessing services hosted by a Cloud Computing Service Provider (CCSP). These include:

- Extend the enterprise monitoring solutions into the public cloud using agents on virtual servers and by using virtual appliances.

- Focus on CCSPs that offer either cloud resource monitoring or application performance management as a service.

- Increase the focus on service delivery and transaction performance by supplementing existing application performance management solutions with capabilities that provide an outside-in service delivery view from the perspective of a client accessing enterprise applications or cloud applications over the Internet or mobile networks.

Application Aware Network Performance Management

There are a number of indications that enterprise networks and the applications that transit these networks are becoming increasingly entwined. One indication is HP's implantation of an SDN app store. The goal of the SDN app store is to enable network operations teams to download applications into their SDN controllers in the same way that smartphone users download apps onto their devices. Another indication is the formation of the Unified Communications Interoperability Forum (UCIF). UCIF's mission is to work with the Open Networking Foundation (ONF) to develop a standardized way for applications to dynamically request services from a software defined network. Yet another indication is the number of network vendors who have announced a network marketecture build around concepts such as *application-aware, application-centric or application driven.*

In response to the fact that enterprise networks and the applications that transit these networks are becoming increasingly entwined, there has been a movement to bring together two management disciplines: Application Performance Management and Network Performance Management. The result of bringing together those two disciplines is a new discipline that is often referred to as Application Aware Network Performance Management (AANPM). An AANPM solution integrates data that has historically been associated with application performance management with data that has historically been associated with network performance management. The result is a system that provides cross-platform visibility that enables IT organizations to monitor, troubleshoot and analyze both network and application systems.

DevOps

The phrase *DevOps* is a result of bringing to together two phrases: *Development* and *Operations*. That's appropriate because the point of adopting DevOps is to establish tight collaboration between a number of the phases of the application development lifecycle, including application development, testing, implementation and ongoing operations. DevOps is not a technology, but an approach. Some of the key characteristics of the approach are that the applications development team writes primarily small incremental pieces of code that are tested on an architecture that reflects the production architecture. Ideally, the network on which the software is tested will reflect not just the architecture but also the same characteristics (i.e., delay, packet loss) as the production network.

Implementing DevOps provides many advantages. For example, DevOps can provide business value by enabling companies to experience sustained innovation. Examples of companies that claim to have experienced sustained innovation as a result of implementing DevOps include Twitter, Netflix and Facebook. Implementing DevOps has other advantages. According to a recent Information Week Report, eighty two percent of the IT organizations that implemented DevOps saw at least some improvement in infrastructure stability and eighty three percent saw at least some improvement in the speed of application development.

A number of service providers who are attempting to become more agile have commented on the need for their organization to move away from slow, manual processes. One such provider is Deutsche Telekom. In a recent article, Deutsche Telekom was quoted as saying: "DT [Deutsche Telekom] needs to build a team that comprises IP, datacenter, programming, and operations specialists that can work in small, empowered, and agile teams, while both the carriers and vendors need to adjust for the migration from hardware-based to software-based business models."

GE Capital is also an advocate of DevOps. In a recent blog, GE Capital's CTO Eric Reed explained some of the impact of DevOps on the IT organization. According to Eric, "Our experience [GE Capital's] on this journey to date has been that the small, self-directed teams required in a DevOps world require an amalgamation of skills spanning everything from IT security to database design and application architecture, plus everything in between. While each individual on the team has a particular strength (say, application design and coding), each one also needs to have working knowledge in other areas (maybe UX or network design)."

Security

How IT Organizations are Implementing 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, can rent botnets, 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.

IT security systems and policies have evolved and developed around the traditional application delivery architecture in which branch offices are connected using a private WAN service to application servers in a central corporate data centers. In this architecture, the central corporate data center is a natural location to implement IT security systems and policies that provide layered defenses as well a single, cost efficient location for a variety of IT security functions. With the adoption of public cloud computing, applications and services are moving out of the central corporate data center and there is no longer a convenient single location for security policies and systems.

