The 2014
Application & Service Delivery Handbook

Part 1: Introduction and Challenges

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

Platinum Sponsors:

Gold Sponsors:

Produced by:

Webtorials

Table of Contents

Executive Summary

The ***2014 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. One goal of this publication is to describe how the ***2014 Application and Service Delivery Handbook*** differs from previous editions in this series. Another goal of this publication is to describe how a variety of traditional and emerging factors are complicating the task of ensuring acceptable application and service delivery.

Subsequent publications of the ***2014 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.
- The management and security of applications and services.

Each of the first three publications will also contain recent market research that identifies how IT organizations are approaching application and service delivery. The fourth and final publication will include an executive summary as well as a copy of the complete document.

Background and Goals of the 2014 Application and Service Delivery Handbook

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

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

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

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

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

A few years ago, a second generation of challenges began to emerge. Those challenges were driven in large part by the:

- Implementation of varying forms of virtualization;
- Adoption of cloud computing;
- Emergence of a sophisticated mobile workforce;
- Shifting emphasis and growing sophistication of cyber-crime.

While all of the first generation of application delivery solutions still add value, a second generation of application delivery solutions is emerging. In many cases these solutions aren't hardware based, but are software based. In a growing number of instances they are provided as part of a managed service or acquired from a public cloud provider. The management component of this new generation of application delivery solutions is less likely to be focused narrowly just on the network and more likely to integrate network and application management.

To further add to the demanding environment, there are two emerging megatrends that have the potential to dramatically impact application delivery. Those megatrends are Software Defined Networking (SDN) and Network Function Virtualization (NFV). As is explained later in The Handbook, these two megatrends are a double edged sword in that they have the potential to make application delivery notably more difficult and, at the same time, they have the potential to ease some of the burden of ensuring acceptable application and service delivery.

The goal of the 2014 Application and Service Delivery Handbook is to help IT organizations ensure acceptable application and/or service delivery when faced with both the first generation, as well as the emerging set of application and service delivery challenges.

Foreword to the 2014 Edition

This year's edition of The Handbook builds on the 2013 Application and Service Delivery Handbook[1]. However, any material in the 2013 edition of The Handbook 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 2014 edition of The Handbook with a URL. Using this approach, this year's edition of The Handbook is significantly reduced in size and it focuses primarily on the changing nature of application delivery.

In early 2014, a survey was given to the subscribers of Webtorials. Throughout this document, the IT professionals who responded to the surveys will be referred to as ***The Survey Respondents***. The survey focused on identifying the optimization and management tasks that are of most interest to IT organizations. With that goal in mind, The Survey Respondents were given a set of optimization tasks and a set of management tasks and asked to indicate how important it was to their IT organization to get better at these tasks over the next year. The survey question used the following five-point scale:

1. Not at all important
2. Slightly important
3. Moderately important

[1] http://www.webtorials.com/main/resource/papers/webtorials/2013-App-Serv-Handbook/2013_Application_and_Service_Delivery_Handbook-Complete.pdf

4. Very Important
5. Extremely important

The answers to the survey will be used throughout the ***2014 Application and Service Delivery Handbook*** to demonstrate both the challenges facing IT organizations as well as the relative importance that IT organizations place on getting better with a wide variety of optimization and management tasks.

First Generation Application & Service Delivery Challenges

There are a number of fairly well understood challenges that have over the years complicated the task of ensuring acceptable application and service delivery. Those challenges are listed below and are described in detail in the document entitled Traditional Application & Service Delivery Challenges[2].

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

[2] http://www.ashtonmetzler.com/Traditional%20App%20Delivery%20Challenges%20V2.0.pdf

Second Generation Application and Service Delivery Challenges

There are a number of second generation challenges that complicate the task of ensuring acceptable application and service delivery. Some of these challenges result from the adoption of application architectures such as SOA. These application architectures tend to be more susceptible to performance problems due to WAN impairments than do traditional application architectures. In addition, the introduction of technologies such as AJAX creates significant security vulnerabilities[3]. Many of the second generation application and service delivery challenges, such as the ones mentioned in the preceding paragraph, are described in the 2012 Application and Service Delivery Handbook[4].

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

- Mobility and BYOD
- Virtualization
- Cloud Computing

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

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

Mobility and BYOD

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

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

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

In the majority of instances, this new generation of employee-provided mobile devices doesn't run the Windows O/S and the existing security and management services for PCs must be extended for mobile devices or, alternatively, additional products and/or services added to perform these functions. Similar to PCs, smartphone and tablet computers are subject to malware and network intrusion attacks. On PCs, there are mature, robust products for malware protection (e.g. anti-virus

[3] http://resources.infosecinstitute.com/ajax-security-issues/
[4] http://www.webtorials.com/content/2012/08/2012-application-service-delivery-handbook-2.html

software) and network intrusion protection (e.g., personal firewall), but these protections are just now emerging for smartphones and tablet computers. Similarly, inventorying and updating installed software on smartphone and tablet computers are emerging capabilities and a critical area for Mobile Device Management solutions.

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

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

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

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

In order to quantify the concern amongst IT organizations about ensuring acceptable application and service delivery to mobile workers, The Survey Respondents were asked two questions. They were asked how important it is for their IT organization over the next year to get better at improving the performance of applications used by mobile workers. They were also asked how important it is for their IT organization over the next year to get better at managing and monitoring the performance of applications used by mobile workers. Their responses are shown in **Table 1**.

