The 2013
Application & Service Delivery Handbook

Executive Summary

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

Platinum Sponsors:

Gold Sponsors:

Produced by:

Executive Summary

Introduction

Throughout the ***2013 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, IT organizations don't want to optimize the performance of malware and spyware. IT organizations must identify this traffic and eliminate it.

At the same time that many IT organizations are still in the process of implementing solutions that respond to the first generation of application delivery challenges such as transmitting large files between a branch office and a data center, a second generation of challenges is emerging. These challenges are 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

The goal of the 2013 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 second generation of application and service delivery challenges. To help to achieve this goal, in early 2013 two surveys were given to the subscribers of Webtorials. Throughout this document, the IT professionals who responded to the surveys will be referred to as The Survey Respondents.

Second Generation Application and Service Delivery Challenges

There are a number of fairly well understood challenges that have over the years complicated the task of ensuring acceptable application and service delivery. These challenges are referred to in The Handbook as First Generation Application and Service Delivery Challenges. Examples of those challenges include:

- Network Latency
- Bandwidth Constraints
- Packet Loss
- Chatty Protocols and Applications

Since these challenges are fairly well understood, they are listed in The Handbook, but not described. The handbook does include a reference to a detailed description of these challenges.

The Handbook describes a number of second-generation challenges that are beginning to complicate the task of ensuring acceptable application and service delivery. Those key challenges that are described in The Handbook are:

- Mobility and BYOD
- Virtualization
- Cloud Computing

Mobility and BYOD

In order to quantify the impact of mobility, The Survey Respondents were asked a couple of questions. One question was: "In some cases employees of a company access business related data and applications by using a mobile device within a company facility and, in some cases, employees access business related data and applications by using a mobile device when they are at an external site. In the typical day, what percentage of your organization's employees use a mobile device at some time during the day to access business related data and applications, either from within a company facility or from an external site?" Their responses are show in **Table 1.**

Table 1: Amount of Mobile Access

	0%	1% to 9.99%	10% to 24.995	25% to 49.99%	50% to 74.99%	75% to 99.99%	100%
Company Facility	**6%**	**14%**	**26%**	**19%**	**22%**	**10%**	**4%**
External Site	**2%**	**23%**	**20%**	**20%**	**14%**	**15%**	**6%**

The data in **Table 1** indicates that the vast majority of employees require mobile access for at least part of their typical day.

The Survey Respondents were also asked to indicate the types of employee owned devices that their organization allows to connect to their branch office networks and which of these devices is actively supported, Their responses are shown in **Table 2.**

Table 2: Support for Employee Owned Devices

	Not Allowed	Allowed but not Supported	Allowed and Supported
Company managed, employee owned laptop	22%	24%	54%
Employee owned and managed laptop	38%	38%	25%
Blackberry	17%	24%	58%
Apple iPhone	14%	30%	55%
Android phone	19%	33%	48%
Windows mobile phone	26%	40%	34%
Apple iPad	18%	40%	52%
Android based tablet	28%	37%	35%
Windows based tablet	28%	36%	37%

The data in **Table 2** indicates that there is wide acceptance BYOD in general and that there is a broad range of mobile devices that IT organizations must support. Unfortunately, this new generation of 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 software) and network intrusion protection (e.g., personal firewall), but these protections are just now emerging for smartphones and tablet computers[1]. Similarly, inventorying and updating installed software on smartphone and tablet computers are emerging capabilities and a critical area for Mobile Device Management solutions.

Virtualization

The Handbook analyzed two forms of virtualization: Server Virtualization and Desktop Virtualization.

Server Virtualization

One of the challenges associated with server virtualization comes from the fact that in most cases, data centers with virtualized servers will have different hypervisors that each has their own management capabilities. Another challenge is the need to integrate the management of virtual servers into the existing workflow and management processes and over half of The Survey Respondents indicated that they consider it to be either very or extremely important over the next year for their organization to get better at performing management tasks such as troubleshooting on a per-VM (Virtual Machine) basis.

