The 2014
Application & Service Delivery Handbook

Part 3: Management & Security

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

Platinum Sponsors:

A10

Akamai

CITRIX

Quali systems

Gold Sponsors:

6WIND

ca technologies

radware

riverbed

Produced by:

Webtorials

Table of Contents

Executive Summary

The ***2014 Application and Service Delivery Handbook (The Handbook)*** will be published both in its entirety and in a serial fashion. The first and second chapters have already been published. The first chapter describes how the ***2014 Application and Service Delivery Handbook*** differs from previous editions in this series. That chapter also describes how a variety of traditional and emerging factors are complicating the task of ensuring acceptable application and service delivery. The second chapter of ***The Handbook*** focuses on describing the technologies, products and services that are available to improve the performance of applications and services.

This is the third of the serial publications and contains the third chapter of ***The Handbook***. The third chapter of ***The Handbook*** will focus on describing the technologies, products and services that are available to improve the management and security of applications and services.

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

Management

Market Research

The previous chapters of ***The Handbook*** discussed a survey that was given in early 2014 to the subscribers of Webtorials. As previously noted, within ***The Handbook*** the respondents to that survey will be referred to as The Survey Respondents. **Table 1** shows how The Survey Respondents answered the 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%	4.8%	11.5%	39.4%	44.2%
Identify the components of the IT infrastructure that support the company's critical business applications	1.2%	4.3%	10.4%	44.8%	39.3%
Obtain performance indicator metrics and granular data that can be used to detect and eliminate impending problems	1.2%	7.8%	16.8%	47.3%	26.9%
Monitor the end user's experience and behavior	0.0%	6.6%	21.0%	42.5%	29.9%

Table 1: The Importance of Getting Better at Key Management Tasks					
Effectively manage SLAs for one or more business critical applications	1.2%	5.5%	18.8%	41.2%	33.3%
Manage the use of VoIP	5.6%	7.4%	24.7%	30.9%	31.5%
Perform traditional management tasks such as troubleshooting and performance management, on a per VM basis	3.0%	5.5%	27.9%	47.3%	16.4%
Monitor and manage the performance of applications delivered to mobile users	4.9%	14.7%	25.8%	32.5%	22.1%
Manage end-to-end in a public cloud computing environment	9.0%	16.0%	25.0%	29.5%	20.5%
Manage end-to-end in a private cloud computing environment	6.3%	6.3%	27.5%	38.8%	21.3%
Effectively monitor and manage an application acquired from a SaaS provider such as Salesforce	13.7%	7.2%	29.4%	33.3%	16.3%
Manage end-to-end in a public cloud computing environment	9.0%	16.0%	25.0%	29.5%	20.5%
Effectively monitor and manage computing services acquired from a IaaS provider such as Rackspace	12.8%	16.2%	32.4%	23.6%	14.9%

Some of the conclusions that can be drawn from the data in **Table 1** include:

Two tasks are in a virtual tie for the most important management task to get better at over the next year: 1) Rapidly identifying the root cause of degraded application performance; 2) Identifying the components of the IT infrastructure that support the company's critical business applications.

The second most important set of management tasks include: 1) Obtain performance indicator metrics and granular data that can be used to detect and eliminate impending problems; 2) Monitor the end user's experience and behavior; 3) Effectively manage SLAs for one or more business critical applications; 4) Manage the use of VoIP.

While managing the use of services acquired from an IaaS provider such as Rackspace is relatively unimportant, it is more important than it was last year.

Forces Driving Change

Previous sections of this handbook described the traditional and emerging service and application delivery challenges. This subsection will identify how some of those challenges are forcing a change in terms of how IT organizations manage applications and services.

Server Virtualization

Until recently, IT management was based on the assumption that the IT organizations 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.

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 Balancing

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 IaaS solutions in general, and the adoption of cloud balancing in particular, demonstrates why 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. The adoption of cloud balancing is also another example of why IT organizations need to adopt an approach to management that is based on the assumption that the components of a service, and the location of those components, can and will change frequently.

Delay Sensitive Traffic

Voice and video are examples of applications that have high visibility and which are very sensitive to transmission impairments. As was also highlighted in **Table 1**, 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. This includes looking at the application payload and measuring the quality of the voice and video communications. In the case of Unified Communications (UC), it also means monitoring the signaling between the components of the UC solution.