Table 1: Importance of Getting Better Delivering Mobile Applications

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

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

Virtualization

Server Virtualization

The vast majority of organizations have made at least some deployment of server virtualization and the deployment of server virtualization will increase over the next several years. Many of the same management tasks that must be performed in the traditional server environment need to be both

extended into the virtualized environment and also integrated with the existing workflow and management processes. One example of the need to extend functionality from the physical server environment into the virtual server environment is that IT organizations must be able to automatically discover both the physical and the virtual environment and have an integrated view of both environments. This view of the virtual and physical server resources must stay current as VMs move from one host to another. The view must also be able to indicate the resources that are impacted in the case of fault or performance issues.

To quantify the impact that managing on a per-VM basis is having on IT organizations, The Survey Respondents were asked how important it is for their IT organization over the next year to get better at performing traditional management tasks such as troubleshooting and performance management on a per-VM basis. Their responses are shown in **Figure 1**.

Figure 1: Managing on a per-VM Basis

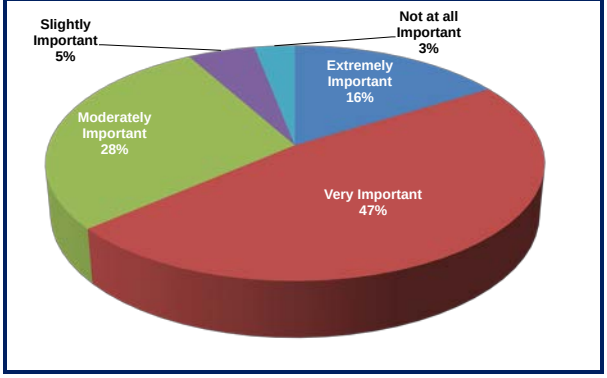

Almost two thirds of the IT organizations consider it to be either very or extremely important over the next year for them to get better performing management tasks such as troubleshooting on a per-VM basis.

Some of the other challenges created by server virtualization include:

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

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

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

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

Desktop Virtualization

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

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

To quantify the concern that IT organizations have relative to supporting desktop virtualization, The Survey Respondents were asked how important it is for their IT organization over the next year to get better at optimizing the performance of virtualized desktops. Their responses are shown in **Figure 2**.

Figure 2: Optimizing the Performance of Virtualized Desktops

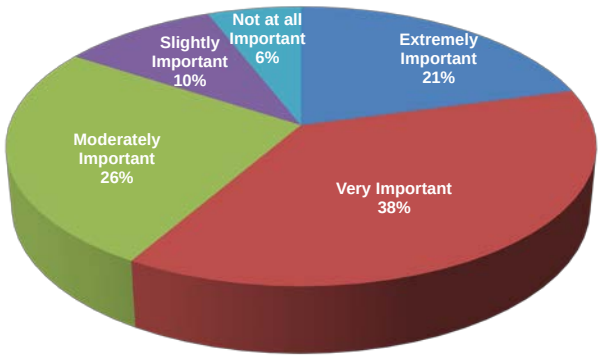

Well over half of The Survey Respondents indicated that getting better at optimizing the performance of virtualized desktops is either extremely or very important to their IT organization. That is a significant increase over the responses to the same question in 2013 and the responses in 2013 were a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

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

Software Defined Data Center (SDDC)

As noted, IT organizations are making increasing use of varying forms of virtualization. SDDC is an emerging concept that is being advocated by a number of vendors. The two primary characteristics of a SDDC are virtualization and automation. In particular, in a SDDC, the entire infrastructure is virtualized and delivered as a service and the control of this datacenter is entirely automated by software. Some vendors also advocate that the software run on commodity hardware. The document entitled The Promise and the Reality of a Software Defined Data Center[5] contains a detailed discussion of SDDCs. As described in that document, while it's true that few if any IT organizations have currently implemented an SDDC, it's also true that the steps that the majority of IT organizations have already taken to implement virtualization and automation are key steps on the path to implementing an SDDC.

[5] http://www.ashtonmetzler.com/ and click on *Journey to a new IT Operational Model*

Cloud Computing

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

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

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

In most instances the SLAs that are associated with public cloud computing services such as Salesforce.com or Amazon's Simple Storage System are weak and, as such, it is reasonable to say that these services are delivered on a best effort basis. For example, the SLA[6] that Amazon offers for its Amazon Web Services (AWS) states that, "AWS will use commercially reasonable efforts to make Amazon EC2 available with an Annual Uptime Percentage of at least 99.95% during the Service Year." As part of the Amazon definition of Annual Uptime Percentage, Amazon excludes any outage of 5 minutes or less. The Amazon SLA also states that if their service doesn't meet the Annual Uptime Percentage commitment, the customer will receive 10% off its bill for the most recent month that the customer included in the SLA claim that it filed.

A key attribute of the vast majority of the SLAs that are associated with public cloud computing services is that they don't contain a goal for the end-to-end performance[7] of the service. The reason for the lack of performance guarantees stems from the way that most public cloud computing services are delivered. As shown in **Figure 3**, one approach to providing public cloud computing services is based on the service being delivered to the customer directly from an independent software vendor's (ISV's) data center via the Internet. Another approach is for an ISV to leverage an IaaS provider such as Amazon to host their application on the Internet. Both of these approaches rely on the Internet and it is not possible to provide end-to-end quality of service (QoS) over the Internet. As a result, neither of these two approaches lends itself to providing an SLA that includes a meaningful commitment to critical network performance metrics such as delay, jitter and packet loss.