[1] http://www.computerworld.com/s/article/9224244/5_free_Android_security_apps_Keep_your_smartphone_safe)

In addition, one of the advantages of a virtualized server is that a production VM can be dynamically transferred to a different physical server without service interruption. However, in the current environment, the supporting network and management infrastructure is still largely static and physical. So while it is possible to move a VM between data centers in a matter of seconds or minutes, it can take days or weeks to get the network and management infrastructure in place that is necessary to enable the VM to be useful.

Desktop Virtualization

There are two primary approaches to server-side virtualization. They are:

- Server Based Computing (SBC)
- Virtual Desktop Infrastructure (VDI)

While there are advantages to both forms of desktop virtualization, the vast majority of virtualized desktops will utilize server side virtualization.

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. Ensuring acceptable performance for desktop virtualization presents some significant challenges. One such challenge is that, as is the case in with any TCP based application, packet loss causes the network to retransmit packets. This can dramatically increase the time it takes to refresh a user's screen.

Cloud Computing

The Handbook details the three primary classes of cloud computing solutions:

- Private Cloud
- Public Cloud
- Hybrid cloud

Private Cloud Computing

One of the primary ways that IT organizations have adopted private cloud computing solutions is by implementing some or all of the key characteristics of public cloud computing solutions in order to be able to provide Infrastructure-as-a-Service (IaaS) solutions that are similar to the solutions offered by IaaS providers such as Rackspace. The Survey Respondents were given a set of 7 possible approaches to IaaS and were asked to indicate which approach best described their company's approach to using IaaS solutions, either provided internally by their own IT organization, or provided externally by a Cloud Computing Service Provider (CCSPs). The survey results indicate that only a small percentage of IT organizations have a strategy for how they will acquire or implement IaaS solutions.

Public Cloud Computing

The Handbook focuses on the two most popular types of public cloud computing solutions: Software-as-a-Service (SaaS) and Infrastructure-as-a-Service (IaaS).

The Survey Respondents were asked to indicate the factors that are driving their organization to adopt SaaS or IaaS solutions. Table 3 lists the factors that are driving the adoption of SaaS in descending order of importance. There is little difference between these factors and the factors that are driving the adoption of IaaS.

Table 3: Factors Driving SaaS Adoption
Factor
Lower cost
Reduce the amount of time it takes to implement an application
Free up resources in the IT organization
Deploy applications that are more robust; e.g., available and scalable
Easier to justify OPEX than CAPEX
Leverage the expertise of the SaaS provider
Reduce risk
Management mandate as our strategic direction
Meet temporary requirements
Other

According to The Survey Respondents, concern about the security and confidentiality of data is by a wide margin the number one factor inhibiting the adoption of any form or public cloud solutions

Hybrid Cloud Computing

The adoption of public and/or hybrid cloud computing solutions creates a new set of management challenges for enterprise IT organizations. Some of these new challenges stem from the fact that IT organizations are typically held responsible for the performance of these solutions even though in most cases they don't have the same access to the enabling IT infrastructure that they would have if the solution was provided internally. Other new management challenges stem from the sheer complexity of the public and hybrid cloud environments. What this complexity means is that in order to manage end-to-end in either a public cloud or a hybrid cloud environment, management data must be gathered from the enterprise, one or more Network Service Providers (NSPs) and one or more CCSPs.

Network and Application Optimization

As shown in **Figure 1**, application response time (R) is impacted by a number of factors including the amount of data being transmitted (Payload), the goodput which is the actual throughput on a WAN link, the network round trip time (RTT), the number of application turns (AppTurns), the number of simultaneous TCP sessions (concurrent requests), the server side delay (Cs) and the client side delay (Cc).

Figure 1: Application Response Time Model

$$R \approx \frac{Payload}{Goodput} + \frac{(\#\ of\ AppsTurns * RTT)}{Concurrent\ Requests} + Cs + Cc$$

The WAN Optimization Controllers, Application Delivery Controllers and WAN Optimization Solutions that are described in this section of The Handbook are intended to mitigate the impact of those factors.