In addition, IT security systems and policies have traditionally distinguished between people who were using IT services for work versus those who were using it for personal use. The use of an employer provided laptop was subject to the employer's IT security policies and systems. In this environment, the use that employees made of personal laptops was generally outside of the corporate IT security policy. With the arrival of smartphones and tablet computers, the ownership, operating systems and security capabilities of the end user devices have changed radically. IT security policies and standards that were developed for PCs are no longer effective nor optimal with these devices. Most corporations have embraced the BYOD movement and end users are less willing to accept strict corporate security policies on devices they own. Additionally, strict separation of work and personal usage on an employee owned device is impractical.

The current and emerging environment creates a set of demanding security challenges.

The demands of governments, industry and customers are another factor that has historically shaped IT security systems and policies. The wide diversity of organizations that create regulations and standards can lead to conflicts. For example, law enforcement requires access to network communications (Communications Assistance for Law Enforcement Act – CALEA) which may in turn force the creation of locations in the network that do not comply with the encryption requirements of other standards (e.g. Health Insurance Portability Accountability Act – HIPPA).

Current State of DDoS Attacks

There is a wide range of ways that a DDoS attack can cause harm to an organization, including the:

- Consumption of computational resources, such as bandwidth, disk space, or processor time;
- Disruption of configuration information, such as routing information;
- Disruption of state information, such as the unsolicited resetting of TCP sessions;
- Disruption of physical network components;
- Obstructing the communication media between the intended users and the victim so that they can no longer communicate adequately.

As noted in the preceding chapter of ***The Handbook***, the majority of the security attacks are to a data center because that's where most of the applications and most of the data resides**.** Every year a number of organizations publish an annual report on the state of security attacks. A number of recent reports have highlighted the evolving threats that are associated with DDoS attacks on data centers.

One of the key findings of those reports is that DDoS attacks are growing in a variety of ways, including:

- Frequency: A 50% increase in DDoS attacks on a year-to-year basis[17];
- Size: One attack out of three is over 20 Gbps, 60 Gbps attacks are common and 100 Gbps attacks are not uncommon[18];
- Severity: The average number of packets per second in a DDoS attack has increased 1,850% to 7.8 Mpps between 2011 and 2013[19];
- Sophistication: 81% of attacks are multi-vector threats[20];
- Persistence: The average duration of a DDoS attack is 17 hours[21].

One example of the state of DDoS attacks is found in the Q1 2014 report from Prolexic. That report mentions a 10 hour long DDoS assault that peaked at over 200 Gbps and 53.5 MPPS.

There are a many components of a DDoS attack, but one component that is common to all such attacks is having a large scale botnet network that sends traffic towards the target at very high rates. As mentioned, it is possible to rent a botnet. The usual way this works is that criminal syndicates and commercially-motivated hackers build botnet networks that can be rented on-demand over the Internet. These on-demand networks are typically known as *booters*[22] and are often marketed at Web performance test tools or *stressers*.

[17] Akamai State of The Internet report
[18] Neustar Annual DDoS Attacks and Impact Report
[19] Verizon data breach report 2014
[20] Incapsula 2013-2014 DDoS Threat Landscape Report
[21] Prolexic Global Attack Report Q1 2014
[22] Safetyskyhacks Top 10 DDosers, Booters, Stressers

Cloud-Based Security

One way that a Cloud-based Security Service (CBSS) could provide value is if it provides protection against the growing number of malware attacks. To effectively protect against malware attacks, a CBSS should be able to identify suspicious content or sites that are either suspicious or are known to distribute malware. In order to be effective, a CBSS that provides Web content filtering or malware protection needs a source of intellectual capital that identifies known and suspected vulnerabilities. This source needs to be both dynamic and as extensive as possible.