[6] http://aws.amazon.com/ec2-sla/
[7] In this context, *performance* refers to metrics such as delay or response time.

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

Figure 3: Distribution Models for Cloud-Based Solutions

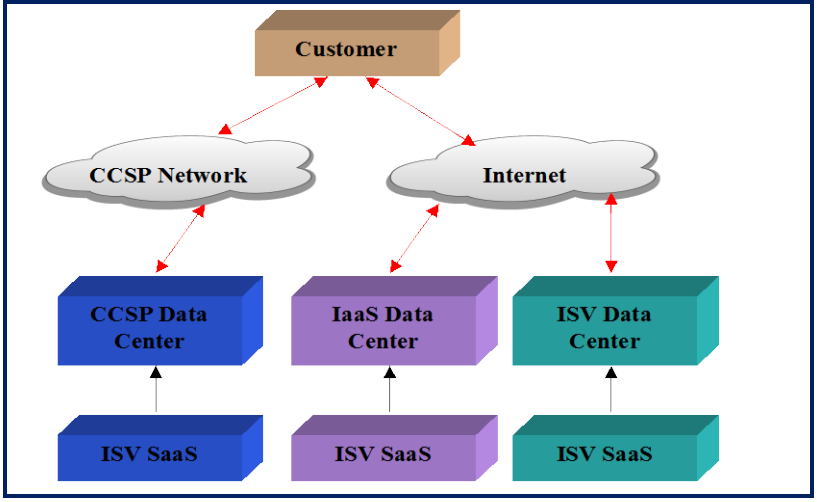

An approach to providing public cloud computing services that does lend itself to offering more meaningful SLAs is based on a CCSP providing these solutions to customers from the CCSP's data center and over a network that is provided by the CCSP and based on a technology such as MPLS.

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

The Survey Respondents were asked to indicate how important it was over the next year for their organization to get better a managing end-to-end in a private cloud environment. **Table 2** shows how The Survey Respondents answered this question in 2014 and also shows how a corresponding set of survey respondents answered this question in 2013.

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

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

The Survey Respondents were also asked to indicate how important it was over the next year for their organization to get better a managing end-to-end in a public cloud environment. **Table 3** shows how The Survey Respondents answered this question in 2014 and also shows how a corresponding set of survey respondents answered this question in 2013.

Table 3: Importance of Getting Better at Managing Public Cloud: 2014 vs. 2013		
	Managing Public Cloud - 2014	**Managing Public Cloud – 2013**
Extremely Important	21%	11%
Very Important	29%	26%
Moderately Important	25%	28%
Slightly Important	16%	14%
Not at all Important	9%	21%

Two key conclusions can be drawn from the data in **Table 2** and **Table 3**:

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

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

Emerging Challenges

Software Defined Networking

A detailed discussion of Software Defined Networking (SDN) can be found in *The 2013 Guide to Network Virtualization and SDN.*[8]

The ONF is the organization that is most associated with the development and standardization of SDN. According to the ONF[9], "Software-Defined Networking (SDN) is an emerging architecture that is dynamic, manageable, cost-effective, and adaptable, making it ideal for the high-bandwidth, dynamic nature of today's applications. This architecture decouples the network control and forwarding functions enabling the network control to become directly programmable and the underlying infrastructure to be abstracted for applications and network services. The OpenFlow™ protocol is a foundational element for building SDN solutions."

Figure 4 contains a graphical representation of the SDN solution architecture as envisioned by the ONF.

Figure 4: ONF's SDN Architecture

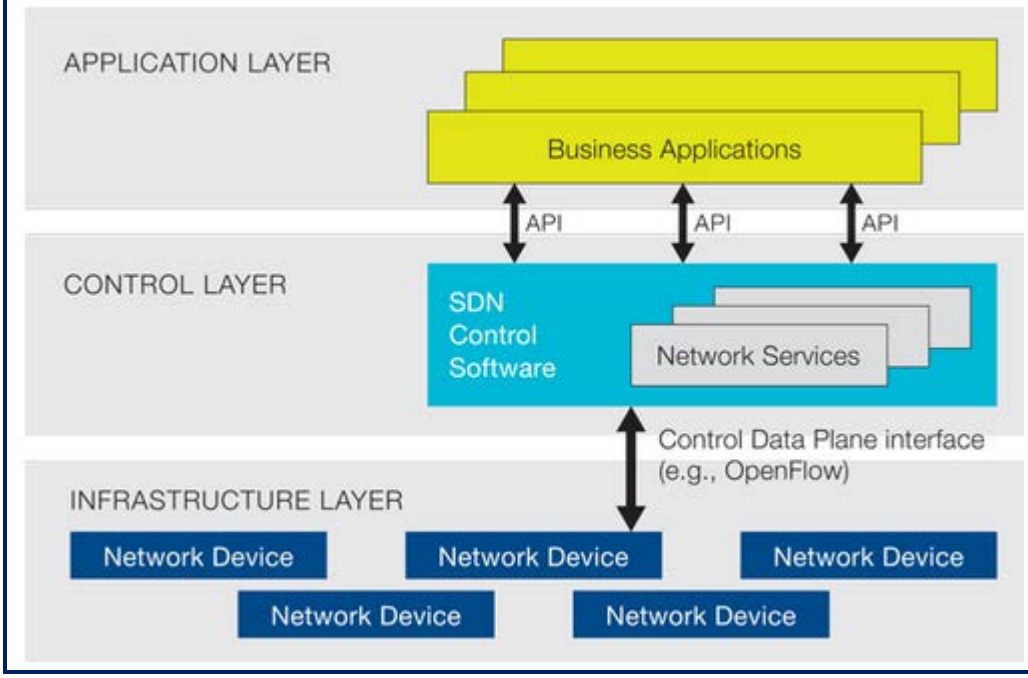

While the use of SDN in data centers receives the majority of attention in the trade press, it is also possible to implement SDN in branch and campus networks as well as in wide area networks (WANs).