In addition to having a single set of tools and more of a focus on application payload, IT organizations need to implement management processes that understand the impact that each application is having on the other applications and that can:

- Analyze voice, video, UC and data applications in consort with the network;
- Support multi-vendor environments;
- Support multiple locations.

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, the 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. In a similar fashion, operations management needs to be consolidated and automated to keep service quality in line with user expectations.

Impact of SDN

SDN management is a combination of good news and bad news. The good news is that SDN has the potential to make network management easier. For example, in theory at least, SDN enables IT organizations to centralize configuration and policy management.

The bad news is that SDN creates some new management challenges. For example, one of the primary benefits of SDN is the ability to support multiple virtual networks that run on top of the physical network. Effective operations management, however, requires tools that give operators clear visibility into the relationships between the virtual and physical networks and their component devices. In particular, when performance or availability problems occur, both root cause analysis and impact analysis require bilateral mapping between the physical and virtual infrastructures. With SDN, the flows between a pair of VMs can be distributed among a number of alternate paths through the network. Mapping a flow to the physical path it takes can be a challenge unless the flow monitoring solution can involve the controller's end-to-end view of the network

With SDN solutions, 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 controller performance in processing flows. One of the characteristics of SDN is that network functions such as load balancing and firewalls are increasingly implemented in software as network services that can be integrated with virtual networks or SDN flow under programmatic control; a.k.a., service chaining. Implementing these functions in software both increases the delay associated with performing these functions and it also increases the variability of that delay. The result is an increased need for insight into the performance of each component of the overall SDN solution.

The document entitled *Mock RFI for Enterprise SDN Solutions* contains a number of questions that IT organizations should ask SDN vendors relative to SDN management.

Application Performance Management

Background

Since any component of a complex service such as Customer Relationship Management (CRM) can cause service degradation or a service outage, in order to effectively perform application performance management IT organizations need a single unified view of all of the components that support a service. This includes the highly visible service components such as servers, storage, switches and routers, in both their traditional stand-alone format as well as in their emerging converged format. It also includes the somewhat less visible network services such as DNS and DHCP, which are significant contributors to application degradation. Multiple organizational units within an IT organization have traditionally provided all of these service components. On an increasing basis, however, one or more network service providers and one or more cloud computing service providers will provide some or all of these service components. As a result, in order to achieve effective service delivery management, management data must be gathered from the enterprise, one or more Network Service Providers (NSPs) and one or more Cloud Computing Service Providers (CCSPs). In addition, in order to help relate the IT function with the business functions, IT organizations need to be able to understand the key performance indicators (KPIs) for critical business processes such as CRM and relate these business level KPIs to the performance of the IT services that support the business processes.

As shown in **Table 1**, being able to monitor the end user's experience and behavior is a very important management task. One of the reasons for that importance is that in spite of all of the effort and resources that have gone into implementing IT management to date:

It is the end user, and not the IT organization who typically is the first to notice when the performance of an application begins to degrade.

An effective approach to application performance management must include the automatic discovery of all the elements in the IT infrastructure that support each service. This functionality provides the basis for an IT organization to being able to create two-way mappings between the services and the supporting infrastructure components. These mappings, combined with event correlation and visualization, can facilitate root cause analysis, significantly reducing mean-time-to-repair. **Table 1** demonstrated how important this functionality is to IT organizations.

If IT organizations can effectively identify which components of the infrastructure support a particular application or service, monitoring can much more easily identify when services are about to begin to degrade due to problems in the infrastructure. As part of this monitoring, predictive technique such as heuristic-based trending of software issues and infrastructure key performance indicators can be employed to identify and alert management of problems before they impact end users. In addition, outages and other incidents that generate alerts can be prioritized based on their potential business impact.

As **Table 1** also demonstrates, getting better at rapidly identifying the causes of application degradation is the most important management task facing IT organizations. Once the components of the infrastructure that support a given application or service have been identified, triage and root cause analysis can be applied at both the application and the infrastructure levels. When applied directly to applications, triage and root cause analysis can identify application issues such as the depletion of threads and pooled resources, memory leaks or internal failures within a Java server or .NET server. At the infrastructure level, root cause analysis can determine the subsystem within the component that is causing the problem.

