The 2015
Application & Service Delivery Handbook

Part 1: The Emerging Application and Service Delivery Environment

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*** will be published both in its entirety and in a serial fashion. This is the first of the serial publications. The primary goal of this publication is to describe how the application and service delivery environment is changing and the challenges and opportunities that the changing environment creates. Subsequent publications of the ***2015 Application and Service Delivery Handbook*** will focus on describing the technologies, products and services that are available to improve the:

- Performance of applications and services.
- 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 to understand the emerging application and service delivery environment and to effectively respond to that environment.

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.

IT organizations need to plan for performance, security and management in an integrated fashion.

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 The Handbook, 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 & 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."

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[2], 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.

[1] http://www.bbc.co.uk/news/business-16611040
[2] 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

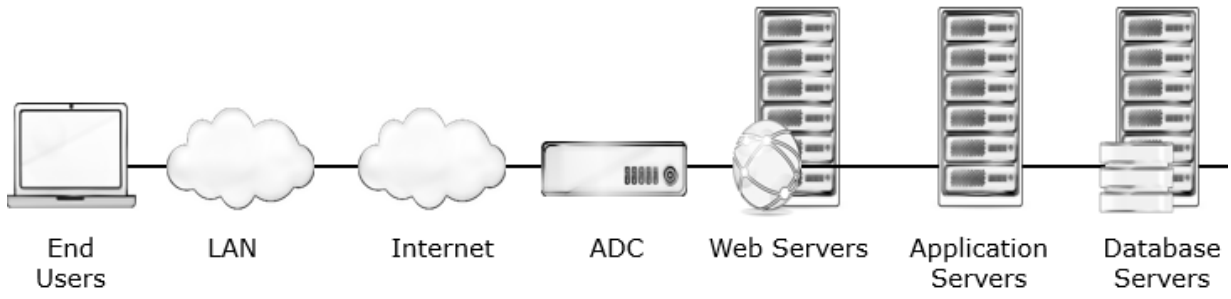

Most of the performance challenges associated with the client server model are included in the previously mentioned list of traditional application and service delivery challenges; i.e., network latency; chatty protocols and applications. One of the key management challenges associated with this delivery 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

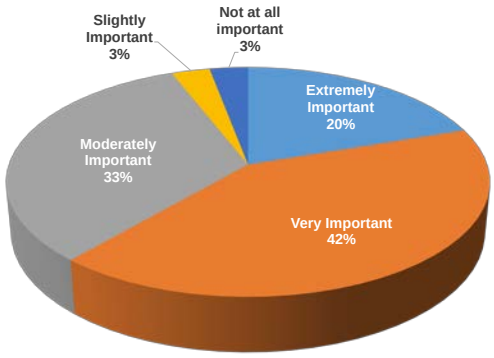

Almost two thirds of the IT organizations consider it to be either very or extremely important over the next year for them to get better performing management tasks such as troubleshooting on a per-VM basis.