One part of the value proposition of a CBSS that provides security functionality is the same as the value proposition of any cloud based service. For example, a security focused CBSS reduces the capital investment in security that an organization would have to make. In addition, a security focused CBSS reduces the amount of time it takes to deploy new functionality. The speed at which changes can be made to a CBSS adds value in a variety of situations, including providing better protection against zero-day attacks. Another part of the value proposition of a security focused CBSS is that unlike a traditional security solution that relies on the implementation of a hardware based proxy, a CBSS can also protect mobile workers. The CBSS does this by leveraging functionality that it provides at its Points of Presence (POPs) as well as functionality in a software agent that is deployed on each mobile device.

In many instances, the best security solution is a hybrid solution that combines traditional on-premise functionality with one or more Cloud-based solutions. For example, in many cases IT organizations already have functionality such as Web filtering or malware protection deployed in CPE at some of their sites. In this case, the IT organization may choose to implement a CBSS just to protect the sites that don't have security functionality already implemented and/or to protect the organization's mobile workers. Alternatively, an organization may choose to implement security functionality in CPE at all of their sites and to also utilize a CBSS as part of a defense in depth strategy.

Web Application Firewall Services

The chapter of this report entitled Network and Application Optimization discussed how a Cloud-based service, such as the one shown in **Figure 12**, can be used to optimize the performance of the Internet.

As will be discussed in this sub-section of ***The Handbook***, that same type of service can also provide security functionality.

Figure 12: Internet Based Security Functionality

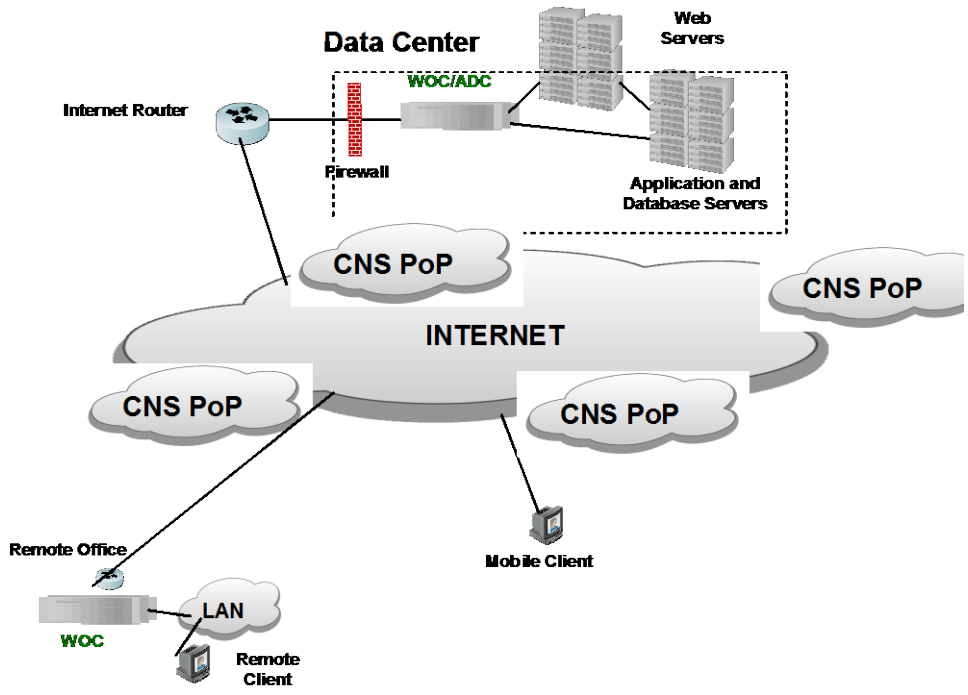

Role of a Traditional Firewall: Protect the Perimeter

Roughly twenty years ago IT organizations began to implement the first generation of network firewalls, which were referred to as packet filters. These devices were placed at the perimeter of the organization with the hope that they would prevent malicious activities from causing harm to the organization.