Application Performance Management in the Private Enterprise Network[1]

Enterprise IT organizations can choose among several types of tools for monitoring and managing application performance over a private enterprise network. These include: application agents, monitoring of real and synthetic transactions, network flow and packet capture, analytics, and dashboard portals for the visualization of results.

At a high level, there are two basic classes of tools. The first class of tool monitors global parameters such as user response time or transaction completion time and provides alerts when thresholds are exceeded. These tools include agents on end user systems and monitoring appliances in the data center. The second class of tool supports triage by monitoring one or more of the components that make up the end-to-end path of the application. These tools include devices that capture application traffic at the flow and packet levels, agents on database, application, and web servers, as well as agents on various network elements.

Each type of individual tool has its strengths and weaknesses. For example, agents can supply the granular visibility that is required for complex troubleshooting but they represent an additional maintenance burden while also adding to the load on the servers and on the network. Monitoring appliances have more limited visibility, but they don't require modification of server configurations and don't add traffic to the network. Taking into consideration these trade-offs, IT organizations need to make tool decisions based on their goals for application performance management, their application and network environment as well as their existing infrastructure and network management vendors.

Independent of the approach that IT organizations take towards application performance management, a critical component of application performance management is end-to-end visibility.

End-to-end visibility refers to the ability of the IT organization to examine every component of IT that impacts communications once users hit ENTER or click the mouse button until they receive a response back from the application.

End-to-end visibility is one of the cornerstones of assuring acceptable application performance. This functionality is important because it:

- Provides the information that allows IT organizations to notice application performance degradation before the end user does.

- Identifies the symptoms of the degradation and as a result enables the IT organization to reduce the amount of time it takes to identify and remove the causes of the degraded application performance.

- Facilitates making intelligent decisions and getting buy-in from other impacted groups.

- Allows the IT organization to measure the performance of a critical application before, during and after a change is made.

[1] This refers to managing the performance of applications that are delivered over WAN services such as Frame Relay, ATM and MPLS.

The value of providing end-to-end visibility is maximized if two criteria are met. One criterion is that all members of the IT organization use the same tool or set of tools. The second criterion is that the tool(s) are detailed and accurate enough to identify the sources of application degradation.

Application Performance Management in Public and Hybrid Clouds

There are a number of possible ways that an IT organization can adjust their application performance management strategies in order to accommodate accessing services hosted by a Cloud Computing Service Provider (CCSP). These include:

- Extend the enterprise monitoring solutions into the public cloud using agents on virtual servers and by using virtual appliances.
- Focus on CCSPs that offer either cloud resource monitoring or application performance management as a service.
- Increase the focus on service delivery and transaction performance by supplementing existing application performance management solutions with capabilities that provide an outside-in service delivery view from the perspective of a client accessing enterprise applications or cloud applications over the Internet or mobile networks.

Application Aware Network Performance Management

There are a number of indications that enterprise networks and the applications that transit these networks are becoming increasingly entwined. One indication is HP's implantation of an SDN app store. The goal of the SDN app store is to enable network operations teams to download applications into their SDN controllers in the same way that smartphone users download apps onto their devices. Another indication is the formation of the Unified Communications Interoperability Forum (UCIF). UCIF's mission is to work with the Open Networking Foundation (ONF) to develop a standardized way for applications to dynamically request services from a software defined network. Yet another indication is the number of network vendors who have announced a network marketecture build around concepts such as *application-aware, application-centric or application driven.*

In response to the fact that enterprise networks and the applications that transit these networks are becoming increasingly entwined, there has been a movement to bring together two management disciplines: Application Performance Management and Network Performance Management. The result of bringing together those two disciplines is a new discipline that is often referred to as Application Aware Network Performance Management (AANPM). An AANPM solution integrates data that has historically been associated with application performance management with data that has historically been associated with network performance management. The result is a system that provides cross-platform visibility that enables IT organizations to monitor, troubleshoot and analyze both network and application systems.

DevOps

The phrase *DevOps* is a result of bringing to together two phrases: *Development* and *Operations*. That's appropriate because the point of adopting DevOps is to establish tight collaboration between a number of the phases of the application development lifecycle, including application development, testing, implementation and ongoing operations. DevOps is not a technology, but an approach. Some of the key characteristics of the approach are that the applications development team writes primarily small incremental pieces of code that are tested on an architecture that reflects the production architecture. Ideally, the network on which the software is tested will reflect not just the architecture but also the same characteristics (i.e., delay, packet loss) as the production network.

