The 2015
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

Introduction

Throughout the ***2015 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 strong relationship between the requirements listed above. For example, in order to implement an appropriate level of security, an IT organization may adopt 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.

The Handbook builds on the 2014 edition of the Application and Service Delivery Handbook. However, any material in the 2014 edition that was deemed to be no longer relevant was removed. Content that was deemed to be relevant but well understood by the majority of IT organizations was removed, stored online and referred to in the 2015 edition of **The Handbook** with a URL. Using this approach, **The Handbook** is of manageable size and focuses primarily on the changing nature of application and service delivery.

In early 2015, multiple surveys were 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***. Because of its key role in application and service delivery, one of the surveys focused on the WAN. The other survey focused on identifying the optimization, management and security tasks that are of most interest to IT organizations. The answers to the surveys will be used throughout the ***2015 Application and Service Delivery Handbook*** to document the current end emerging state of application and service delivery.

Traditional 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 listed below and are described in detail in the document entitled Traditional Application & Service Delivery Challenges.

- Limited focus on performance during 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.

The Changing Application & Service Delivery Environment

There are a number of factors that are driving fundamental change in the application and service delivery environment. Those factors include the:

- Fast paced business environment;
- Evolving application and service delivery models;
- Virtualization of networks and network functions;
- Use of policy;
- Expectations of business unit managers;
- Adoption of DevOps.

The Fast Paced Business Environment

One of the key characteristics of the current business environment is the quickening pace of change. One measure of the quickening pace of business was provided by Dr. Richard Foster of Yale University[1] who stated that "The average lifespan of an S&P 500 company has decreased by more than 50 years in the last century, from 67 years in the 1920s to just 15 years today." Foster added that "By 2020, more than three-quarters of the S&P 500 will be companies that we have not heard of yet." One of the opportunities for the IT function to be perceived as a driver of change comes from the movement to become a *Digital Business*. It would be a mistake to think of *Digital Business* only in the context of companies like Google and Amazon, as virtually all companies are making at least some movement to become a digital business. According to an article in CIO magazine[2], The US retailer Home Depot is an example of a traditional company that is in the process of transitioning to become a digital business.

The Evolving Application and Service Delivery Model

A decade ago the most common application and service delivery model was the client server model. Most of the performance challenges associated with the client server model are included in the previously mentioned list of traditional application and service delivery challenges.

Over the last decade, a number of factors have caused the traditional client server-based application and services delivery model to evolve and become more challenging. Those factors include:

Guest Workers

Many companies want to provide internet access to guest workers, whether they are short term visitors or longer term temporary employees. While it would be technically possible to carry this traffic on the company's LAN, in many cases concerns over security have caused a number of companies to implement a separate network to carry the traffic generated by guest workers.

[1] http://www.bbc.co.uk/news/business-16611040
[2] Home Depot vs. Lowe's, April 1, 2014

Mobile Workers

In the current environment the vast majority of employees require mobile access for at least part of their day, whether they are within a company facility or at an external site. In the majority of instances the IT department isn't able to put any kind of an agent on the employees' mobile devices in order to facilitate optimizing performance or enabling effective management and security.

Server Virtualization

The vast majority of organizations have made at least some deployment of server virtualization and the deployment of server virtualization will continue to increase over the next several years. One of the potential advantages of server virtualization is the ability to dynamically move virtual machines (VMs) between physical servers. When VMs are migrated, however, the source and destination servers typically have to be on the same VLAN. 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.

Dynamic Multi-Pathing in the WAN

Some of the primary advantages of dynamic multi-pathing in the WAN are explained in *The 2015 Guide to WAN Architecture and Design*. A WAN that features dynamic multi-pathing has all of the traditional performance, management and security challenges. An additional management challenge is being able to identify the end-to-end path that traffic took in order to be able to troubleshoot degraded network or application performance.

Public Cloud Applications and Services

In the vast majority of instances, the use of public cloud computing services doesn't come with an SLA for the end-to-end performance[3] of the application or service because the service is virtually always delivered over the Internet. In addition, particularly when accessing SaaS-based applications, IT organizations often have little if any 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.

Private Cloud Computing

Similar to public cloud computing, two of the primary characteristics of private cloud computing are virtualization and automation. All of the traditional application and service delivery challenges apply as IT organizations begin to implement a private cloud computing model. In addition, all of the challenges that are introduced by the adoption of server virtualization also apply. Further complicating the management of an application and service delivery model that includes private cloud computing is that few companies will virtualize all of their data center functionality in the near term. Hence, IT organizations need the ability to manage applications and services that are delivered over a combination of physical and virtual resources.

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

The Virtualization of Networks and Network Functions

Neither Software Defined Networking (SDN) nor Network Functions Virtualization (NFV) are currently having a major impact on application and service delivery. However, both SDN and NFV have the potential in the near term to fundamentally impact application and service delivery by enabling the automated implementation of sophisticated forms of virtualized networks and virtualized network functions.

SDN

The initial discussion of SDN focused on the data center. However, there is a large and growing interest in implementing a software defined WAN (SD-WAN). As is the case with any software defined network, a SD-WAN centralizes the control function into a SDN controller. The controller abstracts the user's private network services from the underlying IP network and it enables the operation of the user's private network services via centralized policy. The controller also enables the automation of management tasks such as configuration and provisioning.

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

Network Functions Virtualization (NFV)

The primary factors that are driving communications service providers to develop and implement NFV are the desire to be more agile in the implementation of new services and the desire to reduce cost, notably OPEX. Driven by those same factors, over the last few years enterprise network organizations have implemented virtualized versions of a range of L4 – L7 network functions including Application Delivery Controllers, WAN Optimization Controllers, Firewalls and Intrusion Detection/Prevention systems. Enterprise network organizations typically don't require the same scale solutions as does a service provider. However, enterprise network organizations require all of the same characteristics for the virtualized network functions they implement as are included in the ETSI vision for NFV[4].

The Use of Policy

There is a broad movement to implement a policy based approach to all aspects of IT, including networking. Policies can be based on hierarchical system of rules designed to deal with the complexities of the environment, and to manage the relationships among users, services, SLAs, and device level performance metrics. One way that policy can be implemented is at the application level. For example, if the performance of an application begins to degrade because the CPU utilization of a physical server hosting a virtualized network function (VNF) that is used by that application becomes excessive, the VNF may be moved to a server with lower utilization, if that is in line with the policy that exists for that application.

The Expectations of Business Unit Managers

While IT has always been under pressure to show business relevance, what has changed over the last couple of years is that in a growing number of companies the role of the CIO is under attack from other

[4] Ibid.

C-level executives. One example of that trend was provided by Gartner who stated that by 2017 the CMO will spend more on IT than the CIO. In addition, driven by the market demand to transition from traditional bricks and mortar business models and adopt emerging digital business models, such as those adopted by Home Depot, a new type of C-Level executive is emerging: The Chief Digital Officer (CDO). The CDO is typically responsible for the development and management of the company's digital business models as well as the management and delivery of the company's digital assets.

The Adoption of DevOps

Some of the key principles of DevOps include:

- Collaboration;
- Continuous integration and delivery;
- Continuous testing and monitoring;
- Automation;
- API centric automated management interfaces.

According to a recent Information Week report, when asked to indicate the level of improvement in application development speed that they have either already gained or expected to gain as a result of adoption DevOps, forty-one percent of respondents indicated "significant improvement" and 42% indicated "some improvement".

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.
- Slightly less important than optimizing the performance of business critical applications is optimizing the transfer of storage associated with business continuity and disaster recovery between different data centers.
- A relatively new challenge, ensuring the performance of applications used by mobile workers, is now one of the most important optimization tasks facing IT organizations.

Traditional Optimization Appliances

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

WAN Optimization Controllers (WOCs)

When WOCs were first introduced in the mid-2000s, they were hardware-based appliances. While that is still an option, it is now possible to implement a software based WOC, which are often referred to a being a virtual WOC (vWOC).

There are some significant technical differences in the vWOCs that are currently available in the marketplace, such as which hypervisors are supported; e.g., hypervisors from the leading vendors such as VMware, Citrix and Microsoft as well as proprietary hypervisors from a cloud computing provider such as Amazon. There are also significant differences in terms of how vendors of virtual appliances structure the pricing of their products. One option, referred to as *pay as you go*, allows IT organizations to avoid the capital costs that are associated with a perpetual license and to acquire and pay for a vWOC or a virtual Application Delivery Controller (vADC) on an annual basis. Another option, referred to as *pay as you grow*, enables an IT organization to get stared by implementing vWOCs or vADCs that have relatively small capacity and are priced accordingly. The IT organization can upgrade to a higher-capacity vWOC or vADC when needed and only pay the difference between the price of the virtual appliance that it already has installed and the price of the virtual appliance that it wants to install.

Application Delivery Controllers (ADCs)

Background

The original purpose of an ADC was to provide load balancing across local servers or among geographically dispersed data centers. ADCs 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.

ADCs and Security

In the case of IT security, the majority of the attacks are to a data center because that's where most of the applications and most of the data resides. Given that the most common deployment of ADCs has them placed in front of application servers in a data center, they are in a strategic position to thwart attacks. In order to be effective thwarting security attacks, ADCs should have an ICSA-certified web application firewall and a DNS application firewall. It should provide protection against DDoS attacks and also support SSL offload and high speed SSL decryption with SSL intercept.

The Requirement for Programmability

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

Virtual ADCs

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;
- Network Appliance O/S Partitioning;
- Network Appliance with OEM Hypervisor;
- Network Appliance with Custom Hypervisor.

The Role of SDN

In a traditional data center implementing L4 – L7 services such as WOCs and Application Delivery Controllers (ADCs) is cumbersome and time consuming as it requires acquiring the requisite network appliances and cabling them together in the correct order. Since each appliance has its own unique interface, configuring these appliances is an error-prone task.

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

NFV Optimization

In order to obtain the potential cost and agility benefits of a software-based approach to providing IT functionality, it must be possible to achieve the same or greater performance in a software-based

environment as is possible in a traditional hardware-based environment. However, that isn't possible without an enabling software architecture because of the bottlenecks that are associated with the hypervisors, virtual switches and virtual machines that are the foundation of the emerging software-based approach to IT.