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

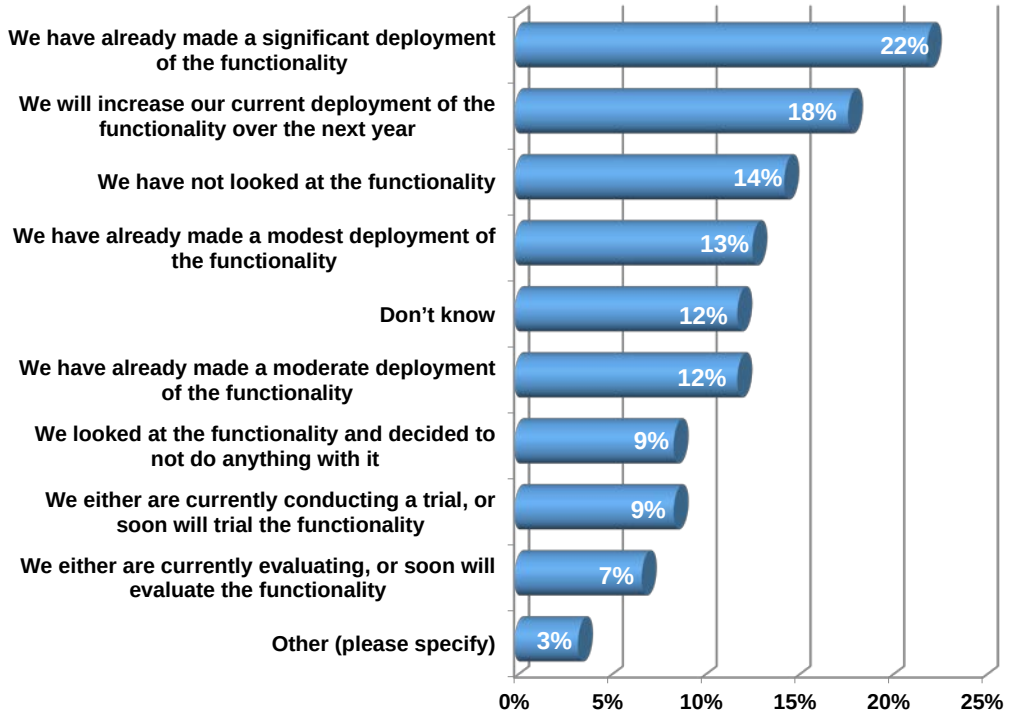

A WAN that features dynamic multi-pathing has all of the traditional performance, management and security challenges. An additional management challenge is being able to identify the end-to-end path that traffic took in order to be able to troubleshoot degraded network or application performance.

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[3] 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**.

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

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**.

Figure 4: Importance of Getting Better at Managing Private Cloud

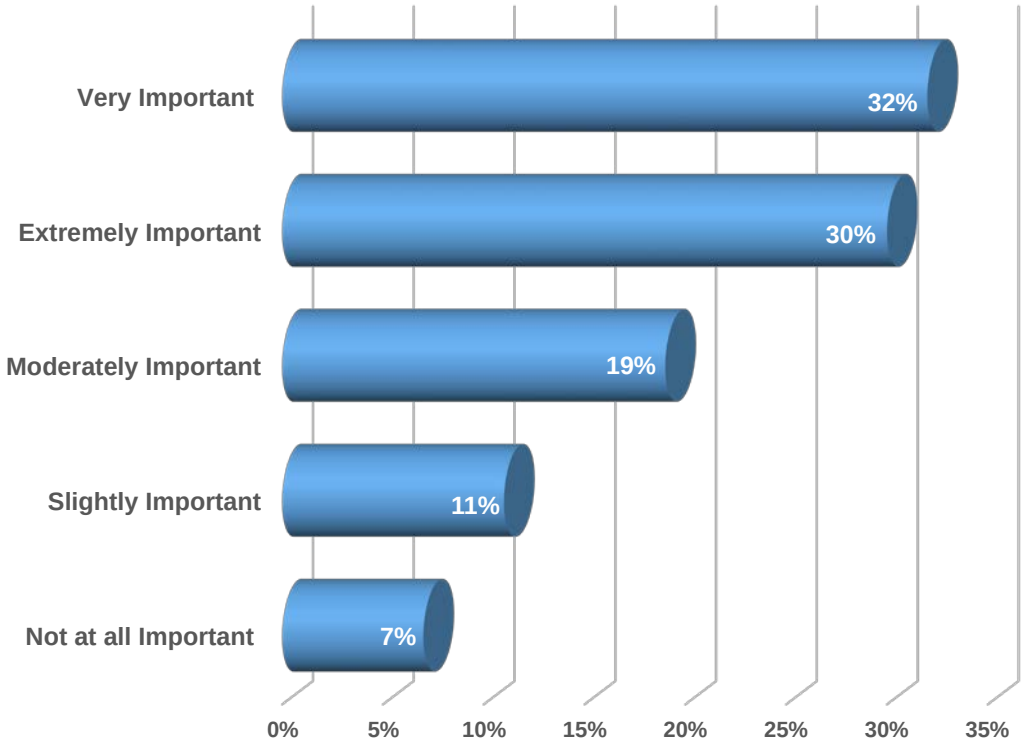

Managing end-to-end in a private cloud environment is slightly more important to IT organizations than is managing end-to-end in a public cloud environment.

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 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.