Implementing DevOps provides many advantages. For example, DevOps can provide business value by enabling companies to experience sustained innovation. Examples of companies that claim to have experienced sustained innovation as a result of implementing DevOps include Twitter, Netflix and Facebook. Implementing DevOps has other advantages. According to a recent Information Week Report, eighty two percent of the IT organizations that implemented DevOps saw at least some improvement in infrastructure stability and eighty three percent saw at least some improvement in the speed of application development.

A number of service providers who are attempting to become more agile have commented on the need for their organization to move away from slow, manual processes. One such provider is Deutsche Telekom. In a recent article, Deutsche Telekom was quoted as saying: "DT [Deutsche Telekom] needs to build a team that comprises IP, datacenter, programming, and operations specialists that can work in small, empowered, and agile teams, while both the carriers and vendors need to adjust for the migration from hardware-based to software-based business models."

GE Capital is also an advocate of DevOps. In a recent blog, GE Capital's CTO Eric Reed explained some of the impact of DevOps on the IT organization. According to Eric, "Our experience [GE Capital's] on this journey to date has been that the small, self-directed teams required in a DevOps world require an amalgamation of skills spanning everything from IT security to database design and application architecture, plus everything in between. While each individual on the team has a particular strength (say, application design and coding), each one also needs to have working knowledge in other areas (maybe UX or network design)."

Security

How IT Organizations are Implementing Security

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

Security is both a first and a second-generation application and service delivery challenge and it will remain a significant challenge for the foreseeable future. Rapid changes in IT, such as those created by the adoption of cloud computing, social networking and the new generation of mobile devices, combined with the ongoing evolution of regulations pose a spate of new challenges for IT security systems and policies in much the same manner that they present challenges to the IT infrastructure.

IT security systems and policies have evolved and developed around the traditional application delivery architecture in which branch offices are connected using a private WAN service to application servers in a central corporate data centers. 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 cloud computing, applications and services are moving out of the central corporate data center and there is no longer a convenient single location for security policies and systems.

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. 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).

Current State of DDoS Attacks

There is a wide range of ways that a DDoS attack can cause harm to an organization, including the:

- Consumption of computational resources, such as bandwidth, disk space, or processor time;
- Disruption of configuration information, such as routing information;
- Disruption of state information, such as the unsolicited resetting of TCP sessions;
- Disruption of physical network components;
- Obstructing the communication media between the intended users and the victim so that they can no longer communicate adequately.

As noted in the preceding chapter of ***The Handbook***, the majority of the security attacks are to a data center because that's where most of the applications and most of the data resides**.** Every year a number of organizations publish an annual report on the state of security attacks. A number of recent reports have highlighted the evolving threats that are associated with DDoS attacks on data centers.

One of the key findings of those reports is that DDoS attacks are growing in a variety of ways, including:

- Frequency: A 50% increase in DDoS attacks on a year-to-year basis[2];
- Size: One attack out of three is over 20 Gbps, 60 Gbps attacks are common and 100 Gbps attacks are not uncommon[3];
- Severity: The average number of packets per second in a DDoS attack has increased 1,850% to 7.8 Mpps between 2011 and 2013[4];
- Sophistication: 81% of attacks are multi-vector threats[5];
- Persistence: The average duration of a DDoS attack is 17 hours[6].

One example of the state of DDoS attacks is found in the Q1 2014 report from Prolexic. That report mentions a 10 hour long DDoS assault that peaked at over 200 Gbps and 53.5 MPPS.

There are a many components of a DDoS attack, but one component that is common to all such attacks is having a large scale botnet network that sends traffic towards the target at very high rates. As mentioned, it is possible to rent a botnet. The usual way this works is that criminal syndicates and commercially-motivated hackers build botnet networks that can be rented on-demand over the Internet. These on-demand networks are typically known as *booters*[7] and are often marketed at Web performance test tools or *stressers*.