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:

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

The Next Generation WAN

The WAN introduces a range of demanding challenges relative to ensuring acceptable application and service delivery.

Background

WAN Evolution

The modern WAN got its start in 1969 with the deployment of the ARPANET which was the precursor to today's Internet. In addition to the continued evolution of the Internet, the twenty-year period that began around 1984 saw the deployment of four distinct generations of enterprise WAN technologies. This included:

- TDM;
- Frame Relay;
- ATM;
- MPLS.

Network organizations currently make relatively little use of WAN services other than MPLS and the Internet and the use they do make of those other services is decreasing somewhat rapidly.

Traditional WAN Design

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

Software Defined WANs

Hybrid WAN

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

Interest in Leveraging SDN in the WAN

The *2015 Guide to SDN and NFV* reported on the results of a survey that was administered in late 2014. The respondents to this survey indicated their belief that three years from now that SDN deployment in data centers will be highly pervasive and that there will also be significant SDN deployment both in the WAN and in campus networks.

Drivers and Inhibitors of SD WAN Adoption

The Survey Respondents were given a set of possible outcomes and were asked to indicate which outcomes would drive their company to implement a SD WAN. The top three responses were:

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

The Survey Respondents were also given a set factors and were asked to indicate which factors would inhibit their company from implementing a SD WAN. The top three responses were:

- The current technologies are unproven and/or immature;
- It would add complexity;
- The current products and/or services are unproven and/or immature.

Management

Key Management Tasks

The Survey Respondents were asked about the importance of a range of management tasks. Their feedback indicates that the most important management tasks to get better at over the next year are:

- Rapidly identifying the root cause of degraded application performance;
- Effectively managing SLAs for one or more business critical applications;
- Identifying the components of the IT infrastructure that support the company's critical business applications;
- Obtaining performance indicator metrics and granular data that can be used to detect and eliminate impending problems.

Existing Trends That Impact Management

Server Virtualization

An assumption that has underpinned the traditional approach to IT management was that the data center environment was static. However, part of the value proposition that is associated with server virtualization is that it is possible to migrate VMs between physical servers. The fact that VMs migrate between physical servers is one of the reasons 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.

Cloud Computing

The adoption of varying forms of cloud computing (i.e., private, public, hybrid) demonstrates that 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.

Real-Time Applications

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. For example, effectively managing voice and video requires looking at the application payload and measuring the quality of the voice and video communications.

Emerging Trends That Impact Management

SDN

One of the management challenges that applies across multiple tiers of the SDN architecture is the requirement to manage the messaging that goes between tiers; e.g., between the application tier and the control tier. At the infrastructure tier, one of the primary challenges is to perform element management potentially of both virtual and physical network elements. One of the management

challenges at the control layer results from the fact that 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.

One of the management challenges that occurs at the application tier is that based on the type of application (e.g., business application vs. a firewall), the service or application needs varying levels of visibility into the underlying network. Another set of management challenges that occurs at the application layer stem from the requirement to ensure acceptable performance. This means that network infrastructure must have visibility into the SLA requirements of the application so that when faced with a spike in demand, a policy-based decision can be made as to whether or not resources should be dynamically allocated to meet those demands.

Looking at network virtualization as an application of SDN, another performance management challenge stems from the fact that one of the primary benefits of overlay-based SDN solutions is the ability to support multiple virtual networks that run on top of a physical network. In order to perform management functions such as root cause analysis and impact analysis, network management organizations need the ability to see the bilateral mapping between the virtual networks and the physical network that supports them.

NFV

Some of the key NFV-related management challenges are described below.

Dynamic relationships between software and hardware components

Due to the mobility of VMs, topology changes can occur in a matter of seconds or minutes rather than the days or weeks required for changing software/hardware relationships in traditional networks. In order to accommodate and leverage virtualization technologies, end-to-end management systems need to be re-architected to be capable of implementing automated processes for virtual resource procurement, allocation, and reconfiguration in accordance with a set of highly granular policies designed to ensure the quality of experience for the user of the network services.

Many-to-Many relationships between network services and the underlying infrastructure

In a virtualized infrastructure a network service can be supported by a number of Virtualized Network Function (VNFs) which may be running on one or several VMs. A single VNF may also support a number of distinct network services. In addition, the group of VNFs supporting a single network service could possibly be running on a number of distinct physical servers. As a result, end-to-end management systems need to support a three-tiered network model based on many-to-many relationships among network services, virtualization infrastructure, and physical infrastructure.

DevOps

One challenge that distinguishes NetOps from DevOps is that since VNFs such as optimization and security are chained together to create an end-to-end service this creates strong dependencies between the VNFs. For example, if an IT organization updates an optimization VNF they need to ensure that it is fully compatible with the security VNF(s). As a result much stronger version control and compatibility testing is needed than would be typical for enterprise applications.

Other challenges created by network services development that must be addressed by NetOps that were not addressed by DevOps include:

- Virtualized services will often be created by integrating services from multiple suppliers. This will require NetOps methodologies and best practices to support concurrent synchronized development and integration across the domains of multiple partners.
- NetOps will need to support dynamic and automated management of service performance and SLAs. This can only be achieved by a policy model that supports end-to-end SLA targets.
- NFV services are often mission critical. This creates a need for high levels of resilience and rapid fallback capabilities.

Security

The Changing Security Environment

The security landscape has changed dramatically in the last few years. In the 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 users are connected to application servers in a central corporate data centers by using an enterprise WAN service such as MPLS. 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 and hybrid cloud computing, applications and services are moving out of the central corporate data center and there is no longer a well-agreed to location for security policies and systems.

The demands of governments, industry and customers are another factor that has historically shaped IT security systems and policies. Unfortunately, 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).

Existing Trends That Impact Security

The *IBM X-Force Threat Intelligence Quarterly, 1Q 2015* identified some of the key security-related trends. Some of the trends that IBM identified are:

- The total number of leaked records (i.e., emails, credit card numbers, passwords and other personally identifiable information) continued to increase on an annual basis. It was a billion leaked records in 2014 which is an increase of 25% over the 800 million records that were leaked in 2013.

- Last year mobile devices were shown to present some unique security vulnerabilities. For example, in 2014 a Computer Emergency Readiness Team-Coordination Center (CERT/CC) researcher discovered security issues in thousands of Android applications. These vulnerabilities can allow an attacker to perform man-in-the-middle attacks against affected mobile applications.

- In 2014, the underlying libraries that handle cryptographic functionality on nearly every common web platform were found to be vulnerable to fairly trivial remote exploitations capable of stealing critical data.

The *IBM X-Force Threat Intelligence Quarterly, 1Q 2015* also presented survey data that identified the percentage of the totality of security incidents in 2014 that were attributable to a particular type of security attack. The top three types of security attacks were:

- Malware;
- DDoS:
- SQL injections.

Emerging Trends That Impact Security

SDN

Some of the security challenges related to SDN are described in *SDN Security Considerations in the Data Center*.

There are many ways that SDN can enhance security. For example, role based access can be implemented by deploying a role-based resource allocation application that leverages the control information and capability of the SDN controller. Another example is that by virtue of Layer 2-4 flow matching capability, OpenFlow access switches can perform filtering of packets as they enter the network, acting as simple firewalls at the edge. With OpenFlow switches that support modification of packet headers, an OpenFlow-enabled controller will also be able to have the switch redirect certain suspicious traffic flows to higher-layer security controls, such as IDS/IPS systems, application firewalls, and Data Loss Prevention (DLP) devices. Other security applications built on an OpenFlow controller can match suspicious flows to databases of malware signatures or divert DDoS attacks.

NFV

A number of organizations are focused on resolving the security issues associated with SDN and NFV. One such organization is the Internet Engineering Task Force (IETF). The IETF has created a security architecture that is based on horizontal (a.k.a., east/west) APIs in addition to the northbound and southbound APIs[5]. One IETF SDN-specific activity focuses on centralized security services (i.e., firewalls and DDoS mitigation systems) designed specifically for SDN environments[6]. Another SDN-specific Internet draft addresses the possible application of DevOps principles to SDNs[7].

ETSI is another organizations focused on resolving the security issues associated with SDN and NFV. In a document entitled *Network Functions Virtualization (NFV); NFV Security; Security and Trust Guidance*, ETSI outlined some high level security goals for NFV.

[5] https://datatracker.ietf.org/doc/draft-bernardo-sec-arch-sdnnvf-architecture/
[6] https://datatracker.ietf.org/doc/draft-jeong-l2nsf-sdn-security-services/
[7] https://datatracker.ietf.org/doc/draft-unify-nfvrg-devops/

The Emerging Application and Service Delivery Environment

Traditional 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 performance during 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.

The Changing Application & Service Delivery Environment

There are a number of factors that are driving fundamental change in the application and service delivery environment. Those factors include the:

- Fast paced business environment;
- Evolving application and service delivery models;
- Virtualization of networks and network functions;
- Use of policy;
- Expectations of business unit managers;
- Adoption of DevOps.

The Fast Paced Business Environment

One of the key characteristics of the current business environment is the quickening pace of change. One measure of the quickening pace of business was provided by Dr. Richard Foster of Yale University[8] who stated that "The average lifespan of an S&P 500 company has decreased by more than 50 years in the last century, from 67 years in the 1920s to just 15 years today." Foster added that "By 2020, more than three-quarters of the S&P 500 will be companies that we have not heard of yet."

As a minimum, the IT function needs to enable rapid business change. Ideally, the IT function is perceived as a driver of that change.

One of the opportunities for the IT function to be perceived as a driver of change comes from the movement to become a *Digital Business*. A digital business has four key pillars:

- Customer centricity;
- Operational agility and effectiveness;
- Agile business models and rapid innovation;
- An agile IT function.

It would be a mistake to think of *Digital Business* only in the context of companies like Google and Amazon, as virtually all companies are making at least some movement to become a digital business. The US retailer Home Depot is an example of a traditional company that is in the process of transitioning to become a digital business. According to an article in CIO magazine[9], Home Depot has already implemented:

- Onmichannel efforts that include BORIS (buy online, return in store) and BOSS (buy online, ship to store) programs, which complement the BOPIS (buy online, pick up in store) system that was previously launched;
- A project to create a mobile mapping app aimed to help customers more easily find items in Home Depot's cavernous stores.