Today most network firewalls are based on *stateful* inspection. A *stateful* firewall holds in memory attributes of each connection. These attributes include such details as the IP addresses and ports involved in the connection and the sequence numbers of the packets traversing the connection. One of the weaknesses associated with network firewalls is that they are typically configured to open up ports 80 and 443 in order to allow passage of all HTTP and SSL traffic. Given that ports 80 and 443 are generally configured to be open, this form of perimeter defense is porous at best.

Whereas network firewalls are focused on parameters such as IP address and port numbers, a more recent class of firewall, referred to as a Web application firewall, analyzes messages at layer 7 of the OSI model. Web application firewalls are typically deployed as a hardware appliance and they sit

behind the network firewall and in front of the Web servers. In some cases, Web application firewall functionality is provided by an Application Delivery Controller (ADC).

Web application firewalls look for violations in the organization's established security policy. For example, the firewall may look for abnormal behavior, or signs of a known attack. It may also be configured to block specified content, such as certain websites or attempts to exploit known security vulnerabilities. Because of their ability to perform deep packet inspection at layer 7 of the OSI model, a Web application firewall provides a level of security that cannot be provided by a network firewall.

Defense in Depth: The Role of a Web Application Firewall Service

As is well known, there are fundamental flaws with an approach to security that focuses only on the perimeter of the organization. To overcome these flaws, most IT organizations have moved to the previously referenced approach to security: *defense in depth*. The concept of defense in depth is not new. What is new in the current environment is the use of a CBSS to provide Web application firewall functionality that is distributed throughout the Internet and to use this functionality to supplement security functionality that is provided on site by devices such as a Web application firewall or an ADC that offers Web application firewall functionality.

Because there are a variety of possible DDoS attacks, IT organizations need to implement a variety of defense in depth techniques. This includes:

- **Minimizing the points of vulnerability**
  If an organization has most or all of its important assets in a small number of locations, this makes the organization more vulnerable to successfully being attacked as the attacker has fewer sites on which to concentrate their attack.

- **Protecting DNS**
  Many IT organizations implement just two or three DNS servers. As such, DNS is an example of what was discussed in the preceding bullet – how IT organization are vulnerable because their key assets are located in a small number of locations.

- **Implementing robust, multi-tiered failover**
  Many IT organizations have implemented disaster recovery plans that call for there to be a stand-by data center that can support at least some of the organization's key applications if the primary data center fails. Distributing this functionality around a global network increases overall availability in general, and dramatically reduces the chance of an outage due to a DDoS attack in particular.

In order to be effective, a CBSS that provides Web application firewall functionality needs to be deployed as broadly as possible, preferably in thousands of locations. When responding to an attack, the service must also be able to:

- Block or redirect requests based on characteristics such as the originating geographic location and whether or not the originating IP addresses are on either a whitelist or a blacklist.

- Direct traffic away from specific servers or regions under attack.

- Issue slow responses to the machines conducting the attack. The goal of this technique, known as tarpits[23], is to shut down the attacking machines while minimizing the impact on legitimate users.

- Direct the attack traffic back to the requesting machine at the DNS or HTTP level.

As noted, a CBSS that provides Web application firewall functionality is complimentary to premise-based security functionality such as that provided by an ADC that offers Web application firewall. That follows because while the Cloud-based Web application firewall service can perform many security functions that cannot be performed by an on premise Web application firewall, there are some security functions that are best performed by an on premise Web application firewall. An example of that is protecting an organization against information leakage by having an onsite Web application firewall perform deep packet inspection to detect if sensitive data such as a social security number or a credit card number is leaving the site. If sensitive data is leaving the site, the onsite Web application firewall, in conjunction with other security devices, can determine if that is authorized and if it is not, it can prevent the data from leaving the site.

Impact of SDN

As was the case with management, SDN poses both security challenges and security opportunities. The primary security challenge is to ensure that an attacker cannot compromise the central SDN controller and hence have access to all of the subtending network elements. In addition to securing the controller itself, all communication between the controller and other devices including switches, network services platforms and management systems must be secured.