[2] Akamai State of The Internet report
[3] Neustar Annual DDoS Attacks and Impact Report
[4] Verizon data breach report 2014
[5] Incapsula 2013-2014 DDoS Threat Landscape Report
[6] Prolexic Global Attack Report Q1 2014
[7] Safetyskyhacks Top 10 DDosers, Booters, Stressers

Cloud-Based Security

One way that a Cloud-based Security Service (CBSS) could provide value is if it provides protection against the growing number of malware attacks.   To effectively protect against malware attacks, a CBSS should be able to identify suspicious content or sites that are either suspicious or are known to distribute malware. In order to be effective, a CBSS that provides Web content filtering or malware protection needs a source of intellectual capital that identifies known and suspected vulnerabilities. This source needs to be both dynamic and as extensive as possible.

One part of the value proposition of a CBSS that provides security functionality is the same as the value proposition of any cloud based service. For example, a security focused CBSS reduces the capital investment in security that an organization would have to make. In addition, a security focused CBSS reduces the amount of time it takes to deploy new functionality. The speed at which changes can be made to a CBSS adds value in a variety of situations, including providing better protection against zero-day attacks. Another part of the value proposition of a security focused CBSS is that unlike a traditional security solution that relies on the implementation of a hardware based proxy, a CBSS can also protect mobile workers. The CBSS does this by leveraging functionality that it provides at its Points of Presence (POPs) as well as functionality in a software agent that is deployed on each mobile device.

In many instances, the best security solution is a hybrid solution that combines traditional on-premise functionality with one or more Cloud-based solutions. For example, in many cases IT organizations already have functionality such as Web filtering or malware protection deployed in CPE at some of their sites. In this case, the IT organization may choose to implement a CBSS just to protect the sites that don't have security functionality already implemented and/or to protect the organization's mobile workers. Alternatively, an organization may choose to implement security functionality in CPE at all of their sites and to also utilize a CBSS as part of a defense in depth strategy.

Web Application Firewall Services

The chapter of this report entitled *Network and Application Optimization*, discussed how a Cloud-based service, such as the one shown in **Figure 1**, can be used to optimize the performance of the Internet. As will be discussed in this sub-section of the handbook, that same type of service can also provide security functionality.

Figure 1: Internet Based Security Functionality

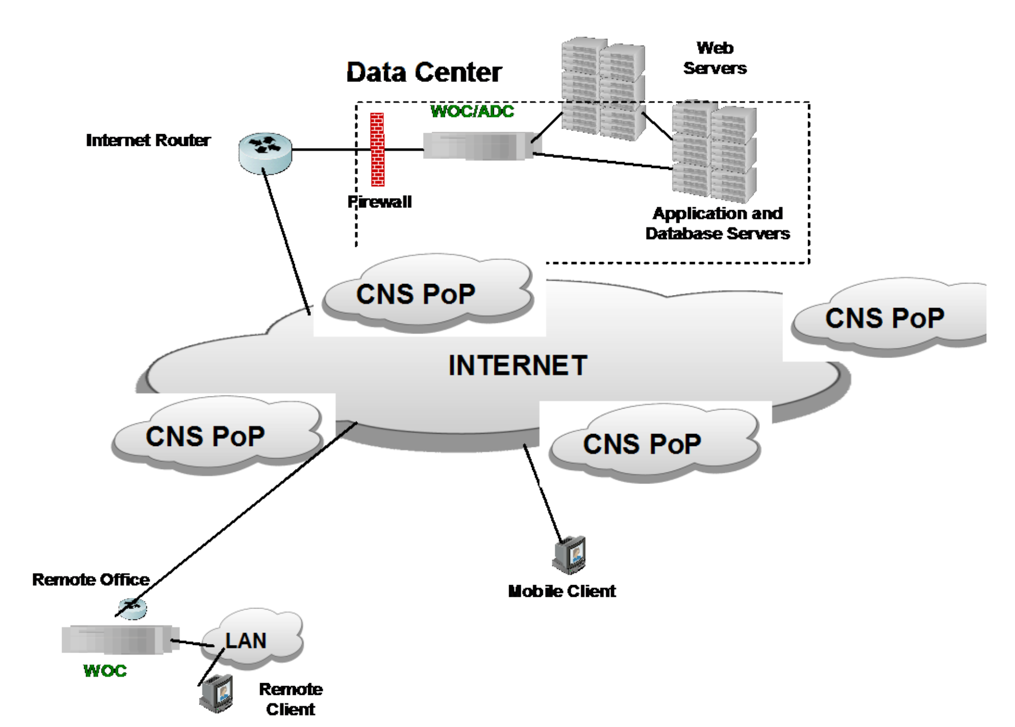

Role of a Traditional Firewall: Protect the Perimeter

Roughly twenty years ago IT organizations began to implement the first generation of network firewalls, which were referred to as packet filters. These devices were placed at the perimeter of the organization with the hope that they would prevent malicious activities from causing harm to the organization.