In addition, Home Depot has started to use analytics to help them with competitive pricing.

[8] http://www.bbc.co.uk/news/business-16611040
[9] Home Depot vs. Lowe's, April 1, 2014

The Evolving Application and Service Delivery Models

This section discusses some of the most common application and service delivery models and the associated issues relative to performance, management and security. This discussion isn't intended to imply the progression that a company takes to go from one model to another, nor is it intended to imply that a company implements just one of these delivery models. In virtually all instances, companies implement multiple application and service delivery models simultaneously.

Client Server

A decade ago the most common application and service delivery model was the hardware-centric client server model shown in **Figure 1**.

Figure 1: Client Server Delivery Model

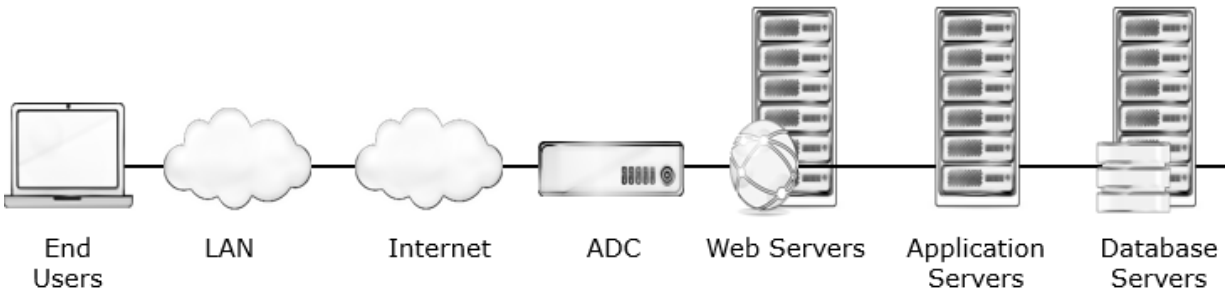

Most of the performance challenges associated with the client server model are included in the previously mentioned list of traditional application and service delivery challenges; i.e., network latency; chatty protocols and applications. One of the key management challenges associated with this delivery model has to do with being able to gather and correlate management data over a wide range of types of equipment, often managed by different groups. The client server delivery model has a range of security vulnerabilities, including susceptibility to the types of DDoS attacks that are caused by a TCP SYN flood.

Guest Workers

Many companies want to provide internet access to guest workers, whether they are short term visitors or longer term temporary employees. While it would be technically possible to carry this traffic on the company's LAN, in many cases concerns over security have caused a number of companies to implement a separate network to carry the traffic generated by guest workers.

Mobile Workers

In the current environment the vast majority of employees require mobile access for at least part of their day, whether they are within a company facility or at an external site. In the majority of instances the IT department isn't able to put any kind of an agent on the employees' mobile devices in order to facilitate optimizing performance or enabling effective management and security. The challenges that result from losing control of the user's access device are exacerbated when the user is using a cellular network due to the high delay and packet loss that is often associated with those networks as well as the increased ease of snooping that traffic.

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 1**.

Table 1: Importance of Getting Better Delivering Mobile Applications

	Improving the Performance	Managing and Monitoring
Extremely Important	22%	28%
Very Important	38%	30%
Moderately Important	23%	23%
Slightly Important	18%	15%
Not at all Important	0%	4%

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

Server Virtualization

The vast majority of organizations have made at least some deployment of server virtualization and the deployment of server virtualization will continue to 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, and 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

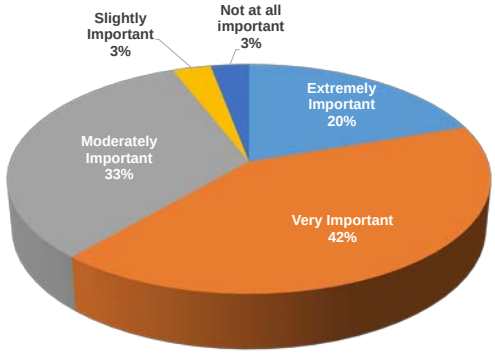

Almost two thirds of the IT organizations consider it to be either very or extremely important over the next year for them to get better performing management tasks such as troubleshooting on a per-VM basis.

There are some significant networking problems associated with server virtualization. For example, one of the potential advantages of server virtualization is the ability to dynamically move virtual machines (VMs) between physical servers, both within a data center and between data centers. 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. 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.

Dynamic Multi-Pathing in the WAN

As shown in **Figure 3**, when asked about their use of dynamic multi-pathing, the most common answer given by The Survey Respondents was that they had already made a significant deployment of the functionality. The second most common answer was that they would increase their current deployment over the next year.

One way to leverage this functionality is to dynamically load balance traffic over both MPLS and Internet links based on centralized policies that indicate the business criticality of the application. This approach has the goal of reducing the capacity, and hence the cost, of the MPLS links and replacing the reduced MPLS bandwidth with relatively inexpensive Internet bandwidth.

Figure 3: Approach to Dynamic Multi-Pathing

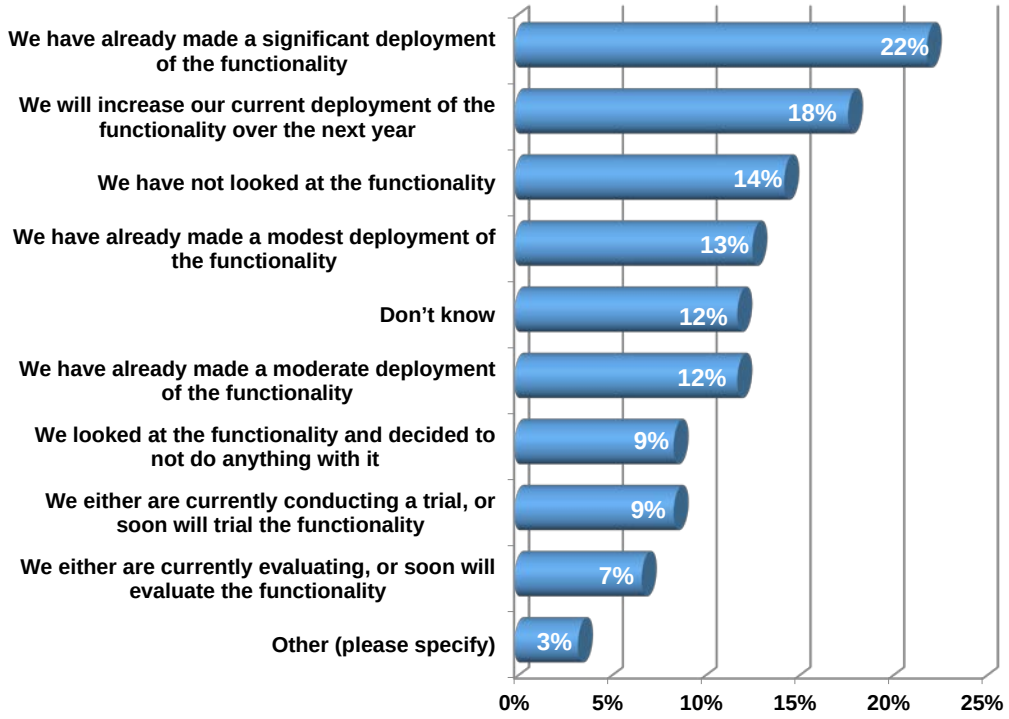

A WAN that features dynamic multi-pathing has all of the traditional performance, management and security challenges. An additional management challenge is being able to identify the end-to-end path that traffic took in order to be able to troubleshoot degraded network or application performance.

Public Cloud Applications and Services

In the vast majority of instances, the use of public cloud computing services doesn't come with an SLA for the end-to-end performance[10] of the application of service because the service is virtually always delivered over the Internet and nobody provides a performance guarantee for the Internet. In addition, particularly when accessing SaaS-based applications, IT organizations often have little if any 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.

In order to quantify the concern amongst IT organizations about ensuring acceptable application and service delivery when accessing public cloud applications and services, 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 optimizing the performance of applications and services acquired from public cloud providers. They were also asked how important it is for their IT organization over the next year to get better at managing end-to-end in a public cloud environment. Their responses are shown in **Table 2**.

Table 2: Importance of Getting Better Delivering Public Cloud Apps and Services

	Improving the Performance	Managing End-to-End
Extremely Important	11%	22%
Very Important	29%	32%
Moderately Important	26%	18%
Slightly Important	21%	19%
Not at all Important	13%	10%

Getting better at improving the performance of applications and services acquired from a public cloud provider is either very or extremely important to well over a third of IT organizations.

Getting better at managing end-to-end in a public cloud environment is either very or extremely important to the majority of IT organizations.

Private Cloud Computing

Similar to public cloud computing, two of the primary characteristics of private cloud computing are virtualization and automation. All of the traditional application and service delivery challenges still apply if IT organizations begin to implement a private cloud computing model. In addition, the management and challenges that are introduced by the adoption of server virtualization also apply. Further complicating the management of an application and service delivery model that includes private cloud computing is that few companies will virtualize all of their data center functionality in the near term. Hence, IT organizations need the ability to manage applications and services that are delivered over a combination of physical and virtual 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. Their answers are shown in **Figure 4**.