As noted, SDN also presents opportunities to improve security by implementing security related applications that leverage the control information that has been centralized in the SDN controller. One such application that has been announced[24] is a network service that provides DDoS protection. Another such example is an application[25] that was designed to combat the security challenges that are associated with BYOD.

The document entitled Mock RFI for Enterprise SDN Solutions contains a number of questions that IT organizations should ask SDN vendors relative to security.

[23] Wikipedia Tarpit (networking)
[24] Radware Defense Flow Report
[25] HP Network Protector SDN Application

Conclusions

The following is a summary of the conclusions that were reached in the preceding sections of ***The Handbook***.

- IT organizations need to plan for optimization, security and management in an integrated fashion.
- The goal of the 2014 Application and Service Delivery Handbook is to help IT organizations ensure acceptable application and/or service delivery when faced with both the first generation, as well as the emerging set of application and service delivery challenges.
- IT managers face the application delivery challenges associated with both the legacy environment and the emerging environment.
- The vast majority of employees require mobile access for at least part of their typical day.
- The BYOD movement has resulted in a loss of control and policy enforcement.
- Adopting BYOD increases a company's vulnerability to security breaches.
- Getting better at managing and optimizing the delivery of mobile application is either very or extremely important to the majority of IT organizations.
- Almost two thirds of the IT organizations consider it to be either very or extremely important over the next year for them to get better performing management tasks such as troubleshooting on a per-VM basis.
- From a networking perspective, the primary challenge in implementing desktop virtualization is achieving adequate performance and an acceptable user experience for client-to-server connections over a WAN.
- Getting better at optimizing the performance of virtualized desktops is becoming increasingly more important.
- The goal of cloud computing is to enable IT organizations to achieve a dramatic improvement in the cost effective, elastic provisioning of IT services.
- Managing end-to-end in both a private and public cloud environment has become notably more important to IT organizations over the last year.
- Managing end-to-end in a private cloud environment is slightly more important to IT organizations than is managing end-to-end in a public cloud environment.
- 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.

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

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

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

- 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.

- 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.

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

- Two tasks are in a virtual tie for the most important management task to get better at over the next year: 1) Rapidly identifying the root cause of degraded application performance; 2) Identifying the components of the IT infrastructure that support the company's critical business applications.

- The second most important set of management tasks include: 1) Obtain performance indicator metrics and granular data that can be used to detect and eliminate impending problems; 2) Monitor the end user's experience and behavior; 3) Effectively manage SLAs for one or more business critical applications; 4) Manage the use of VoIP.

- While managing the use of services acquired from an IaaS provider such as Rackspace is relatively unimportant, it is more important than it was last year.

- IT organizations need to adopt an approach to management that is based on the assumption that the components of a service, and the location of those components, can and will change frequently.

- It is the end user, and not the IT organization who typically is the first to notice when the performance of an application begins to degrade.

- End-to-end visibility refers to the ability of the IT organization to examine every component of IT that impacts communications once users hit ENTER or click the mouse button until they receive a response back from the application.

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

- The current and emerging environment creates a set of demanding security challenges.

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.