Today most network firewalls are based on *stateful* inspection. A *stateful* firewall holds in memory attributes of each connection. These attributes include such details as the IP addresses and ports involved in the connection and the sequence numbers of the packets traversing the connection. One of the weaknesses associated with network firewalls is that they are typically configured to open up ports 80 and 443 in order to allow passage of all HTTP and SSL traffic. Given that ports 80 and 443 are generally configured to be open, this form of perimeter defense is porous at best.

Whereas network firewalls are focused on parameters such as IP address and port numbers, a more recent class of firewall, referred to as a Web application firewall, analyzes messages at layer 7 of the OSI model. Web application firewalls are typically deployed as a hardware appliance and they sit behind the network firewall and in front of the Web servers. In some cases, Web application firewall functionality is provided by an Application Delivery Controller (ADC).

Web application firewalls look for violations in the organization's established security policy. For example, the firewall may look for abnormal behavior, or signs of a known attack. It may also be configured to block specified content, such as certain websites or attempts to exploit known security vulnerabilities. Because of their ability to perform deep packet inspection at layer 7 of the OSI model, a Web application firewall provides a level of security that cannot be provided by a network firewall.

Defense in Depth: The Role of a Web Application Firewall Service

As is well known, there are fundamental flaws with an approach to security that focuses only on the perimeter of the organization. To overcome these flaws, most IT organizations have moved to the previously referenced approach to security: *defense in depth*. The concept of defense in depth is not new. What is new in the current environment is the use of a CBSS to provide Web application firewall functionality that is distributed throughout the Internet and to use this functionality to supplement security functionality that is provided on site by devices such as a Web application firewall or an ADC that offers Web application firewall functionality.

Because there are a variety of possible DDoS attacks, IT organizations need to implement a variety of defense in depth techniques. This includes:

- **Minimizing the points of vulnerability**
  If an organization has most or all of its important assets in a small number of locations, this makes the organization more vulnerable to successfully being attacked as the attacker has fewer sites on which to concentrate their attack.

- **Protecting DNS**
  Many IT organizations implement just two or three DNS servers. As such, DNS is an example of what was discussed in the preceding bullet – how IT organization are vulnerable because their key assets are located in a small number of locations.

- **Implementing robust, multi-tiered failover**
  Many IT organizations have implemented disaster recovery plans that call for there to be a stand-by data center that can support at least some of the organization's key applications if the primary data center fails. Distributing this functionality around a global network increases overall availability in general, and dramatically reduces the chance of an outage due to a DDoS attack in particular.

In order to be effective, a CBSS that provides Web application firewall functionality needs to be deployed as broadly as possible, preferably in thousands of locations. When responding to an attack, the service must also be able to:

- Block or redirect requests based on characteristics such as the originating geographic location and whether or not the originating IP addresses are on either a whitelist or a blacklist.

- Direct traffic away from specific servers or regions under attack.

- Issue slow responses to the machines conducting the attack. The goal of this technique, known as tarpits[8], is to shut down the attacking machines while minimizing the impact on legitimate users.

- Direct the attack traffic back to the requesting machine at the DNS or HTTP level.

As noted, a CBSS that provides Web application firewall functionality is complimentary to premise-based security functionality such as that provided by an ADC that offers Web application firewall. That follows because while the Cloud-based Web application firewall service can perform many security functions that cannot be performed by an on premise Web application firewall, there are some security functions that are best performed by an on premise Web application firewall. An example of that is protecting an organization against information leakage by having an onsite Web application firewall perform deep packet inspection to detect if sensitive data such as a social security number or a credit card number is leaving the site. If sensitive data is leaving the site, the onsite Web application firewall, in conjunction with other security devices, can determine if that is authorized and if it is not, it can prevent the data from leaving the site.

Impact of SDN

As was the case with management, SDN poses both security challenges and security opportunities. The primary security challenge is to ensure that an attacker cannot compromise the central SDN controller and hence have access to all of the subtending network elements. In addition to securing the controller itself, all communication between the controller and other devices including switches, network services platforms and management systems must be secured.