[8] http://www.webtorials.com/content/2013/10/2013-guide-to-software-defined-networking-network-virtualization.html
[9] https://www.opennetworking.org/sdn-resources/sdn-definition

The OpenFlow Protocol

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

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

The most current version of OpenFlow that vendors have begun to support is OpenFlow 1.3. A complete, detailed description of the functionality provided by OpenFlow 1.3 can be found at the ONF Web site[10].

Network Function Virtualization

A detailed description of Network Function Virtualization (NFV) can be found at *An NFV Reality Check*[11].

Figure 5: The Virtualization of Network Appliances
Source: NFV ISG

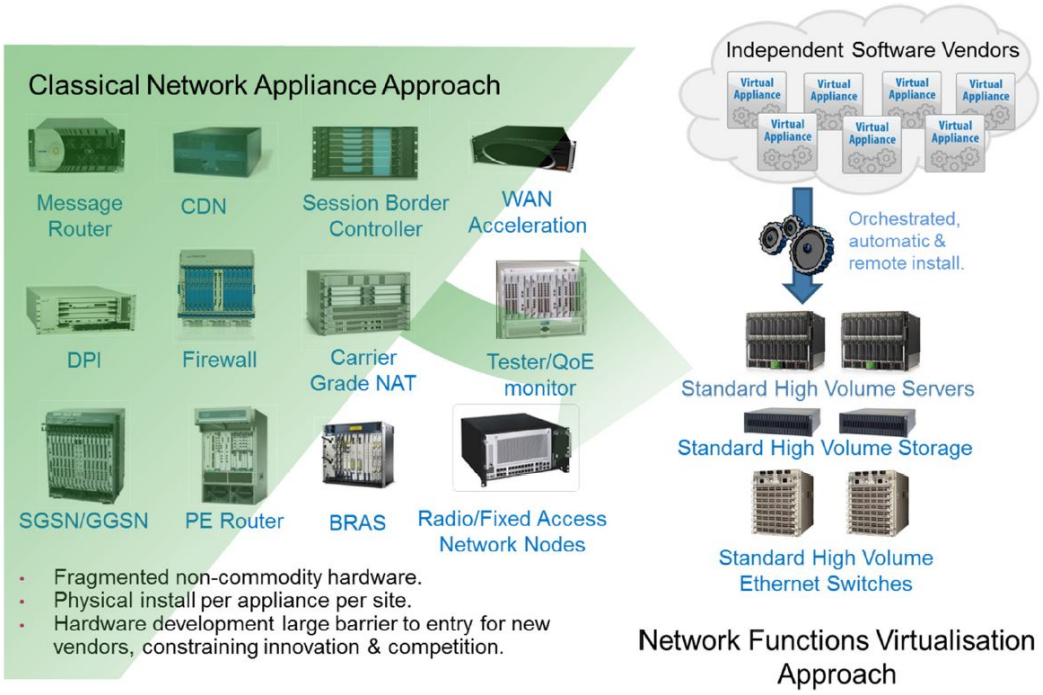

[10] https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow/openflow-spec-v1.3.0.pdf
[11] http://www.ashtonmetzler.com/NFV%20Reality%20Check.pdf

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

Telecommunications service providers feel that they can greatly simplify their operations and reduce expense if all network functions were available as virtual appliances that can be easily provisioned and integrated regardless of the vendor who provided the appliance or the hypervisor(s) on which it runs. In order to bring this vision to fruition, an Industry Specifications Group for Network Functions Virtualization (NFV ISG) was formed under the auspices of the European Telecommunications Standards Institute (ETSI). Their vision for the transition from hardware appliances of today to a fully virtualized appliance environment is depicted in **Figure 5**.

The approach that the NFV ISG is taking is that the virtualization of network functionality is applicable to any data plane packet processing and control plane function in both fixed and mobile networks. In October 2013, ETSI published a set of high level reference documents that are openly available on the ETSI website[12]. One of those documents discussed a framework for conducting a NFV Proof of Concept (POC). ETSI currently has eighteen POCs underway.

Until recently, the conventional wisdom in the IT industry in general, and on the part of the ONF and ETSI in particular, was that was that SDN and NFV were separate topics and didn't need to be formally coordinated. That conventional wisdom changed in March 2014 when the ONF and ETSI announced the signing of a Memorandum of Understanding (MOU). As part of the announcing the MOU[13], the ONF and ETSI said that "Together the organizations will explore the application of SDN configuration and control protocols as the base for the network infrastructure supporting NFV, and conversely the possibilities that NFV opens for virtualizing the forwarding plane functions."

Some of the challenges associated with NFV include:

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

[12] http://www.etsi.org/nfv
[13] http://www.rethink-wireless.com/2014/03/19/etsi-nfv-group-closer-operator-sdn.htm

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 at *jim@webtorials.com* or Steven Taylor at *taylor@webtorials.com*.