WAN Optimization Controllers (WOCs)

Table 4 lists some of WAN characteristics that impact application delivery and identifies WAN optimization techniques that a WOC can implement to mitigate the impact of those characteristics. These techniques are described in detail in The Handbook and The Handbook also provides a suggested approach for evaluating WOCs.

Table 4: Techniques to Improve Application Performance

WAN Characteristics	WAN Optimization Techniques
Insufficient Bandwidth	Data Reduction: • Data Compression • Differencing (a.k.a., de-duplication) • Caching
High Latency	Protocol Acceleration: • TCP • HTTP • CIFS • NFS • MAPI Mitigate Round-trip Time • Request Prediction • Response Spoofing
Packet Loss	Congestion Control Forward Error Correction (FEC) Packet Reordering
Network Contention	Quality of Service (QoS)

As described in The Handbook, WOCs come in a variety of form factors including:

- **Standalone Hardware/Software Appliances**
  These are typically server-based hardware platforms that are based on industry standard CPUs with an integrated operating system and WOC software.

- **Client software**
  WOC software can also be provided as client software for a PC, tablet or Smartphone to provide optimized connectivity for mobile and SOHO workers.

- **Integrated Hardware/Software Appliances**
  This form factor corresponds to a hardware appliance that is integrated within a device such as a LAN switch or WAN router via a card or other form of sub-module.

- **Virtual WOCs**
  The phrase virtual WOC refers to optimizing the operating system and the WOC software to run in a VM on a virtualized server.

Application Delivery Controllers

Among the functions users can expect from an ADC are the following:

- **Traditional SLB**
  ADCs can provide traditional load balancing across local servers or among geographically dispersed data centers based on Layer 4 through Layer 7 intelligence.

- **SSL Offload**
  One of the primary new roles played by an ADC is to offload CPU-intensive tasks from data center servers. A prime example of this is SSL offload, where the ADC terminates the SSL session by assuming the role of an SSL Proxy for the servers.

- **XML Offload**
  XML is a verbose protocol that is CPU-intensive. Hence, another function that can be provided by the ADC is to offload XML processing from the servers by serving as an XML gateway.

- **Application Firewalls**
  ADCs may also provide an additional layer of security for Web applications by incorporating application firewall functionality.

- **Denial of Service (DOS) Attack Prevention**
  ADCs can provide an additional line of defense against DOS attacks, isolating servers from a range of Layer 3 and Layer 4 attacks that are aimed at disrupting data center operations.

- **Asymmetrical Application Acceleration**
  ADCs can accelerate the performance of applications delivered over the WAN by implementing optimization techniques such as reverse caching, asymmetrical TCP optimization, and compression.

- **Response Time Monitoring**
  The application and session intelligence of the ADC also presents an opportunity to provide real-time and historical monitoring and reporting of the response time experienced by end users accessing Web applications.

The Handbook describes the techniques used within ADCs and also provides a suggested approach for evaluating ADCs.

WAN Optimization Solutions

Cloud-Based Optimization Solutions

As shown in **Figure 2**, it is now possible to acquire a number of IT-centric functions, such as network and application optimization from a cloud service provider.