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

Figure 4: Importance of Getting Better at Managing Private Cloud

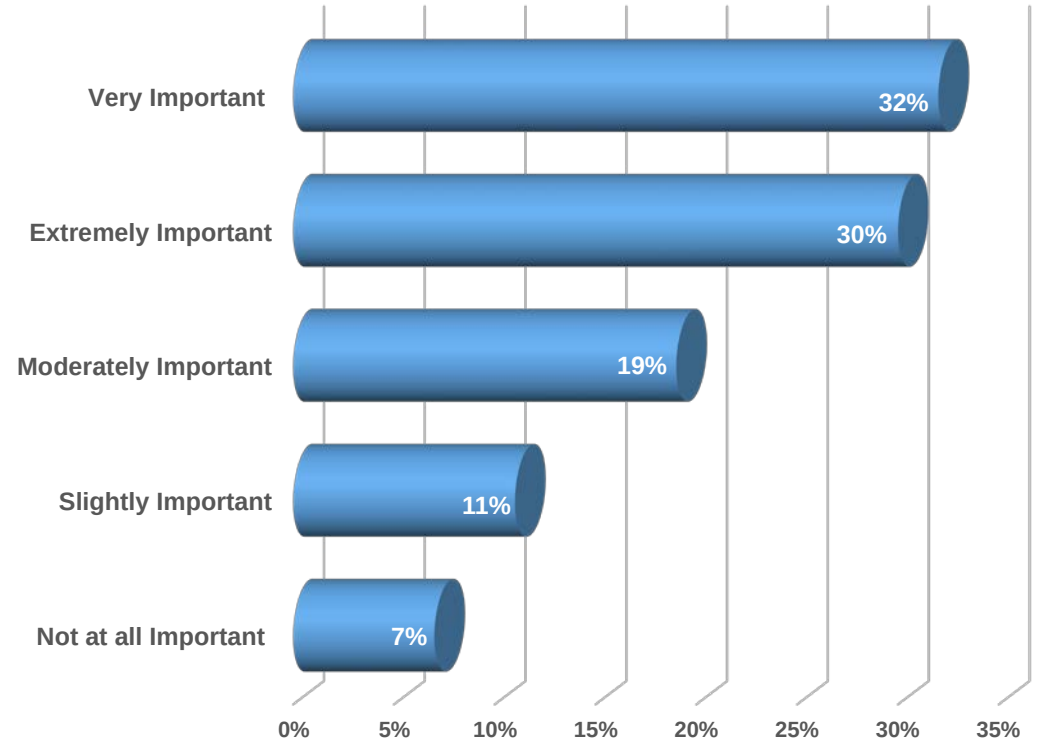

Managing end-to-end in a private cloud environment is slightly more important to IT organizations than is managing end-to-end in a public cloud environment.

Hybrid Cloud Computing

The phrase hybrid cloud refers to a composition of two or more clouds that remain distinct entities but are bound together, offering the benefits of multiple deployment models. One of the use cases for a hybrid cloud is cloud balancing whereby a collection of individual data centers appear to both users and administrators as a single cloud data center, with the physical location of application resources as transparent as possible. The goal of having the location of application resources be transparent creates a number of requirements. This includes:

- VLAN Extension
  Hybrid clouds depend heavily on VM migration among geographically dispersed servers. The VLANs within which VMs are migrated must be extended over the WAN between and amongst the private and public data centers. This involves the creation of an overlay network that allows the Layer 2 VLAN traffic to be bridged or tunneled through the WAN.

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

- Application Performance Optimization
  Application performance must meet user expectations regardless of the location of the users or the IT resources that the users are accessing.

The Virtualization of Networks and Network Functions

Unlike the factors discussed in the preceding section of **The Handbook**, neither Software Defined Networking (SDN) nor Network Functions Virtualization (NFV) are currently having a major impact on application and service delivery. However, both SDN and NFV have the potential in the near term to fundamentally impact application and service delivery in part by enabling the automated implementation of sophisticated forms of virtualized networks and virtualized network functions.

SDN

Network virtualization isn't a new topic. IT organizations have implemented various forms of network virtualization for years; i.e., VLANs, VPNs, VRF. However, in the context of SDN the phrase *network virtualization* refers to the creation of logical, virtual networks that are decoupled from the underlying network hardware to ensure the network can better integrate with and support increasingly virtual environments.

SDN has the potential to provide numerous benefits, including the ability to support the dynamic movement of VMs between physical servers without requiring any manual intervention.

There are two fundamental architectural approaches to create a SDN and the associated virtual networks. These two approaches are the:

- Overlay-based model;
- Fabric-based or underlay model.

The overlay-based model focuses on the hypervisor and it uses tunneling and encapsulation. The use cases associated with this model focus on the challenges and opportunities associated with virtual servers; i.e., support network virtualization, microsegmentation. Whereas the overlay-based model focuses on the hypervisor and uses tunneling and encapsulation, the underlay-based model focuses on a range of virtual and physical network elements and relies on the SDN controller manipulating flow tables in the network elements. The use cases that are associated with the underlay-based model are broader in scope than those that are associated with the overly-based model; i.e., support network virtualization, ease the burden of configuring and provisioning both physical and virtual network elements.

The initial discussion of SDN focused on the data center. However, as discussed in detail in the next chapter of **The Handbook**, there is a large and growing interest in implementing a software defined WAN (SD-WAN). As is the case with any software defined network, a SD-WAN centralizes the control function into a SDN controller. The controller abstracts the user's private network services from the underlying IP network and it enables the operation of the user's private network services via centralized policy (see below). The controller also enables the automation of management tasks such as configuration and provisioning.

Leveraging the underlying WAN platforms, which may include physical or virtual routers, the controller sets up virtual overlays that are both transport and technology agnostic. Under the direction of the

controller, the WAN platforms implement functionality such as quality of service, path selection, optimization and security, often using dynamic multi-pathing over multiple WAN links.

Network Functions Virtualization (NFV)

Many people associate NFV exclusively with service providers. That's understandable because the organizations that are most closely associated with the definition and development of NFV, such as the European Telecommunications Standards Institute (ETSI) and the TM Forum, focus almost exclusively on service providers.

The primary factors that are driving service providers to develop and implement NFV are the desire to be more agile in the implementation of new services and the desire to reduce cost, notably OPEX. Driven by those same factors, over the last few years enterprise network organizations have implemented virtualized versions of a range of L4 – L7 network functions including Application Delivery Controllers, WAN Optimization Controllers, Firewalls and Intrusion Detection/Prevention systems. Enterprise network organizations typically don't require the same scale solutions as does a service provider. However, enterprise network organizations require all of the same characteristics for the virtualized network functions they implement as are included in the ETSI vision for NFV.

The 2015 Guide to SDN and NFV contains the results of a survey in which over 200 survey respondents, most of whom work in enterprise IT organizations, were asked what they thought about the applicability of NFV. Only 5% of the respondents indicated that NFV is applicable only in a service provider environment. Eighty-two percent of the respondents indicated that NFV is either applicably equally in a service provider and enterprise environment or that it is applicable primarily in a service provider environment but that it does provide value in an enterprise environment.

The concepts and principles that are associated with NFV apply equally well in a service provide or enterprise environment.

An extensive discussion of SDN and NFV can be found in the 2015 Guide to SDN and NFV. This includes a discussion of the associated performance, management and security opportunities and challenges.

The Use of Policy

There is a broad movement to implement a policy based approach to all aspects of IT, including networking. Policies can be based on hierarchical system of rules designed to deal with the complexities of the environment, and to manage the relationships among users, services, SLAs, and device level performance metrics. One way that policy can be implemented is at the application level. For example, if the performance of an application begins to degrade because the CPU utilization of a physical server hosting a virtualized network function (VNF) that is used by that application becomes excessive, the VNF may be moved to a server with lower utilization, if that is in line with the policy that exists for that application. As was mentioned in the discussion of dynamic multi-pathing, another way to implement policy-based networking is to control which WAN link application traffic transits based in part on centralized policies that indicate among other things, the business criticality of that application.

The Survey Respondents were given a number of alternatives and asked to indicate which alternative best describes their company's interest in implementing a policy-based model in order to enhance the performance, management and/or security of their applications and services. Their responses are shown in **Table 3**.

Table 3: Approach to Implementing a Policy-Based Model	
Alternative	**Percentage of Respondents**
We have already begun to deploy such a model in production	27%
We have not evaluated such models, but are likely to evaluate them in the next year	18%
We have not evaluated such models and are unlikely to evaluate them in the next year	12%
We are currently trialing and/or testing such models from one or more vendors	8%
We are currently performing a paper evaluation of such models from one or more vendors	5%
We have already done a paper evaluation of such models and over the next year we will likely take steps towards implementation	4%
We have already evaluated such models and decided to not do anything with them at least for now	4%

There is broad interest in implementing a policy-based model in order to enhance the performance, management and/or security of their applications and services.

The Expectations of Business Unit Managers

If you looked inside of virtually any company a decade ago, the IT organization was either affectionately regarded as being the technology gurus, or perhaps less affectionately regarded as being the technology nerds. In either case, the vast majority of the company's employees didn't regard themselves as being tech savvy. In the current environment it is very common for a company's employees to have a lot of experience using IT functionality in their personal lives, and as a result, they consider themselves to be very tech savvy.

One of the biggest impacts of the growing use of IT functionality amongst a company employees is that it has dramatically changed the expectations of the company's business and functional managers. In a growing number of cases these managers don't want to be told that it will takes months for the IT organization to implement the functionality they need and are pushing IT organizations to become much more agile than they ever have been. If IT organizations can't keep up it is not much of a leap for managers who are culturally conditioned to downloading applications on their smart phone to bypass the IT organization and make use of public cloud computing solutions.

IT organizations either exhibit more agility or risk becoming irrelevant.

In the past, the way that business unit managers often dealt with the IT organization inability to meet their needs in a timely fashion was by building their own shadow IT organization. The business unit managers hired or assigned responsibility to people on their staff whose role was to provide the IT services that the business unit manager was unable to obtain from the IT organization. In the current environment, public cloud providers play the role of a shadow IT organization when a company's

business and functional managers go around the company's IT organization to obtain services or functionality that they either can't get internally or they can't get in a timely or cost effective manner.

In many cases public cloud providers play the role of a shadow IT function.

One way to relieve this pressure is for the IT organization to modify their traditional role of being the exclusive provider of IT services and to adopt a role in which they provide some IT services themselves and/or act as a broker between the company's business unit managers and cloud computing service providers. In addition to contract negotiations, the IT organization can add value by ensuring that the acquired application or service doesn't create any security or compliance issues, can perform well, can be integrated with other applications as needed, is scalable, cost effective and can be managed effectively and efficiently.

IT organizations need to play the role of honest broker between which applications and services are provided internally and which are acquired from a third party.

The evolution of the CMO and the CDO

In part to respond to the quickening pace of business and in part to respond to the ongoing digitization of business, CIOs, and the IT function that they manage are under intense pressure to show the business relevance of IT. While IT has always been under pressure to show business relevance, what has changed is that in a growing number of companies the role of the CIO is under attack from other C-level executives. One example of that trend was provided by Gartner who stated that by 2017 the CMO will spend more on IT than the CIO. In addition, driven by the market demand to transition from traditional bricks and mortar business models and adopt emerging digital business models, such as those adopted by Home Depot, a new type of C-Level executive is emerging: The Chief Digital Officer (CDO). The CDO is typically responsible for the development and management of the company's digital business models as well as the management and delivery of the company's digital assets. Starbucks is an example of a company with a CDO. The Starbucks' CDO, Adam Brotman, is responsible for Starbucks core digital businesses, including web, mobile, social media, card, loyalty, e-commerce and Wi-Fi."