Published by Webtorials Editorial/Analyst Division
www.Webtorials.com

Division Cofounders:
Jim Metzler
jim@webtorials.com
Steven Taylor
taylor@webtorials.com

Professional Opinions Disclaimer

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

Copyright © 2014 Webtorials
For editorial and sponsorship information, contact Jim Metzler or Steven Taylor. The Webtorials Editorial/Analyst Division is an analyst and consulting joint venture of Steven Taylor and Jim Metzler.

6WINDGate™

Packet Processing Software

Increase Data Plane Performance
No Change To Linux Environments
Available Across All Major Platforms
Support Extensive Set Of Protocols

L2-L4 Acceleration
IPsec VPN Gateways
TCP / UDP Termination
Virtual Switching
DPDK
And More ...

6WIND.com

SPEED MATTERS

ARE YOUR USERS WAITING MORE THAN DOING?

A10 helps more than 3,000 Enterprises and Service Providers deliver quality experiences

MAKE YOUR APPS RUN FASTER

Experience the Performance

www.a10networks.com

Enterprises today aspire to grow revenue by expanding globally and acquiring new customers, while also cutting costs and finding ways to become more agile. To realize their goals, every Enterprise has a core set of applications that they rely on to run their business operations.

Current Application Delivery Landscape

The user requirements for accessing and business applications is changing dramatically, and Enterprises must support more applications across a broader user base including customers, suppliers, partners, and employees. In order to leverage their applications to achieve their business goals, Enterprises must optimize the delivery of their applications to support fast, reliable, and secure access to ensure all users, both inside and outside of their organization, have the best possible experience.

In the past, Enterprises would resort to optimizing their application delivery using a physical hardware box or a virtual appliance that was deployed within a data center and any offices where users were located. While costly to deploy and manage, this approach did a good job of optimizing application delivery between the data center and branch office locations that were connected via a private network. Today, this approach is no longer effective due to several factors including:

- The complexity of having more users outside the organization's private network
- Applications distributed across multiple data centers and in the cloud
- End-users located all over the world using all sorts of different devices and networks, and
- A growing list of critical business applications such as CRM, collaboration, product lifecycle management, and support portals that users rely on every day.

It's not realistic for IT organizations to establish private network connections between all their users and all the data centers where their applications are hosted, or implement an application delivery box or virtual appliance in every data center, cloud environment, and every location where their end-users are located today.

In order to leverage their applications to achieve their business goals, organizations today cannot only rely exclusively on their private WAN to deliver their applications, but they must also leverage the ubiquity and scale of the Internet in order to embrace the trends of globalization and consumerization within their organizations.

Considering Akamai's Cloud-based Application Delivery Platform

Akamai's Terra Alta solution is a cloud-based Application Delivery Platform that enables Enterprises to leverage the Internet to deliver all their web-based applications in a fast, reliable, secure, and cost-effective way. Terra Alta is a managed service that empowers Enterprises to overcome the challenges related to delivering their applications over the Internet by placing all of the application delivery capabilities within Akamai's cloud-based Intelligent Platform, instead of requiring IT organizations to take on the burden of deploying and managing these critical capabilities on their own in the form of hardware boxes or virtual appliances. With Akamai, application optimizations are distributed globally across our Intelligent Platform, not constrained within the four walls of a few data centers, or restricted only to those users on a private network connection.

Akamai's Intelligent Platform is deployed on over 150,000+ servers which are embedded deeply into thousands of networks worldwide, which means we are very close to nearly all of the world's Internet users and datacenters. This means that users can benefit from fast, reliable, and secure business applications regardless of where they are located in the world! In addition to being a cloud-based platform, Akamai is device agnostic and does not require any application changes, which means it's quick and easy to implement and allows organizations to lower their IT costs and reduce complexity as compared to alternative application delivery optimization solutions. Akamai's unique cloud-based architecture also means that applications can be seamlessly migrated across data centers or cloud providers at will, and the application delivery optimizations will automatically move with the application. Terra Alta empowers Enterprises to embrace their cloud, mobile, and big data initiatives without the fear of increased costs or low application adoption.

Conclusion

By overcoming the new realm of global application delivery challenges, Akamai's cloud-based Application Delivery Platform empowers organizations to meet the demands of globalization and consumerization and instantly enter new markets, acquire new customers, improve customer interactions, do business via lower-cost online channels, enable end-users to get more done in less time, and achieve their goal of increasing revenue and reducing costs.

We make the Internet fast, reliable, and secure.