As noted, SDN also presents opportunities to improve security by implementing security related applications that leverage the control information that has been centralized in the SDN controller. One such application that has been announced[9] is a network service that provides DDoS protection. Another such example is an application[10] that was designed to combat the security challenges that are associated with BYOD.

The document entitled *Mock RFI for Enterprise SDN Solutions* contains a number of questions that IT organizations should ask SDN vendors relative to security.

[8] Wikipedia Tarpit(networking)
[9] Radware Defense Flow Report
[10] HP Network Protector SDN Application

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

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.

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

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.

fast path
6WIND
FORD

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

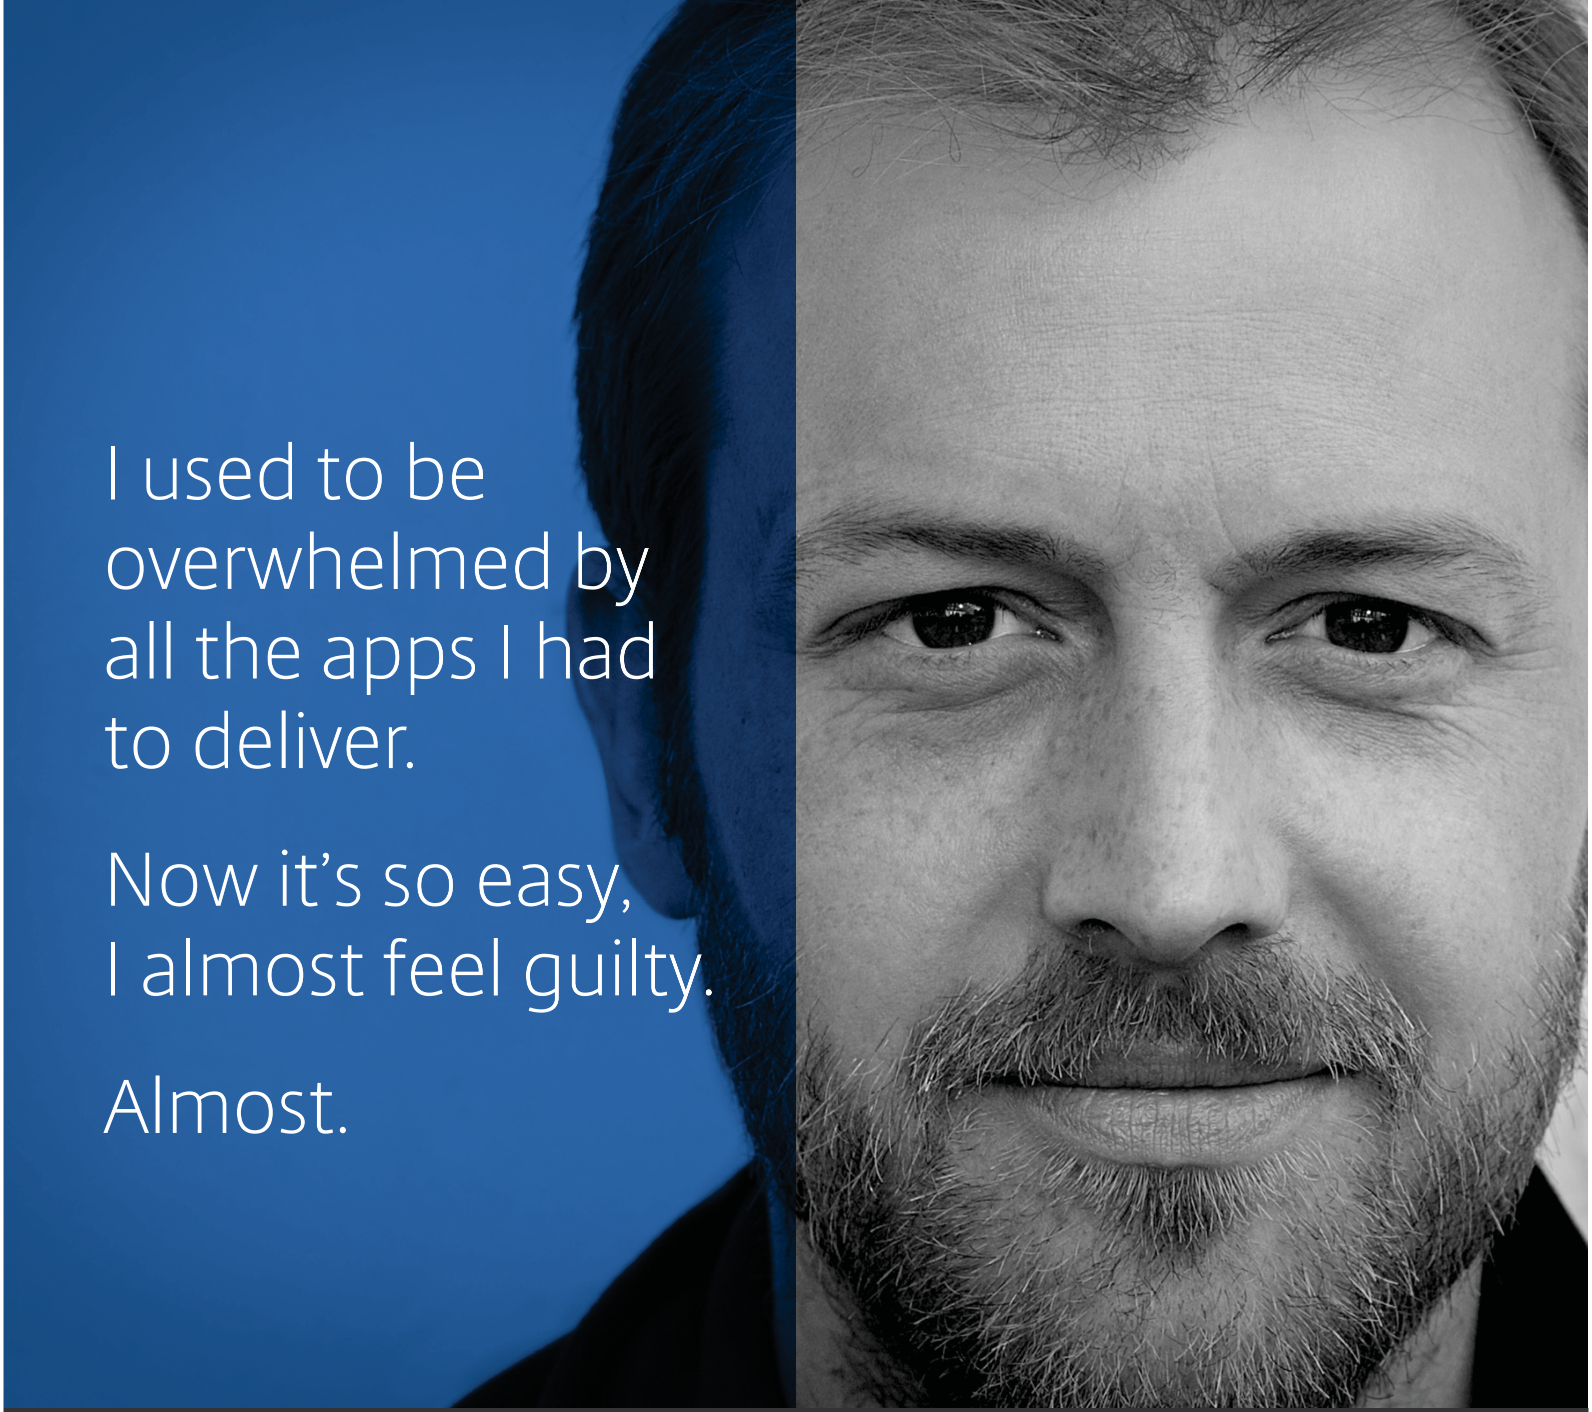
I used to be overwhelmed by all the apps I had to deliver.
Now it's so easy, I almost feel guilty.
Almost.

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.

radware

Alteon NG

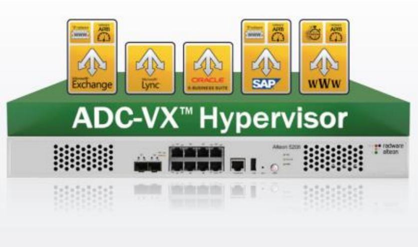

Predictable Application SLA, Guaranteed. Only with Alteon NG.

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.