The movement to Digital Business is both an opportunity and a threat to IT organizations.

The Adoption of DevOps

Since the phrase *DevOps* can be interpreted in many ways, to avoid confusion this e-book will use the following definition:

> ***DevOps*** *is a concept dealing with, among other things: software development, operations, and services. It emphasizes communication, collaboration, and integration between software developers* http://en.wikipedia.org/wiki/Software_developer*and information technology (IT) operations personnel. DevOps is a response to the interdependence of software development and IT operations. It aims to help an organization rapidly produce software products and services.*

According to a recent Information Week report, when asked to indicate the level of improvement in application development speed that they have either already gained or expected to gain as a result of adopting DevOps, forty-one percent of respondents indicated "significant improvement" and 42% indicated "some improvement".

The adoption of DevOps leads to more rapid application development.

A subsequent section of **The Handbook** will discuss how to apply the key principles of DevOps to increase the agility of network operations groups. Those principles include:

- Collaboration
  A key aspect of DevOps is to create a culture of collaboration amongst all the groups that have a stake in the delivery of new software.

- Continuous integration and delivery
  With continuous integration, software changes are added to a large code base immediately after development so that new capabilities can be continuously delivered to the entire release chain for testing and monitoring in production-style environments.

- Continuous testing and monitoring
  With DevOps, testing is performed continuously at all stages of the release process and not just by the QA organization. Developers do testing and provide test data and procedures that can be used by collaborating groups downstream in the process. The operations group is also typically involved in the test and monitoring processes. Part of their value add is that operations groups can specify load patterns to make testing by other groups more in line with actual usage conditions.

  Operations groups perform continuous monitoring to identify problems with the services being delivered so that they can be fixed in near real-time. Monitoring relies on an appropriate set of tools. The same tools that monitor the production environment can also be employed in development to identify performance problems prior to production deployment.

- Automation
  With DevOps all stages of software delivery are highly dependent on automated tools. Automation is essential because it enhances agility and provides the productivity required to support the continuous nature of integration, delivery, testing, and monitoring of many small increments to the code base.

- API centric automated management interfaces
  Software Defined Environments (SDEs) are an emerging core capability of DevOps that allow organizations to manage the scale and the speed with which environments need to be provisioned and configured to enable continuous delivery. SDEs use technologies such as API-centric automated management interfaces that define entire systems made up of multiple components. These interfaces are based on information models that define the characteristics, behaviors, configurations, roles, relationships, workloads, and work- load policies, for all the entities that comprise the system.

Network and Application Optimization

Key Optimization Tasks

The previous chapter of **The Handbook** discussed two surveys that were given in early 2015 to the subscribers of Webtorials. As previously noted, within **The Handbook** the respondents to those surveys will be referred to as The Survey Respondents. The Survey Respondents were given a number of optimization-related tasks and were asked to indicate how important it was for their IT organization to get better at each task at over the next year. Their responses are shown in **Table 4.**

Table 4: The Importance of Key Optimization Tasks

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

Table 4: The Importance of Key Optimization Tasks					
	Not at All	**Slightly**	**Moderately**	**Very**	**Extremely**
Controlling the cost of the WAN by reducing the amount of traffic by techniques such as compression	5.4%	9.7%	35.5%	34.4%	15.1%
Ensuring acceptable performance of traditional video traffic	3.3%	17.4%	27.2%	28.3%	23.9%
Optimizing the performance of applications and services acquired from public cloud providers	12.6%	20.7%	26.4%	28.7%	11.5%
Optimizing the transfer of virtual machines between data centers	14.6%	15.7%	33.7%	25.8%	10.1%
Optimizing the performance of chatty protocols such as CIFS	14.1%	17.6%	35.3%	22.4%	10.6%

Some of the conclusions that can be drawn from the data in **Table 4** are:

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

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

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

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

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

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

Traditional Optimization Appliances

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

WAN Optimization Controllers (WOCs)

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

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

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

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

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

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

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

Many IT organizations choose to implement a proof-of-concept (POC) trial prior to acquiring WOCs. The purpose of these trials is to enable the IT organization to quantify the performance improvements provided by the WOCs and to understand related issues such as the manageability and transparency of

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

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

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

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

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

Application Delivery Controllers (ADCs)

Background

The original purpose of an ADC was to provide load balancing across local servers or among geographically dispersed data centers based on Layer 4 through Layer 7 intelligence. By providing this functionality, an ADC maximized the efficiency and availability of servers through intelligent allocation of application requests to the most appropriate server. ADCs, however, have assumed, and will most likely continue to assume, a wider range of more sophisticated roles that enhance server efficiency and provide asymmetrical functionality to accelerate the delivery of applications from the data center to individual remote users. In particular, the ADC can allow a number of compute-intensive functions, such as SSL processing and TCP session processing, to be offloaded from the server. Server offload can increase the transaction capacity of each server and hence can reduce the number of servers that are required for a given level of business activity.

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

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

ADCs and Security

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

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

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

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

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

The Requirement for Programmability

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

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

Virtual ADCs

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

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

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

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

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

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

The Role of SDN

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

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

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

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

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

NFV Optimization

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

Performance bottlenecks are not unique to virtualized environments. For example, some of the bottlenecks that occur in a physical environment are shown in **Figure 5.**

Figure 5: Bottlenecks in a Physical Environment

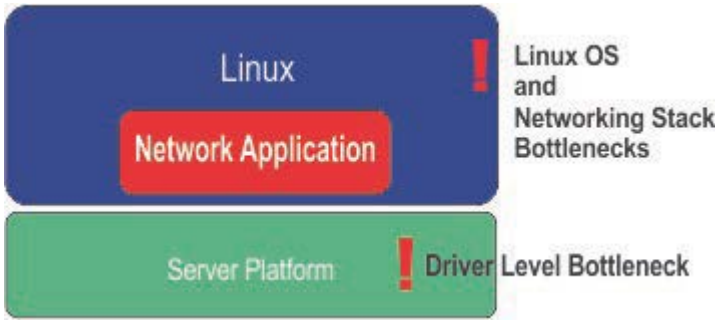

Figure 6: Bottlenecks in a Virtual Environment

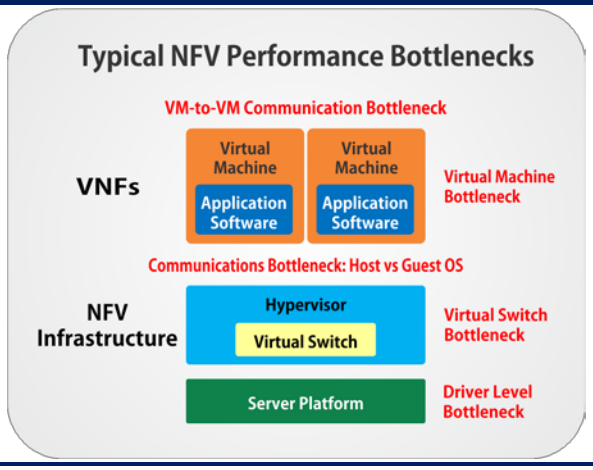

Unfortunately, as shown in **Figure 6**, as IT organizations adopt a virtualized environment the performance bottlenecks increase.

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

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

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

The Next Generation WAN

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

Background

WAN Evolution

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

In addition to the continued evolution of the Internet, the twenty-year period that began around 1984 saw the deployment of four distinct generations of enterprise WAN technologies. For example, in the mid to late 1980s, it became common for enterprise IT organizations to deploy integrated TDM-based WANs to carry both voice and data traffic. In the early 1990s, IT organizations began to deploy Frame Relay-based WANs. In the mid to late 1990s, some IT organizations replaced their Frame Relay-based WANs with WANs based on ATM (Asynchronous Transfer Mode) technology. In the 2000s, many IT organizations replaced their Frame Relay or ATM-based WANs with WANs based on MPLS. Cost savings was the primary factor that drove the adoption of each of the four generations of WAN technologies.

WAN Services

As discussed in The 2014 State of the WAN Report, network organizations currently make relatively little use of WAN services other than MPLS and the Internet and the use they do make of those other services is decreasing somewhat rapidly. That report also identified the concerns that network organizations have with those two services. Those concerns are shown in **Table 5** in descending order of importance.

Table 5: Concerns with WAN Services

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

Traditional WAN Design

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

One alternative to the traditional approach to designing a branch office WAN is to supplement the T1 access link(s) in a branch office with direct Internet access and to also leverage technology such as Policy Based Routing (PBR). PBR allows network administrators to create routing policies to allow or deny paths based on factors such as the identity of a particular end system, the protocol or the application.