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

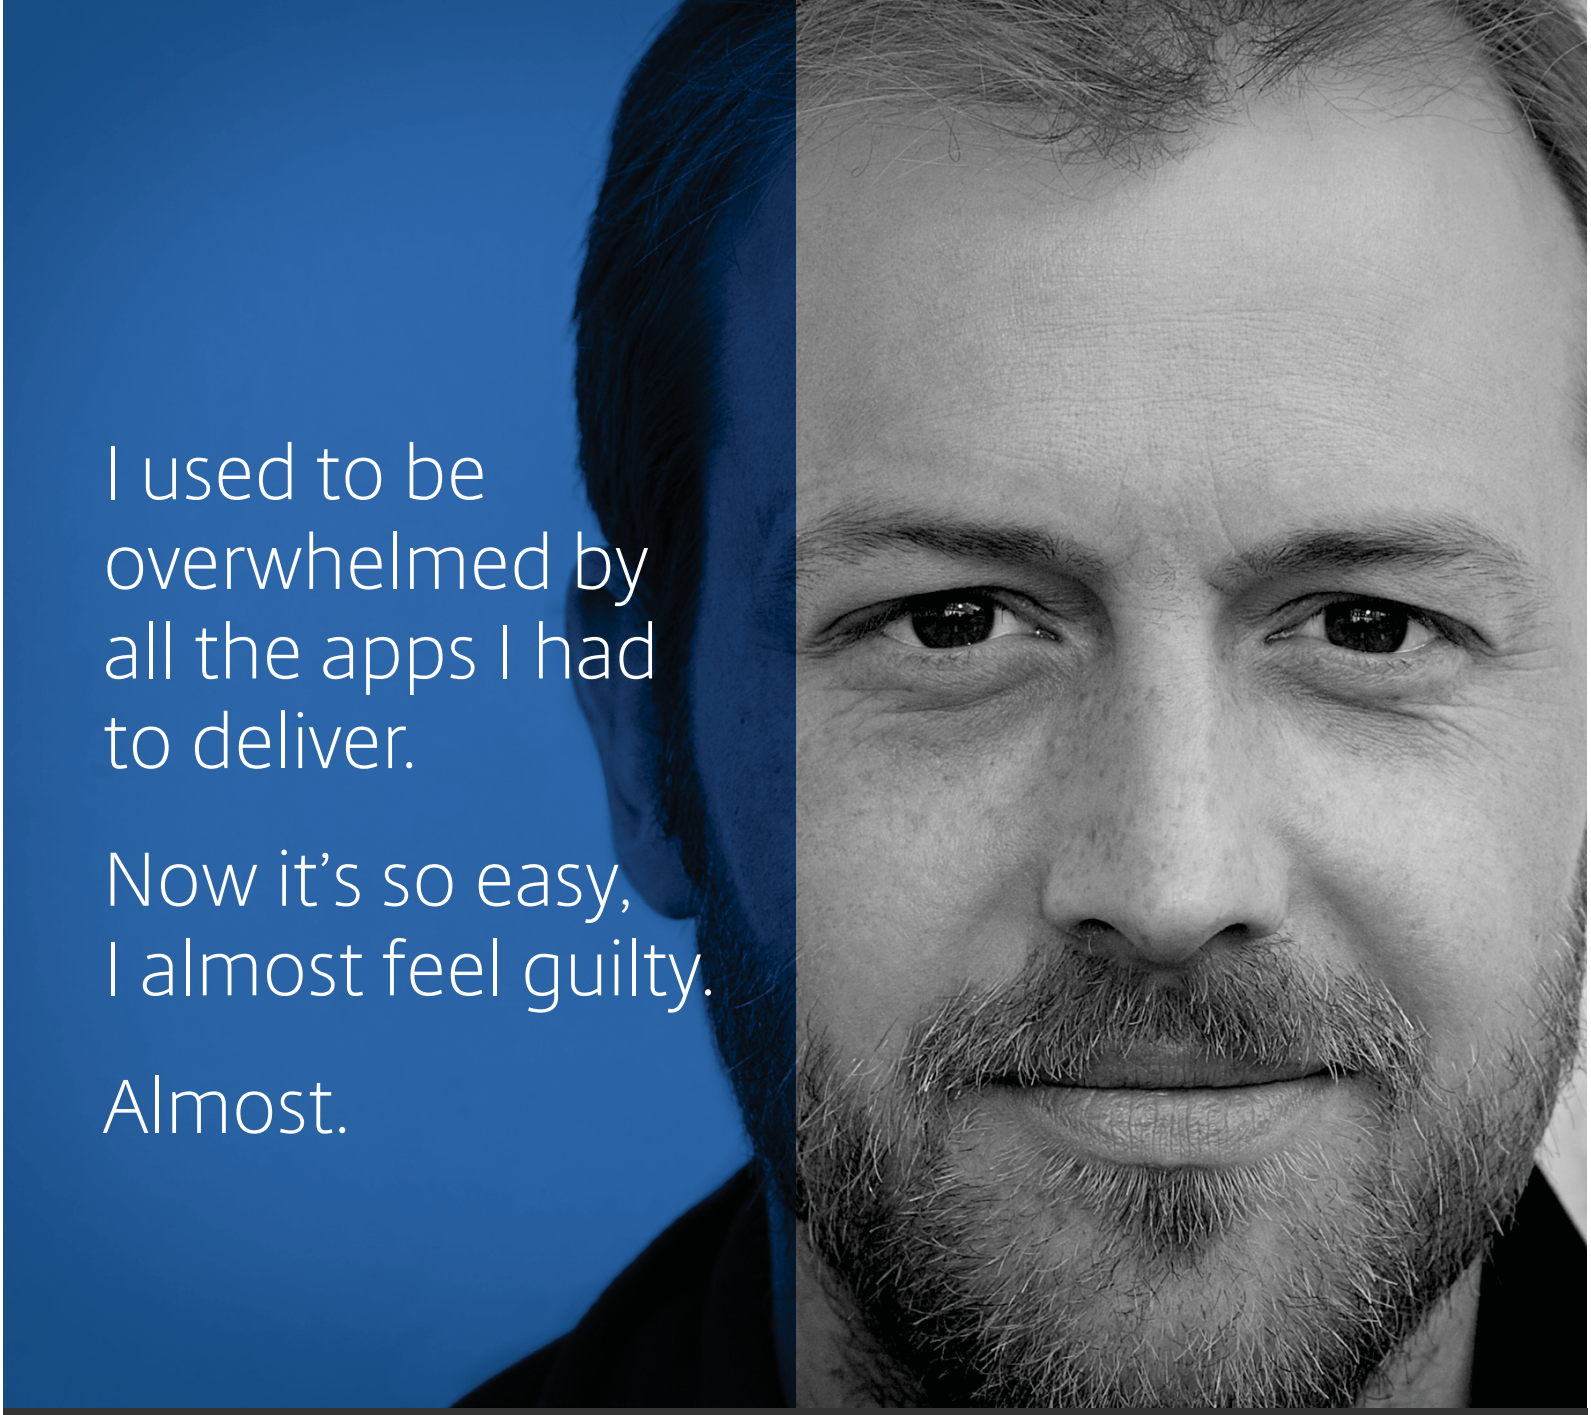

NetScaler with TriScale harnesses the power of software so you can effortlessly customize your app delivery for any business need.

NetScaler with TriScale
SOFTWARE SMART. HARDWARE STRONG.

CITRIX®

www.citrix.com/netscaler

CITRIX

Customer Spotlight

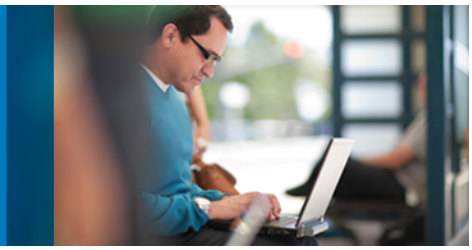

City Index

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.

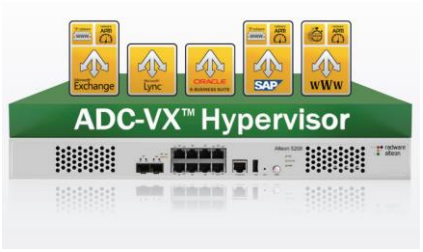

Predictable Application SLA, Guaranteed. Only with Alteon NG.

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.