Akamai

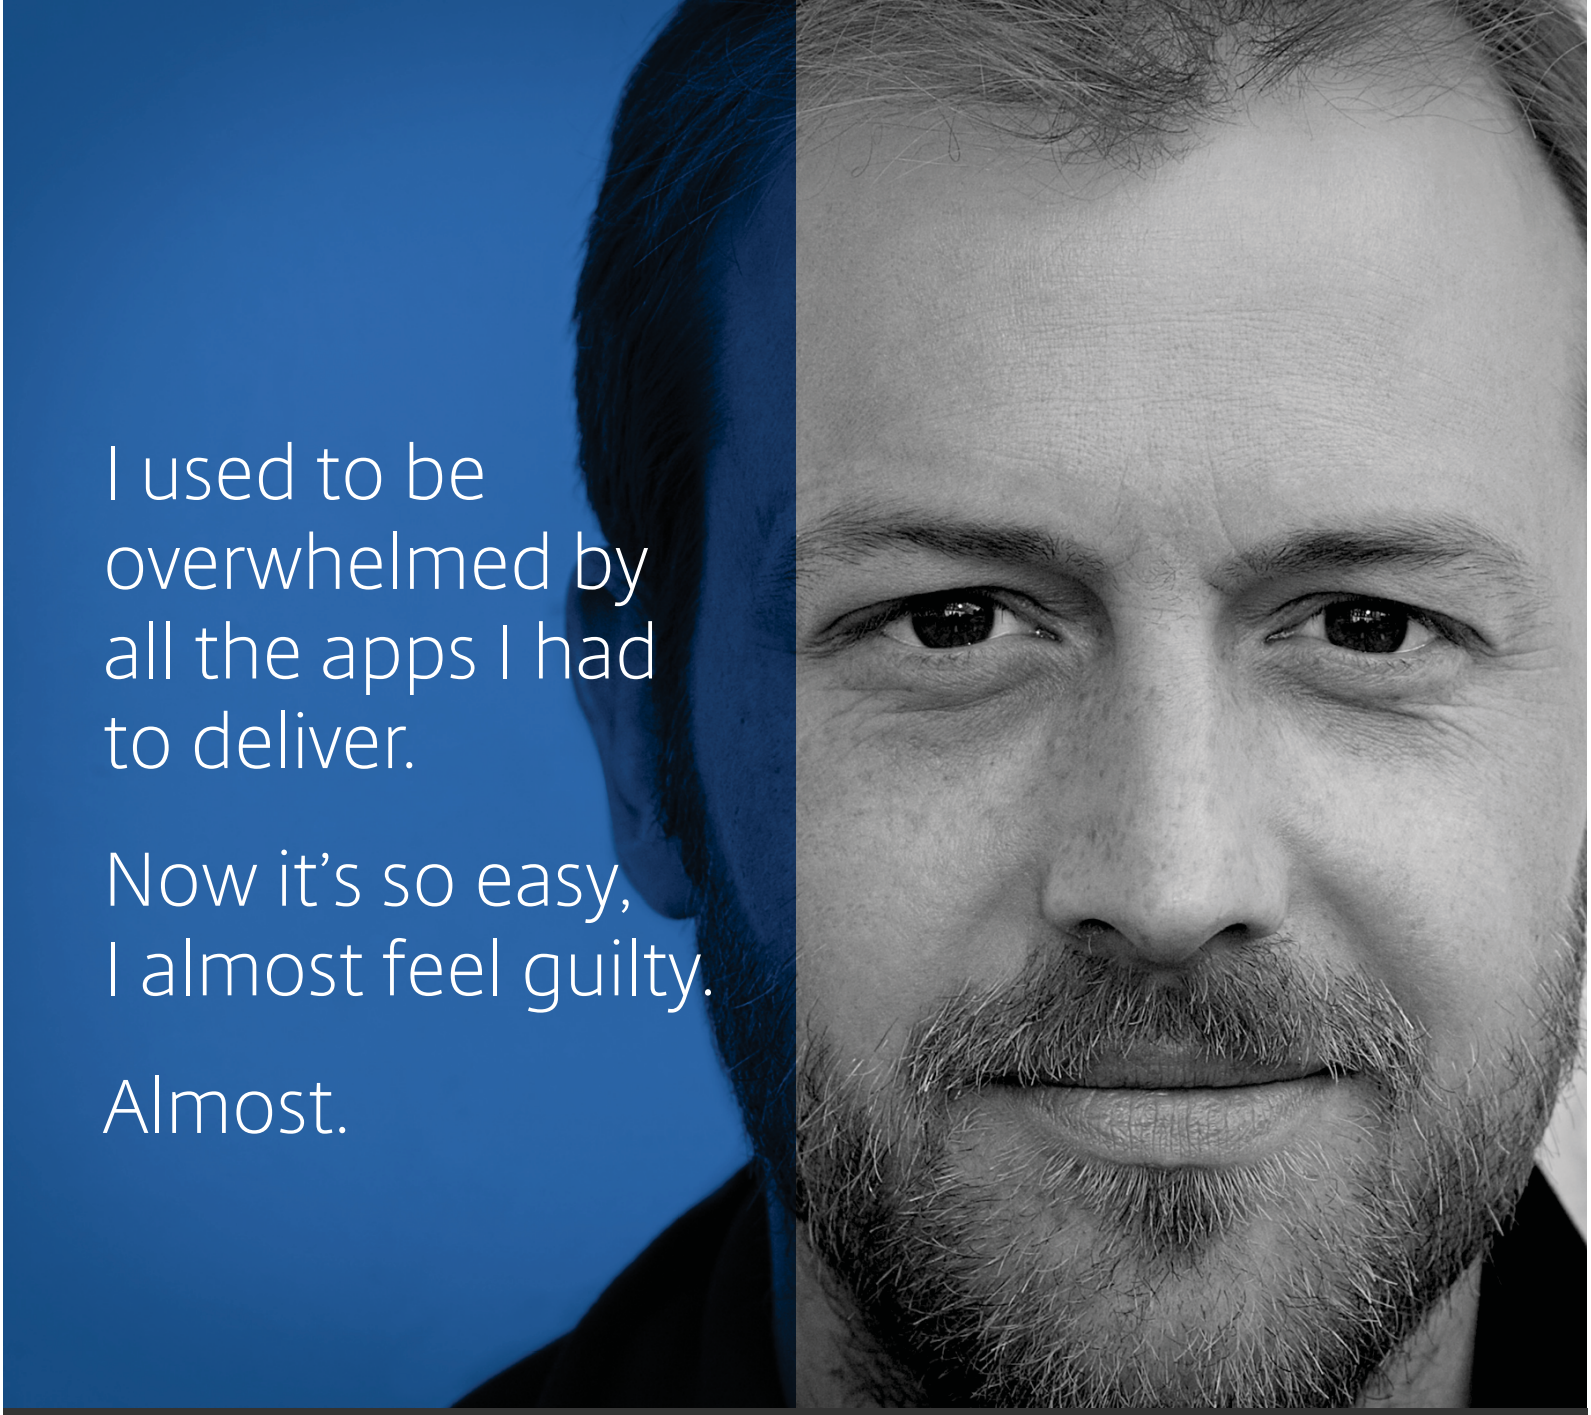

NetScaler with TriScale harnesses the power of software so you can effortlessly customize your app delivery for any business need.

NetScaler with TriScale
SOFTWARE SMART. HARDWARE STRONG.

CITRIX®

www.citrix.com/netscaler

CITRIX

Customer Spotlight

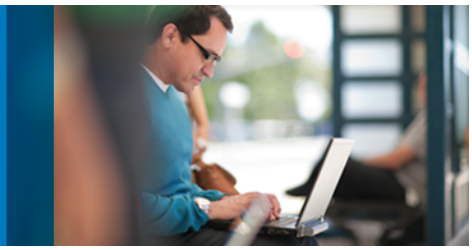

City Index

Industry:

Financial Services

Key Benefits:

- Reduces capital and operating expenses
- Provides an extended feature set on demand
- Ensures uninterrupted availability of trading platforms and applications

Citrix Products:

- Citrix NetScaler SDX
- Citrix NetScaler VPX

NetScaler reduces expenses with data center consolidation

City Index (www.cityindex.co.uk), is a leading global provider of retail trading services, including Spread Betting (UK only), Contracts for Difference (CFDs) and margined foreign exchange (FX). With offices in London (HQ), Warsaw, Tel Aviv, Singapore, Sydney and Shanghai, City Index supports its global clients on various trading platforms running on multiple backend applications. City Index is committed to providing a market-leading client service, transparent prices and innovative technology. All IT operations are centralized from two datacenters in London and one in Florida.

The challenge: need for a cost effective application delivery controller

Marc Morgan-Davies, Infrastructure Manager for City Index explained, our existing load balancing F5 solution that consisted of six BIG-IP 3400s and eight BIG-IP 6400s reached the end of support. There were two legacy Citrix Access Gateways (CAG) that we wanted to replace as well. We wanted a more cost effective and consolidated solution for our two London datacenters with an extended feature set that included global load balancing, DNS responder and rewrite, SSL offload, compression and caching.

Our choices were F5, Citrix, A10 Networks and Riverbed and we narrowed it down to F5 VIPRION 4400 and Citrix NetScaler 11500 SDX based on features and reputation. Our emphasis for the solution was more on the available feature set rather than raw processing power due to the nature of our platforms. The licensing model employed by Citrix is much simpler and more cost effective in my opinion than the competitors. For instance, if you want to enable a fourth module on the F5 you require another physical blade." Marc Morgan-Davies continued, "NetScaler on the other hand provided all of our required features on a single appliance with a simpler licensing model as well as allowing us to consolidate the existing CAGs onto the new devices further reducing our physical footprint and operating expenses. NetScaler gave us more features at a lower cost so was our chosen solution."