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

Factors Driving Change in the WAN

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

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

Figure 7: Factors Driving Change in the WAN

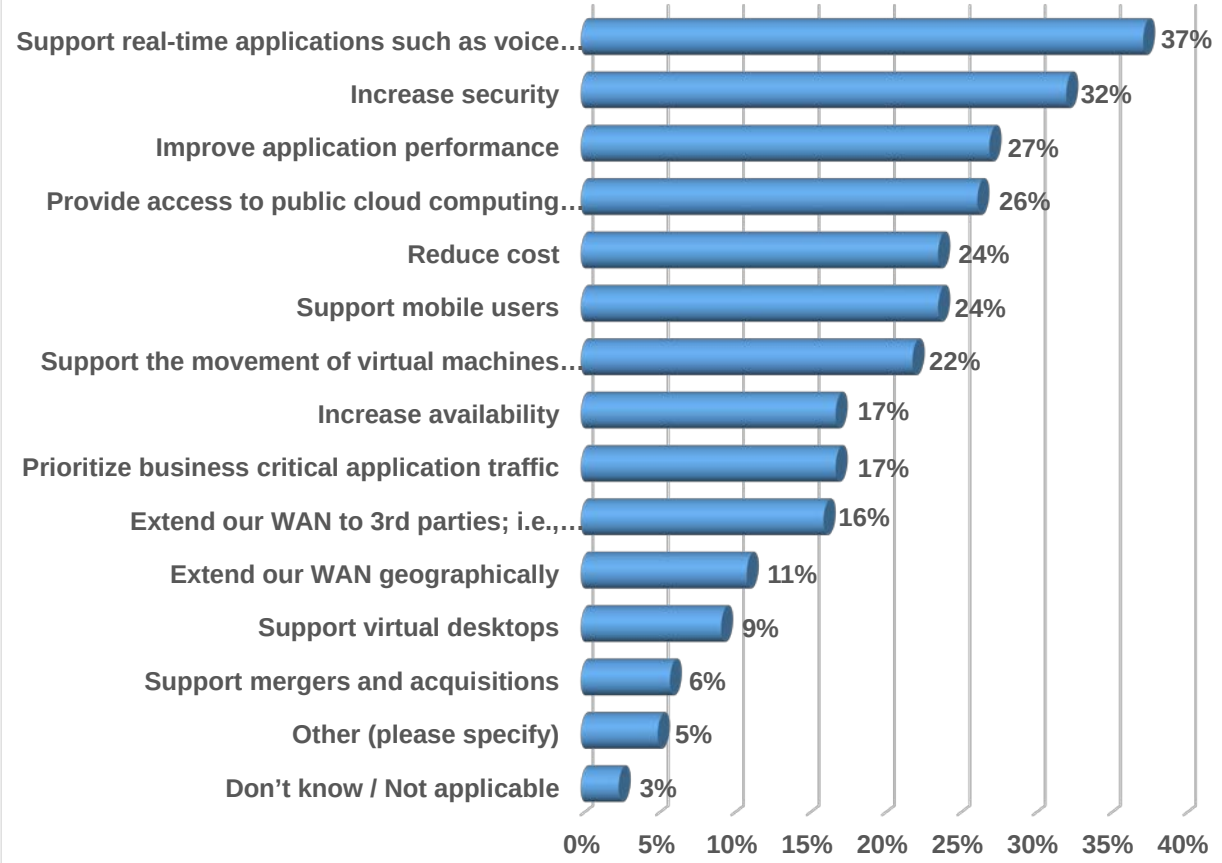

The data in **Figure 7** indicates that there is a wide range of WAN challenges that are important for network organizations to respond to.

Software Defined WANs

Hybrid WAN

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

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

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

Interest in Leveraging SDN in the WAN

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

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

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

The respondents to this survey further demonstrated their interest in implementing SDN in the WAN when they indicated how broadly they expected their campus, WAN and data center networks would be based on SDN three years from now. Their responses (**Table 6**) show that IT organizations believe that three years from now that SDN deployment in data centers will be highly pervasive and that there will also be significant SDN deployment both in the WAN and in campus networks.

Table 6: Anticipated SDN Deployment

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

Definition of a Software Defined WAN

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

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

Over half of The Survey Respondents (**Figure 8**) indicated that they either just learned about a SD-WAN from the definition that was in the survey instrument or they were just somewhat familiar with the concept. This lack of familiarity isn't surprising given that a SD-WAN is an emerging concept. It does, however, highlight the need for more education on this topic.

Figure 8: Familiarity with SD-WAN

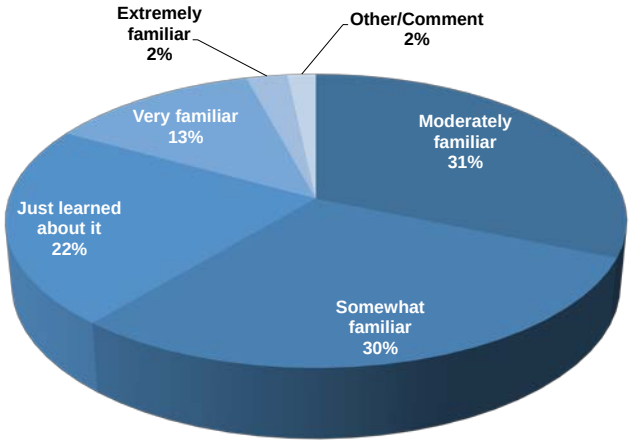

Drivers and Inhibitors of SD-WAN Adoption

The Survey Respondents were given a set of possible outcomes and were asked to indicate which outcomes would drive their company to implement a SD-WAN. Their responses are shown in **Table 7**.

Table 7: Drivers of SDN WAN Adoption

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

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

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

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

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

The Survey Respondents were also given a set of factors and were asked to indicate which factors would inhibit their company from implementing a SD-WAN. Their responses are shown in **Table 8**.

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

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

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

Management & Security

Management

Market Research

The first chapter of **The Handbook** discussed two surveys that were given in early 2015 to the subscribers of Webtorials. As previously noted, within **The Handbook** the respondents to those surveys will be referred to as The Survey Respondents.

Table 9 shows how The Survey Respondents answered a survey question about the management tasks that their IT organizations are most interested in getting better at over the next year.

Table 9: 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%	5.7%	14.3%	36.2%	43.8%
Identify the components of the IT infrastructure that support the company's critical business applications	1.9%	1.9%	18.3%	42.3%	35.6%
Obtain performance indicator metrics and granular data that can be used to detect and eliminate impending problems	0.9%	6.6%	16.0%	45.3%	31.1%
Monitor the end user's experience and behavior	0.9%	8.4%	19.6%	44.9%	26.2%
Effectively manage SLAs for one or more business critical applications	1.0%	8.6%	13.3%	33.3%	43.8%
Manage the use of VoIP	5.9%	12.87%	32.7%	23.8%	24.8%
Perform traditional management tasks such as troubleshooting and performance management, on a per VM basis	2.8%	2.8%	33.0%	41.5%	19.8%
Monitor and manage the performance of applications delivered to mobile users	3.9%	14.7%	22.6%	30.4%	28.4%
Manage end-to-end in a public cloud computing environment	10.3%	18.6%	17.5%	32.0%	21.7%
Manage end-to-end in a private cloud computing environment	7.2%	11.3%	19.6%	32.0%	29.9%

Some of the conclusions that can be drawn from the data in **Table 9** include:

The most important management tasks to get better at over the next year are:

- ***Rapidly identifying the root cause of degraded application performance;***
- ***Effectively managing SLAs for one or more business critical applications;***
- ***Identifying the components of the IT infrastructure that support the company's critical business applications;***
- ***Obtaining performance indicator metrics and granular data that can be used to detect and eliminate impending problems.***

A relatively new management task, monitor and manage the performance of applications delivered to mobile users, continues to increase in importance.

Existing Trends That Impact Management

Chapter 1 of **The Handbook** described the emerging service and application delivery challenges. This subsection will identify how some of the existing trends are forcing a change in terms of how IT organizations manage applications and services; e.g., how do critical tasks, such as identifying the root cause of degraded application performance, need to change as the IT environment changes?

Server Virtualization

Until recently, IT management was based on the assumption that the IT organization performed tasks such as monitoring, baselining and troubleshooting on a server-by-server basis. Now, as highlighted by the data in **Table 1**, 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. The fact that VMs migrate between physical servers is one of the reasons why so many of The Survey Respondents indicated that it is important to their organization to get better at identifying the components of the IT infrastructure that support the company's critical business applications.

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 Computing

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 varying forms of cloud computing (i.e., private, public, hybrid) demonstrates that:

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.

Real-Time Applications

Voice and video are examples of real-time applications that have high visibility and which are sensitive to transmission impairments. Chapter 2 contained survey data that showed how important supporting real time applications across the WAN was to The Survey Respondents. In addition **Table 1**, shows that 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.

Effectively managing voice and video requires looking at the application payload and measuring the quality of the voice and video communications.

In the case of Unified Communications (UC), effective management also requires monitoring the signaling between the components of the UC solution.

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:

A 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.

Emerging Trends That Impact Management

Because of the breadth and depth of their potential impact, this subsection will look at the management issues brought about by the adoption of Software Defined Networks (SDN) and Network Functions Virtualization (NFV).

SDN

One of the promises of SDN is that it will ease the administrative burden of management tasks such as configuration and provisioning. However:

In SDN environments the challenges associated with end-to-end service management are more demanding than they are in traditional network environments.

This follows in part because in a SDN environment there is a need to monitor additional components, such as SDN controllers, in an environment that is a combination of physical and virtual resources and which is changing dynamically.

SDN creates both management opportunities and management challenges.

One of the management challenges that applies across multiple tiers of the SDN architecture is the requirement to manage the messaging that goes between tiers; e.g., between the application tier and the control tier as well as between the control tier and the infrastructure tier. At the infrastructure tier, one of the primary challenges is to perform element management potentially of both virtual and physical network elements. One of the management challenges at the control layer results from the fact that 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 into how the SDN controller is processing flows.

One of the management challenges that occurs at the application tier is that based on the type of application (e.g., business application vs. a firewall) the service or application needs varying levels of visibility into the underlying network. Another set of management challenges that occurs at the application layer stem from the requirement to ensure acceptable performance. One thing this means is that:

Network infrastructure must have visibility into the SLA requirements of the application so that when faced with a spike in demand, a policy-based decision can be made as to whether or not resources should be dynamically allocated to meet those demands.

Looking at network virtualization as an application of SDN, another performance management challenge stems from the fact that one of the primary benefits of overlay-based SDN solutions is the ability to support multiple virtual networks that run on top of a physical network. In order to perform management functions such as root cause analysis and impact analysis, network management organizations need the ability to see the bilateral mapping between the virtual networks and the physical network that supports them.

NFV

The adoption of NFV poses a number of significant challenges that must be overcome in order to ensure the ability to continue to implement effective end-to-end management. In recognition of that fact, the European Telecommunications Standards Institute (ETSI) has established a management and orchestration framework for NFV entitled Network Function Virtualization Management and Orchestration. Some of the key concepts contained in that framework were summarized in an ETSI document. According to that document:

> "In addition to traditional Fault, Configuration, Accounting, Performance, and Security (FCAPS) Management, the NFV Management and Orchestration framework introduces a new set of management functions associated with the lifecycle management of a VNF. The NFV ISG has focused on detailing these new sets of management functions, which include, but are not limited to: on-board a VNF, instantiate a VNF, scale a VNF, update a VNF, and terminate a VNF. A difference also worth highlighting relates to fault and performance management - in a virtualized environment this is the responsibility of different functional blocks at different layers. As a result, the correlation of faults, alarms and other monitored data such as performance metrics and resource usage, and the consequent fault resolution needed to operate the service in a reliable manner, will typically be distributed.