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

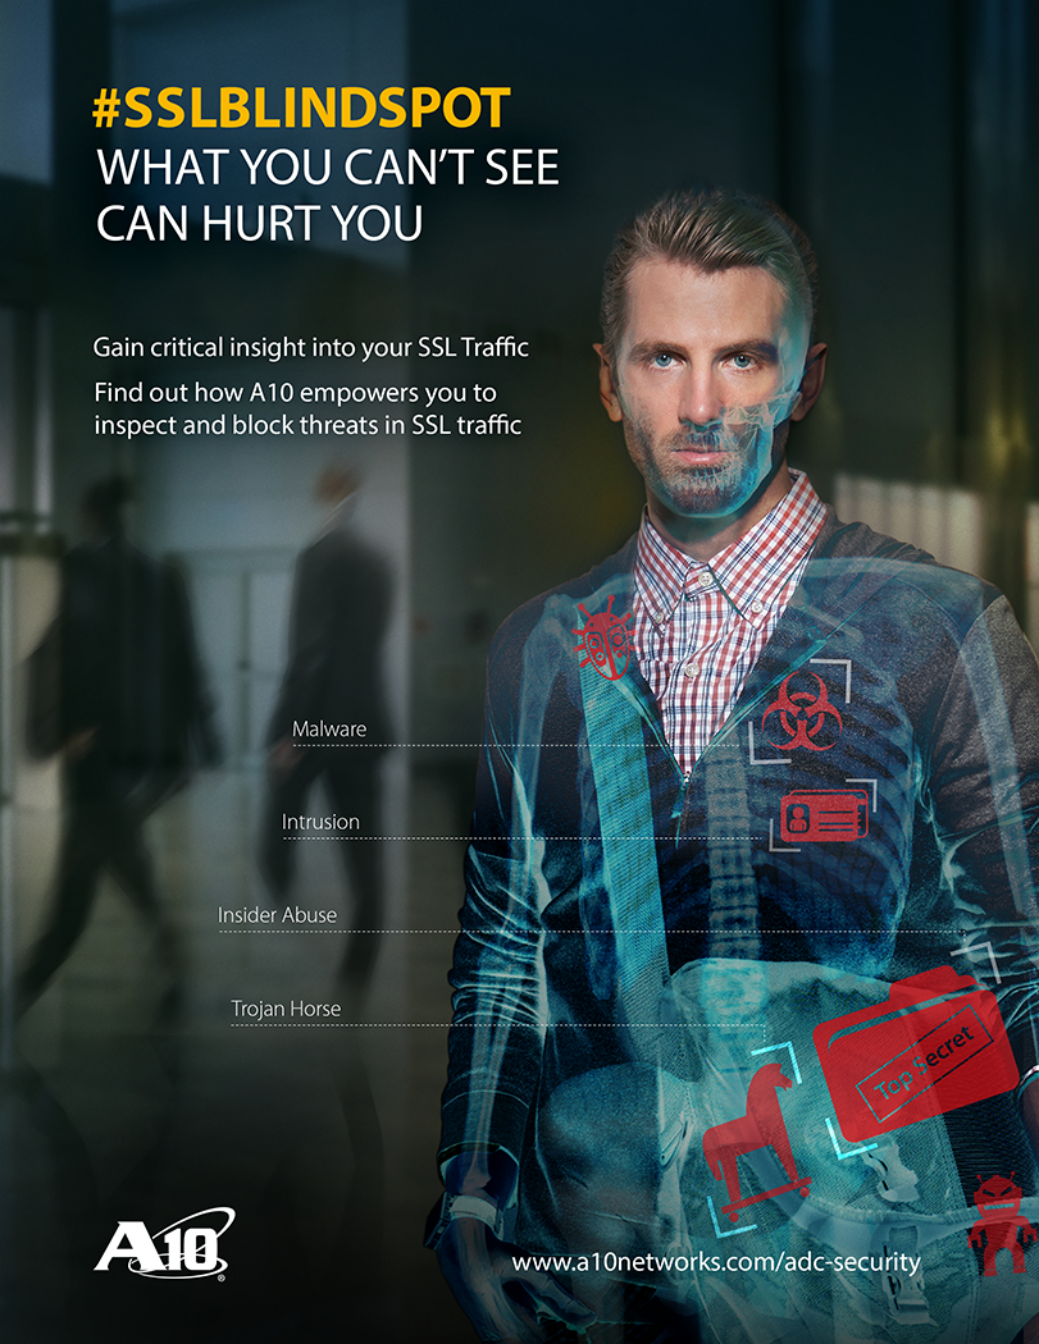
#SSLBLINDSPOT
WHAT YOU CAN'T SEE
CAN HURT YOU
Gain critical insight into your SSL Traffic
Find out how A10 empowers you to
inspect and block threats in SSL traffic
Malware
Intrusion
Insider Abuse
Trojan Horse
Top Secret
A10
www.a10networks.com/adc-security

SDN Today:

Delivered by Citrix NetScaler and Cisco ACI

Learn more at citrix.com/netscaler/cisco

CITRIX®

cisco™

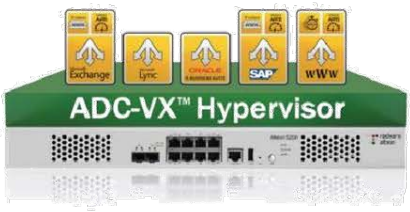

Predictable Application Service Levels, Guaranteed—Only with Alteon NG

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.