The solution: NetScaler SDX

Citrix NetScaler is an Application Delivery Controller (ADC) that optimizes the security, availability, scalability and performance of web-based applications and is available as a physical or virtual appliance. Citrix NetScaler

SDX is a true service delivery networking platform for enterprises and cloud datacenters. NetScaler SDX provides an advanced virtualized architecture that supports multiple NetScaler instances on a single hardware appliance, while an advanced control plane unifies provisioning, monitoring and management to meet the most demanding multi-tenant requirements**.**

NetScaler VPX is a software-based virtual appliance built for cloud scale. As an easy-to-deploy application delivery solution that runs on multiple virtualization platforms, the simplicity and flexibility of NetScaler VPX make it simple and cost-effective to fully optimize every web application and more effectively integrate networking services with application delivery. Performance capacities can be upgraded in production with the simple addition of a pay-as-you-grow license. NetScaler VPX helps organizations control costs by leveraging processing capacity already in place, including existing virtualized servers and associated resources.

"In October 2012, we installed 2 NetScaler SDXs as HA pair in production in each of our London datacenters. Each SDX box have 2 VPX instances that have discrete security layers In addition, we installed 2 SDXs as HA pair for staging in London with each SDX running 9 VPX instances. We are extremely pleased with NetScaler's ease of configuration and use." said Morgan Davies.

Key benefit: reduces capital and operating expenses

Using NetScaler we were able to prevent appliance sprawl by upgrading and consolidating 14 F5 Big-IP appliances and 2 Citrix Access Gateways to just 6 Citrix NetScaler appliances. This helped reduce support costs, rack space, ongoing power and cooling requirements drastically." Marc Morgan-Davies emphasized.

Key benefit: provides an extended feature set on demand

According to Marc Morgan-Davies, "Citrix NetScaler helped upgrade infrastructure while controlling costs. NetScaler provided the complete ADC feature set we required with the ability to enable features on demand."

Key benefit: ensures uninterrupted availability of trading platforms and applications

NetScaler ensured 24X7 availability of City Index's trading platforms and applications by providing global load balancing and SSL offloading between the London datacenters.