This subsection of **The Handbook** expands on some of the key NFV-related management challenges that ETSI and others are working to address.

Dynamic relationships between software and hardware components

In traditional networks, application software and network function software generally run on dedicated hardware that is statically provisioned by manual processes. With the current approach to virtualization, software running on virtual machines (VMs) can readily be moved among physical servers or replicated to run on newly created VMs in order to dynamically maintain availability, expand/shrink capacity, or balance the load across physical resources. Many of these changes in the infrastructure can be automated and programmatically activated to conform to configured policies under specific sets of circumstances.

Due to the mobility of VMs, topology changes can occur in a matter of seconds or minutes rather than the days or weeks required for changing software/hardware relationships in traditional networks. In order to accommodate and leverage virtualization technologies:

End-to-end management systems need to be re-architected to be capable of implementing automated processes for virtual resource procurement, allocation, and reconfiguration in accordance with a set of highly granular policies designed to ensure the quality of experience for the user of the network services.

Dynamic changes to physical/virtual device configurations

To accommodate the dynamic nature of virtualized networks, end-to-end management systems will need to be able to adjust the configuration of devices to react to changing conditions in the network.

SDN holds the potential to enable IT organizations to dynamically change the environment in order to meet SLAs.

Many-to-Many relationships between network services and the underlying infrastructure

In a traditional network infrastructure there is 1-to-1 relationship between a network service and a set of dedicated physical resources. In a virtualized infrastructure a network service can be supported by a number of Virtualized Network Function (VNFs) which may be running on one or several VMs. A single VNF may also support a number of distinct network services. In addition, the group of VNFs supporting a single network service could possibly be running on a number of distinct physical servers. As a result:

End-to-end management systems need to support a three-tiered network model based on many-to-many relationships among network services, virtualization infrastructure, and physical infrastructure.

Hybrid physical/virtual infrastructures

As virtualization is gradually adopted, IT organizations will need to be able to integrate virtual environments into their existing end-to-end traditional/legacy monitoring infrastructures. Therefore:

End-to-end management systems developed for the virtual infrastructure will need to be compatible with legacy infrastructure.

DevOps

Chapter 1 of **The Handbook** defined the term *DevOps* and discussed some of its key principles. All of the key principles of DevOps are applicable in a network operations (NetOps) setting. However DevOps is generally applied to discreet services that are frequently delivered over the web on a best effort basis. The network environment is different than that and as a result virtualized network services development creates challenges that are not addressed by DevOps.

The approach that most IT organizations take to DevOps needs to be modified before it can be applied to NetOps.

One challenge that distinguishes NetOps from DevOps is that since VNFs such as optimization and security are chained together to create an end-to-end service this creates strong dependencies between the VNFs. For example, if an IT organization updates an optimization VNF they need to ensure that it is fully compatible with the security VNF(s). As a result much stronger version control and compatibility testing is needed than would be typical for enterprise applications.

Other challenges created by network services development that must be addressed by NetOps that were not addressed by DevOps include:

- Since for the foreseeable future the vast majority of environments will be a combination of hardware-based and software-based functionality, the NetOps methodology must accommodate services that depend on network functions running on dedicated hardware platforms as well as on virtualized platforms.

- Virtualized services will often be created by integrating services from multiple suppliers. This will require NetOps methodologies and best practices to support concurrent synchronized development and integration across the domains of multiple partners.

- Unlike what happens when delivering an application over the Web, NetOps will need to support dynamic and automated management of service performance and SLAs. This can only be achieved by a policy model that supports end-to-end SLA targets.

- Again in contrast to what often happens when delivering an application over the Web, NFV services are often mission critical. This creates a need for high levels of resilience and rapid fallback capabilities.

Security

The Changing Security Environment

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

IT security systems and policies have evolved and developed around the traditional application delivery architecture in which branch offices users are connected to application servers in a central corporate data centers by using an enterprise WAN service such as MPLS. 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 and hybrid cloud computing, applications and services are moving out of the central corporate data center and there is no longer a well-agreed to location for security policies and systems. This topic is explored in detail in The 2015 Guide to WAN Architecture and Design.

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. Unfortunately, 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).

In 2014 the department store Target announced that thieves had stolen massive amounts of credit and debit card information and they also stole the names, addresses and phone numbers of 70 million of Target's customers. As a result of the security breach, Target's profits dropped by almost 50%[11]. As

[11] http://www.forbes.com/sites/maggiemcgrath/2014/02/26/target-profit-falls-46-on-credit-card-breach-and-says-the-hits-could-keep-on-coming/

sometimes happens when there is a breach of this magnitude, Target fired their CIO[12]. However, because of the impact on profits, Target also fired their CEO[13].

Security breaches can have a very negative impact on the career of CIOs and CEOs.

Existing Trends That Impact Security

The *IBM X-Force Threat Intelligence Quarterly, 1Q 2015* identified some of the key security-related trends. Some of the trends that IBM identified are:

- Target was not the only company that was hacked in 2014. The total number of leaked records (i.e., emails, credit card numbers, passwords and other personally identifiable information) continued to increase on an annual basis. It was a billion leaked records in 2014 which was an increase of 25% over the 800 million records that were leaked in 2013.

- In addition to an overall concern about breaches, security incidents and malware, last year mobile devices were shown to present some unique security vulnerabilities. For example, in 2014 a Computer Emergency Readiness Team-Coordination Center (CERT/CC) researcher discovered security issues in thousands of Android applications. These vulnerabilities can allow an attacker to perform man-in-the-middle attacks against affected mobile applications.

- In 2014, the underlying libraries that handle cryptographic functionality on nearly every common web platform, including Microsoft Windows, Mac OS X and Linux, were found to be vulnerable to fairly trivial remote exploitations capable of stealing critical data.

- A family of vulnerabilities affecting cryptographic systems which was named Padding Oracle on Downgraded Legacy Encryption (POODLE) was discovered in 2014. These vulnerabilities allows attackers to perform a man in the middle attack to silently intercept a secure session.

The *IBM X-Force Threat Intelligence Quarterly, 1Q 2015* presented survey data that identified the percentage of the totality of security incidents in 2014 that were attributable to a particular type of security attack. The data in the IBM report is shown in **Table 10**.

[12] http://www.forbes.com/sites/howardbaldwin/2014/03/11/the-other-shoe-drops-for-targets-cio/
[13] http://www.forbes.com/sites/ericbasu/2014/06/15/target-ceo-fired-can-you-be-fired-if-your-company-is-hacked/

Table 10: Most Common Security Attacks in 2014	
Type of Attack	**Percentage**
Undisclosed	40.2%
Malware	17.2%
DDoS	17.2%
SQLi	8.4%
Phishing	4.6%
Watering Hole	4.2%
Misconfiguration	3.4%
Brute Force	1.9%
Cross-site Scripting	0.8%
Heartbleed	0.8%

One of the conclusions that can be drawn from **Table 10** is that:

Malware and DDoS attacks are the two most dominant forms of a security attack.

In order to identify the degree to which the types of security attacks shown in Table 2 are of concern to IT organizations, The Survey Respondents were asked to indicate how important it was to their organization that over the year that they get better at defending against each type of attack. Their responses are shown in **Table 11**.

Table 11: The Importance of Getting Better at Defending Against Specific Security Attacks					
	Not at All	**Slightly**	**Moderately**	**Very**	**Extremely**
Malware	0.0%	7.1%	7.1%	40.4%	45.5%
DDoS	0.0%	5.2%	17.7%	29.2%	47.9%
Phishing	0.0%	6.3%	18.9%	36.8%	37.9%
Misconfigurations	0.0%	5.2%	23.7%	38.1%	33.0%
Cross-site scripting	1.1%	7.5%	21.5%	39.8%	30.1%
Heartbleed	2.2%	7.5%	26.9%	36.6%	26.9%
Brute force	1.1%	7.6%	27.2%	40.2%	23.9%
Watering hole	3.4%	9.1%	30.7%	35.2%	21.6%
SQL injections	3.3%	12.0%	22.8%	35.9%	26.1%

Some of the conclusions that can be drawn from **Table 11** include:

While some forms of a security attack (i.e., a malware attack) are more concerning than other forms (i.e., SQL injections), in general IT organizations feel that they need to get better at thwarting a wide range of security attacks.

Emerging Trends That Impact Security

Because of the breadth and depth of their potential impact, this subsection will look at the management issues brought about by the adoption of Software Defined Networks (SDN) and Network Functions Virtualization (NFV).

SDN

There are many ways that SDN can enhance security. For example, role based access can be implemented by deploying a role-based resource allocation application that leverages the control information and capability of the SDN controller. Another example is that by virtue of Layer 2-4 flow matching capability, OpenFlow access switches can perform filtering of packets as they enter the network, acting as simple firewalls at the edge. With OpenFlow switches that support modification of packet headers, an OpenFlow-enabled controller will also be able to have the switch redirect certain suspicious traffic flows to higher-layer security controls, such as IDS/IPS systems, application firewalls, and Data Loss Prevention (DLP) devices. Other security applications built on OpenFlow controller can match suspicious flows to databases of malware signatures or divert DDoS attacks.

Some of the security challenges related to SDN are described in *SDN Security Considerations in the Data Center*. As pointed out in that document:

- The centralized controller emerges as a potential single point of attack and failure that must be protected from threats;
- The southbound interface between the controller and underlying networking devices is vulnerable to threats that could degrade the availability, performance, and integrity of the network;
- The underlying network infrastructure must be capable of enduring occasional periods where the SDN controller is unavailable, yet ensure that any new flows will be synchronized once the devices resume communications with the controller.

Similar to the situation with management, SDN creates both security opportunities and security challenges.

Some security-related considerations that IT organizations should consider include:

- Implement measures to deal with possible control flow saturation (controller DDoS) attacks;
- Harden the SDN controller's operating system to ensure availability of the controller function;
- Implement effective authentication and authorization procedures that govern operator access to the controller.

