The 2015
Application & Service Delivery Handbook

Part 3: Management and Security

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

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

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

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 1 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 1: 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 1** 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%[1]. As

[1] 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[2]. However, because of the impact on profits, Target also fired their CEO[3].

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 2**.

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

Table 2: 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 2** 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 3**.

Table 3: 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 3** 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.

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.

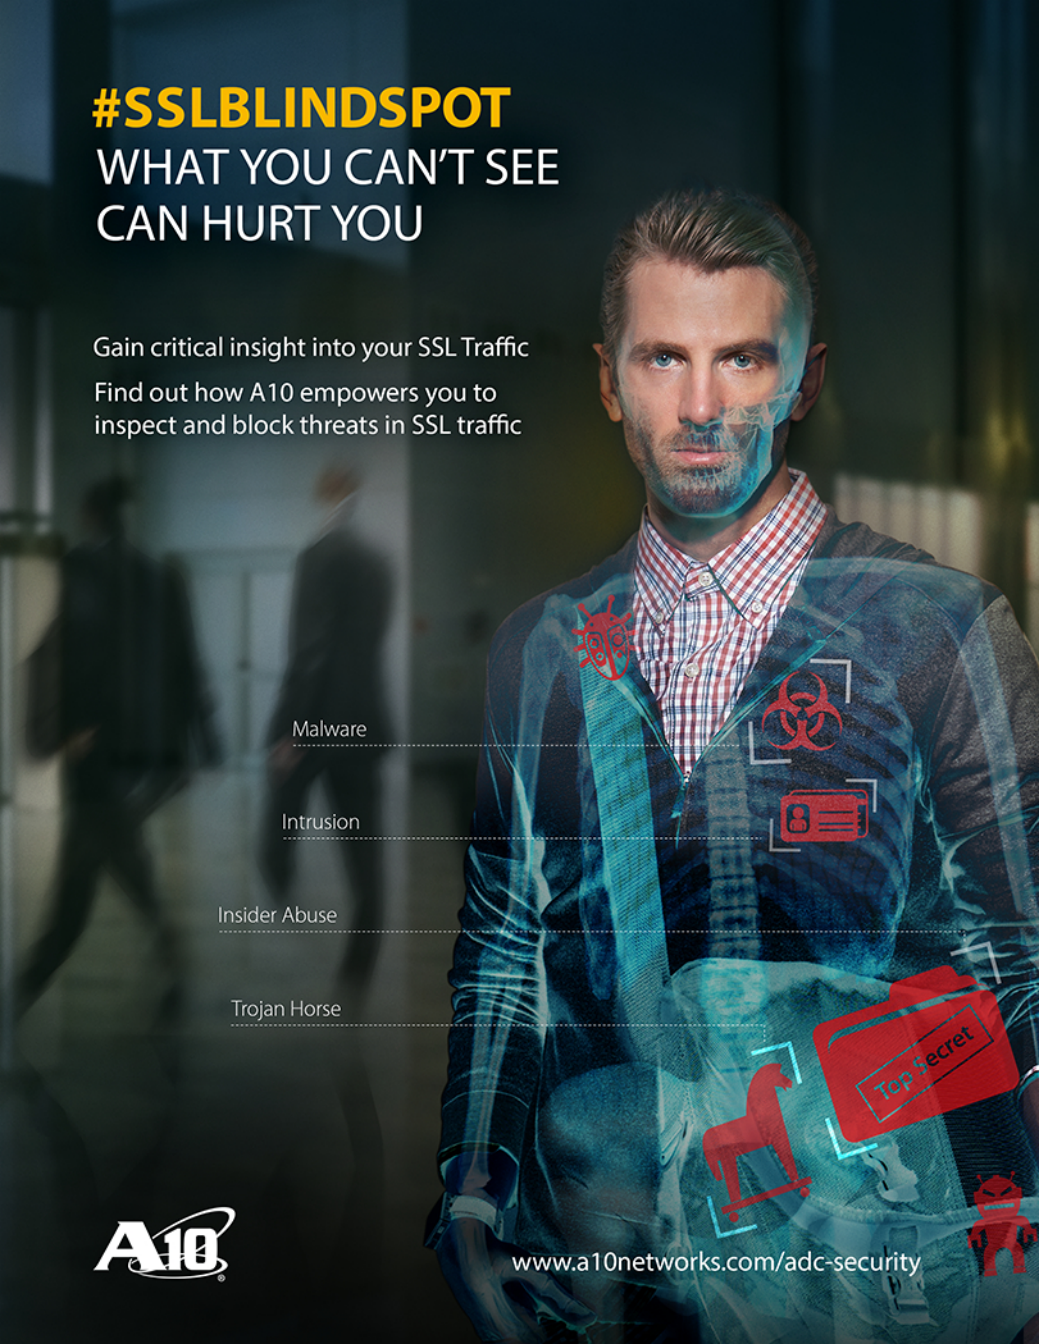
#SSLBLINDSPOT
WHAT YOU CAN'T SEE
CAN HURT YOU
Gain critical insight into your SSL Traffic
Find out how A10 empowers you to
inspect and block threats in SSL traffic
Malware
Intrusion
Insider Abuse
Trojan Horse
Top Secret
A10
www.a10networks.com/adc-security

SDN Today:

Delivered by Citrix NetScaler and Cisco ACI

Learn more at citrix.com/netscaler/cisco

CITRIX®

cisco™

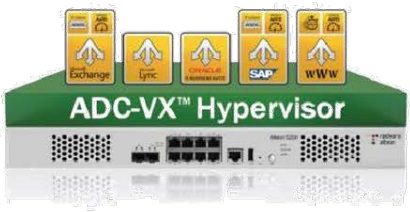

Predictable Application Service Levels, Guaranteed—Only with Alteon NG

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.