The 2015
Application & Service Delivery Handbook

Executive Summary

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

Platinum Sponsors:

Gold Sponsors:

Produced by:

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/

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

Professional Opinions Disclaimer

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

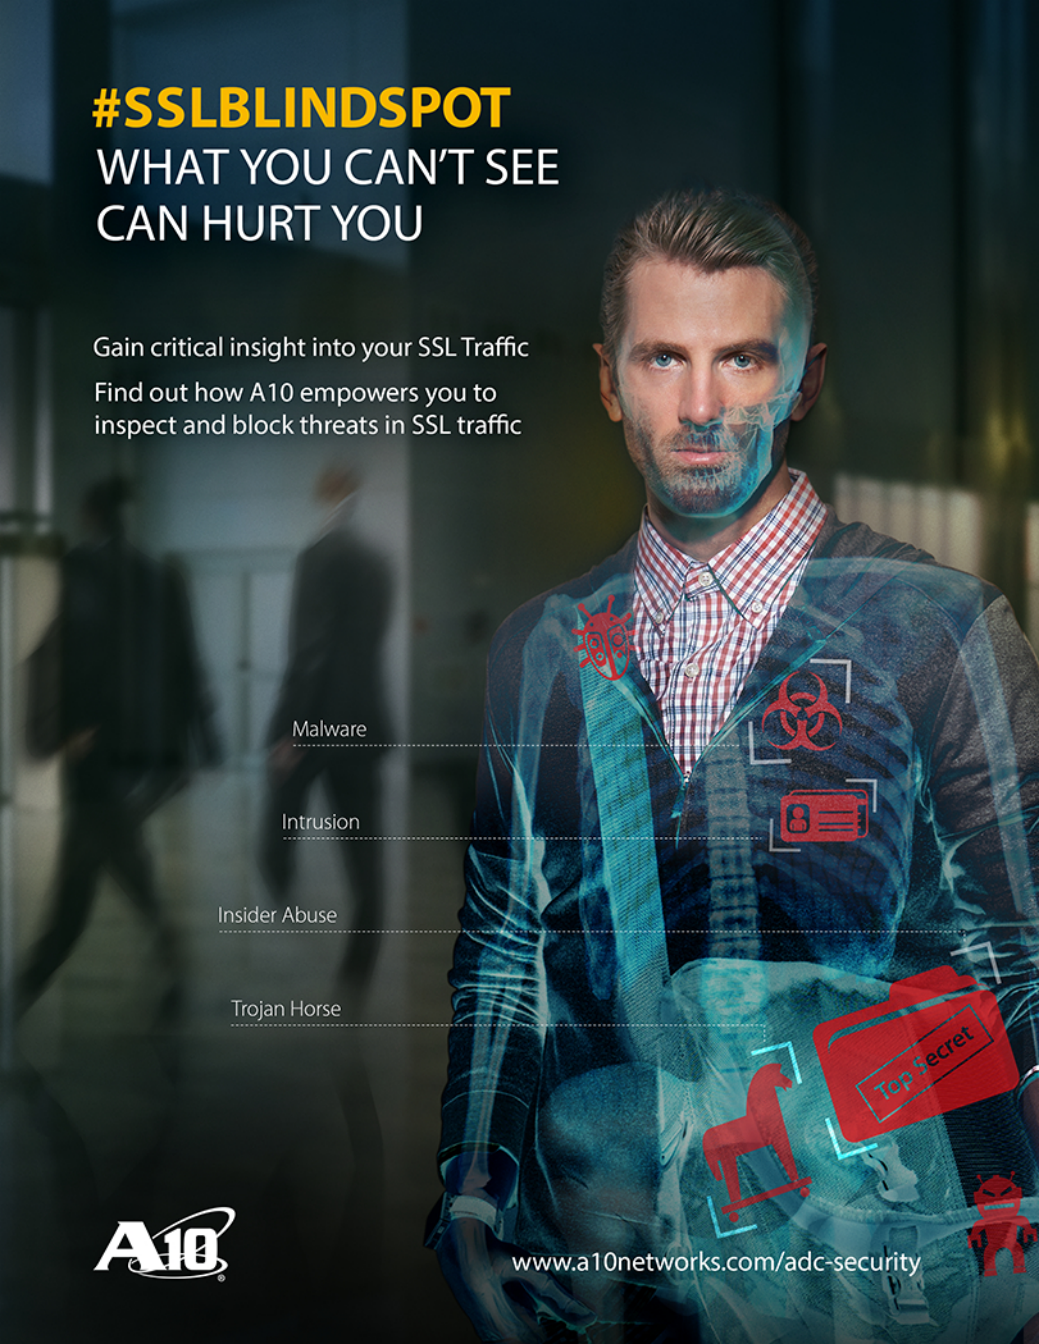
#SSLBLINDSPOT
WHAT YOU CAN'T SEE
CAN HURT YOU
Gain critical insight into your SSL Traffic
Find out how A10 empowers you to
inspect and block threats in SSL traffic
Malware
Intrusion
Insider Abuse
Trojan Horse
Top Secret
A10
www.a10networks.com/adc-security

SDN Today:

Delivered by Citrix NetScaler and Cisco ACI

Learn more at citrix.com/netscaler/cisco

CITRIX®

cisco™

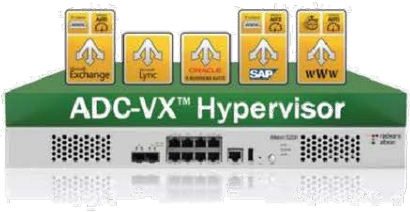

Predictable Application Service Levels, Guaranteed—Only with Alteon NG

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

The Standard ADC: Not Good Enough Anymore

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

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

Alteon NG: Complete Application Service Level Assurance

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

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

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

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

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

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

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

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

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