NFV

A number of organizations are focused on resolving the security issues associated with SDN and NFV. One such organization is the Internet Engineering Task Force (IETF). The IETF has created a security architecture that is based on horizontal (a.k.a., east/west) APIs in addition to the northbound and

southbound APIs. One IETF SDN-specific activity focuses on centralized security services (i.e., firewalls and DDoS mitigation systems) designed specifically for SDN environments. Another SDN-specific Internet draft addresses the possible application of DevOps principles to SDNs.

ETSI is another organizations focused on resolving the security issues associated with SDN and NFV. In a document entitled *Network Functions Virtualization (NFV); NFV Security; Security and Trust Guidance*, ETSI outlined some high level security goals for NFV. According to that ETSI document:

> The dynamic nature of Network Function Virtualization demands that security technologies, policies, processes and practices are embedded in the genetic fabric of NFV. Additional high-level security goals for NFV include:
>
> - Establish a secured baseline of guidance for NFV operation, while highlighting optional measures that enhance security to be commensurate with risks to confidentiality, integrity and availability;
>
> - Define areas of consideration where security technologies, practices and processes have different requirements than non-NFV systems and operations;
>
> - Supply guidance for the operational environment that supports and interfaces with NFV systems and operations, but avoid redefining any security considerations that are not specific to NFV.

The ETSI document also summarizes a number of NFV security-related use cases.

Conclusions

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

- The goals of the 2015 Application and Service Delivery Handbook are to help IT organizations understand the emerging application and service delivery environment and to effectively respond to that environment.
- IT organizations need to plan for performance, security and management in an integrated fashion.
- As a minimum, the IT function needs to enable rapid business change. Ideally, the IT function is perceived as a driver of that change.
- 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.
- Getting better at improving the performance of applications and services acquired from a public cloud provider is either very or extremely important to well over a third of IT organizations.
- Getting better at managing end-to-end in a public cloud environment is either very or extremely important to the majority of IT organizations.
- 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.
- SDN has the potential to provide numerous benefits, including the ability to support the dynamic movement of VMs between physical servers without requiring any manual intervention.
- The concepts and principles that are associated with NFV apply equally well in service provide and enterprise environments.
- There is broad interest on the part of IT organizations to implement a policy-based model in order to enhance the performance, management and/or security of their applications and services.
- IT organizations either exhibit more agility or risk becoming irrelevant.
- In many cases public cloud providers play the role of a shadow IT function.
- IT organizations need to play the role of honest broker between what applications and services are provided internally and which are acquired from a third party.

- The movement to become a digital business is both an opportunity and a threat to IT organizations.

- The adoption of DevOps leads to more rapid application development.

- The most important management tasks to get better at over the next year are:

    - Rapidly identifying the root cause of degraded application performance;
    - Effectively managing SLAs for one or more business critical applications;
    - Identifying the components of the IT infrastructure that support the company's critical business applications;
    - Obtaining performance indicator metrics and granular data that can be used to detect and eliminate impending problems.

- A relatively new management task, monitor and manage the performance of applications delivered to mobile users, continues to increase in importance.

- 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 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.

- Effectively managing voice and video requires looking at the application payload and measuring the quality of the voice and video communications.

- A converged infrastructure requires a management system and management processes that have the same level of integration and cross-domain convergence that the infrastructure has.

- In SDN environments the challenges associated with end-to-end service management are more demanding than they are in traditional network environments.

- SDN creates both management opportunities and management challenges.

- Performance management systems need visibility not only into application performance but also into how the SDN controller is processing flows.

- Network infrastructure must have visibility into the SLA requirements of the application so that when faced with a spike in demand, a policy-based decision can be made as to whether or not resources should be dynamically allocated to meet those demands.

- End-to-end management systems need to be re-architected to be capable of implementing automated processes for virtual resource procurement, allocation, and reconfiguration in accordance with a set of highly granular policies designed to ensure the quality of experience for the user of the network services.

- SDN holds the potential to enable IT organizations to dynamically change the environment in order to meet SLAs.

- End-to-end management systems need to support a three-tiered network model based on many-to-many relationships among network services, virtualization infrastructure, and physical infrastructure.
- End-to-end management systems developed for the virtual infrastructure will need to be compatible with legacy infrastructure.
- The approach that most IT organizations take to DevOps needs to be modified before it can be applied to NetOps.
- 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.
- Security breaches can have a very negative impact on the career of CIOs and CEOs.
- Malware and DDoS attacks are the two most dominant forms of a security attack.
- While some forms of a security attack (i.e., a malware attack) are more concerning than other forms (i.e., SQL injections), in general IT organizations feel that they need to get better at thwarting a wide range of security attacks.
- Similar to the situation with management, SDN creates both security opportunities and 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
Steven Taylor

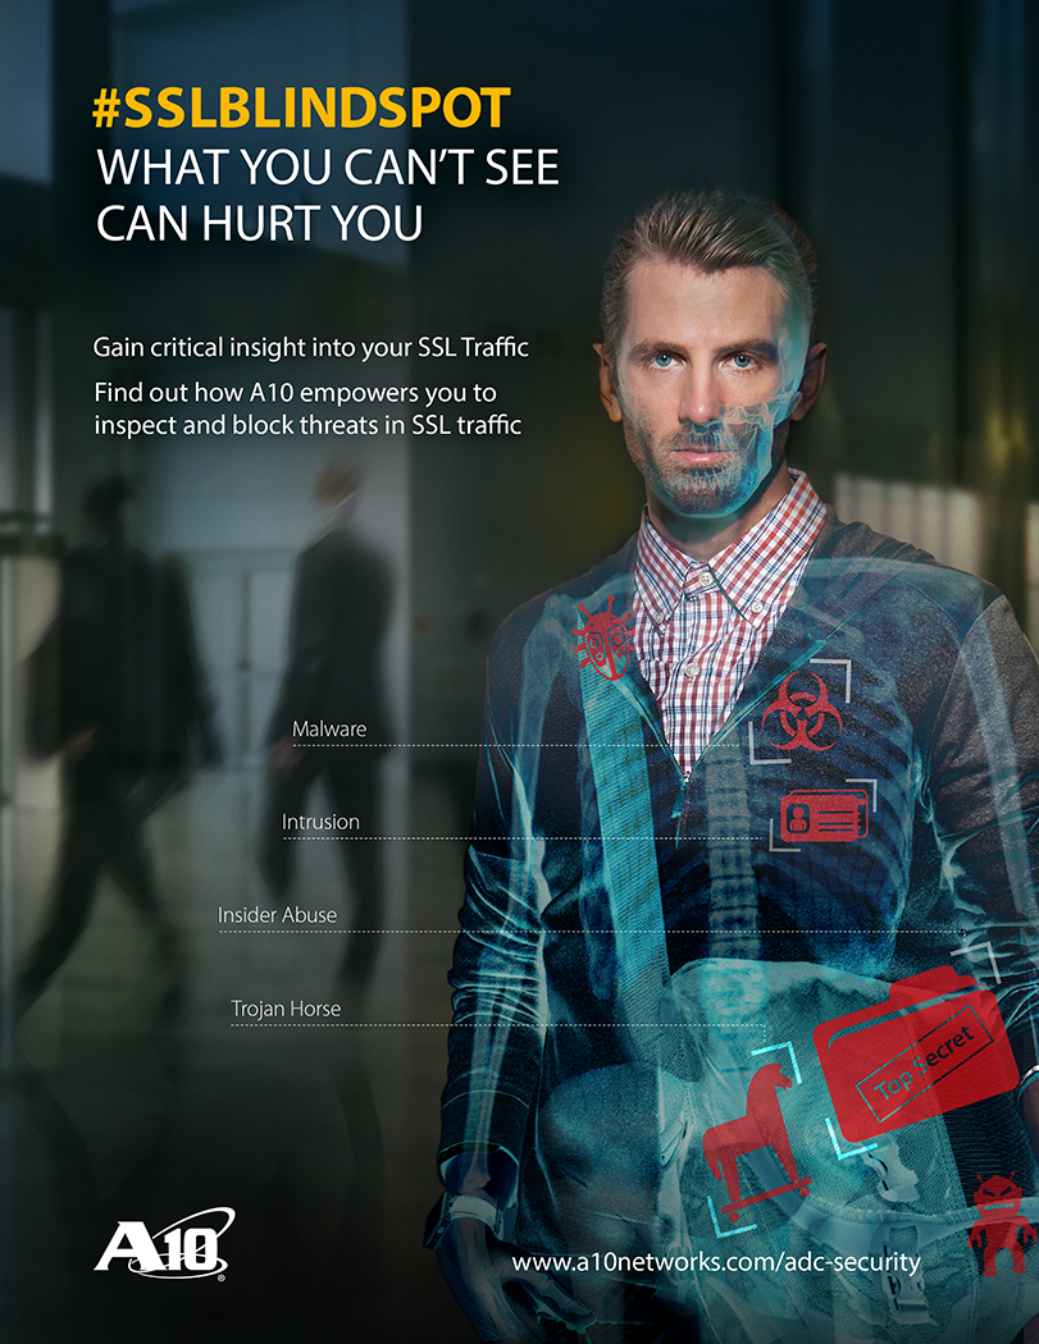
#SSLBLINDSPOT
WHAT YOU CAN'T SEE
CAN HURT YOU
Gain critical insight into your SSL Traffic
Find out how A10 empowers you to
inspect and block threats in SSL traffic
Malware
Intrusion
Insider Abuse
Trojan Horse
Top Secret
A10
www.a10networks.com/adc-security

SDN Today:

Delivered by Citrix NetScaler and Cisco ACI

Learn more at citrix.com/netscaler/cisco

CITRIX®

cisco™

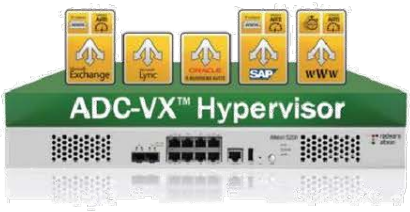

Predictable Application Service Levels, Guaranteed—Only with Alteon NG

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

The Standard ADC: Not Good Enough Anymore

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

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

Alteon NG: Complete Application Service Level Assurance

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

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

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

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

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

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

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

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

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