Looking ahead to the future

"NetScaler is a requirement to deploy Citrix XenMobile and MDM that would satisfy our requirement for a secure solution for users to access the corporate resources from any location using any device. Citrix mobility technologies are now very much on our scope for implementation in the near future. We are also looking into NetScaler App Firewall feature as well." Marc Morgan-Davies concluded.

About Citrix

Citrix (NASDAQ:CTXS) is the cloud company that enables mobile workstyles—empowering people to work and collaborate from anywhere, easily and securely. With market-leading solutions for mobility, desktop virtualization, cloud networking, cloud platforms, collaboration and data sharing, Citrix helps organizations achieve the speed and agility necessary to succeed in a mobile and dynamic world. Citrix products are in use at more than 330,000 organizations and by over 100 million users globally. Annual revenue in 2012 was $2.59 billion. Learn more at www.citrix.com.

©2014 Citrix Systems, Inc. All rights reserved. Citrix®, NetScaler®, SDX™, VPX™ XenMobile and App Firewall are trademarks or registered trademarks of Citrix Systems, Inc. and/or one or more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries. All other trademarks and registered trademarks are property of their respective owners.

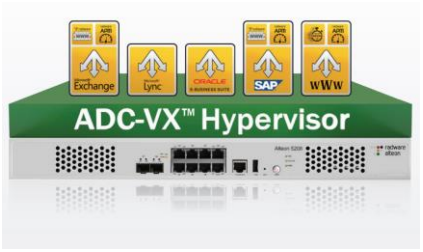

Predictable Application SLA, Guaranteed. Only with Alteon NG.

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

The Standard ADC: Not Good Enough Anymore

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

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

Alteon NG: Complete Application SLA Assurance

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

- ☑ Alteon NG is **architecturally designed to ensure application SLA** by delivering full resource isolation per application, service, or department. Each virtual ADC (vADC) instance is completely isolated from neighboring instances with independent CPU cores, memory, network stack, management control, and operating system. Our unique solution is designed to dynamically scale to add more throughput, services, and vADCs without hardware modification resulting in fast provisioning of additional vADC instances and no service degradation, interruption, or resource overcapacity.
- ☑ Radware's **Application Performance Monitoring (APM)** module provides real-time tracking of application SLAs by measuring real-user transactions and errors. Embedded in Alteon NG, Radware's APM is an out-of-the-box solution which doesn't require synthetic transaction scripting or additional installation - reducing deployment time and costs. Radware's APM intuitively tracks SLA by location, user, application and transaction type to expedite root cause analysis. In addition, it provides historical reports based on user-defined SLA that feature granular analysis allowing the measurement of the delay per transaction phase including data center time, network latency and browser rendering time.
- ☑ Alteon NG integrates FastView® the industry's most advanced **Web Performance Optimization (WPO)** technology – which accelerates application response by up to 40% – for higher conversion rates, revenues, productivity, and customer loyalty. FastView acceleration treatments are optimized according to each user, end-user device and browser - with specific optimization for mobile devices. In addition, FastView automatically optimizes new applications,

new application versions and new application modules – reducing manual code optimization while letting you focus on core business competencies.

- ☑ Alteon NG is part of Radware's unique **Attack Mitigation System (AMS)**, which enables accurate detection and mitigation of the most advanced cyber-attacks. Leveraging a unique Defense Messaging™ mechanism, AMN efficiently mitigates attacks by signaling attack information to Radware DefensePipe cloud service and Radware DefensePro data center attack mitigator, located in the cloud or the network perimeter, respectively.
- ☑ Integrating advanced **Web Application Firewall (WAF)** capabilities, Alteon NG enables risk-free implementation thanks to a unique out-of-path WAF deployment mode along with auto-policy generation capabilities. Moreover, as ADC resources are ensured via full instance isolation and resource reservation, even when WAF policies are updated there's no impact on application availability and performance. This results in secured web applications with SLA guarantee.
- ☑ Alteon NG features a built-in authentication gateway with **Single Sign On (SSO)** capabilities by supporting Radius, Active Directory, LDAP and RSA SecurID – simplifying the user experience without compromising on application security.
- ☑ Alteon NG employs Radware's **AppShape™** offering configuration templates for leading business applications (e.g. Microsoft, Oracle, SAP). This helps customers roll out ADC-optimized applications in a simple, fast risk-free manner. In addition, Radware's AppShape++ scripting technology lets customers customize any ADC service per specific application flow/scenario. Using the AppShape++ script library, customers can refine various Layer 4-7 policies including HTTP, HTTPS, TCP, UDP, SSL and more – without application modifications to reduce cost and risk.

Complete Load Balancing/Layer 4-7 Feature Set

Alteon NG delivers a complete set of layer 4-7 services to ensure the availability, performance and security of mission-critical applications in the local and cloud data centers. These extend to traffic redirection, content modification, persistency, redundancy, advanced health monitoring and global server load balancing (GSLB). In addition, Alteon NG integrates advanced modules such as bandwidth management and link load balancing – reducing data center footprint and simplifying deployment. The combination of these advantages – along with an industry unique 5-year longevity guarantee, "pay-as-you-grow" approach in throughput, number of vADCs and services, plus performance leadership in all layer 4-7 metrics – makes Alteon simply your best application delivery choice.

Want to see more for yourself? We invite you to download our Radware ADC solution white paper here or contact us at: info@radware.com.