Figure 2: Network and Application Optimization Solution

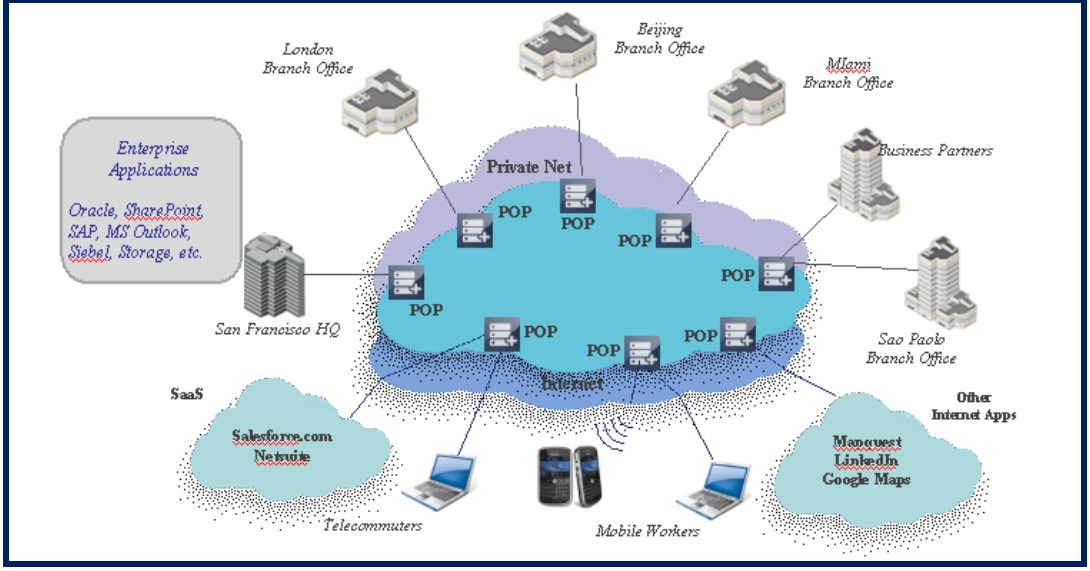

As shown in **Figure 2**, a variety of types of users (e.g., mobile users, branch office users) access WAN optimization functionality at the service provider's points of presence (POPs). Ideally these POPs are inter-connected by a dedicated, secure and highly available network. To be effective, the solution must have enough POPs so that there is a POP in close proximity to the users. One use case for a service such as this is the previously mentioned mandate to support mobile workers.

The Optimization of Internet Traffic

WOCs make the assumption that performance characteristics within the WAN are not capable of being optimized because they are determined by the relatively static service parameters controlled by the WAN service provider. This assumption is reasonable in the case of private WAN services such as MPLS. However, this assumption does not apply to enterprise application traffic that transits the Internet because there are significant opportunities to optimize performance within the Internet itself. Throughout The Handbook, a service that optimizes Internet traffic will be referred to as an Internet Optimization Service (IOS).

An IOS would, out of necessity, leverage service provider resources that are distributed throughout the Internet in order to optimize the performance, security, reliability, and visibility of the enterprise's Internet traffic. The servers at the IOS provider's PoPs perform a variety of optimization functions. Some of the functions provided by the IOS include:

- **Route Optimization**
  A route optimization solution leverages the intelligence of the IOS servers that are deployed in the service provider's PoPs to measure the performance of multiple paths through the Internet and to choose the optimum path from origin to destination.

- **Transport Optimization**
  TCP performance can be optimized by setting retransmission timeout and slow start parameters dynamically based on the characteristics of the network such as the speed of the links and the distance between the transmitting and receiving devices.

- **HTTP Protocol Optimization**
  HTTP inefficiencies can be eliminated by techniques such as compression and caching at the edge IOS server with the cache performing intelligent pre-fetching from the origin.

- **Content Offload**
  Static content can be offloaded out of the data-center to caches in IOS servers and through persistent, replicated in-cloud storage facilities.

An Integrated Private-Public WAN

The traditional approach to providing Internet access to branch office employees has been to backhaul that Internet traffic on the organization's enterprise network (e.g., their MPLS network) to a central site where the traffic was handed off to the Internet. The advantage of this approach is that it enables IT organizations to exert more control over their Internet traffic and it simplifies management. Disadvantages of this approach are that it results in extra traffic transiting the enterprise's WAN and that it usually adds additional delay to the Internet traffic.

In order to improve performance of backhauled traffic, an IT organization might use WOCs to optimize the performance of the traffic as it flows from the branch office to the central site over their enterprise WAN. However, once the traffic is handed off to the Internet, the traffic is not optimized and the organization gets little value out of optimizing the traffic as it flows over just the enterprise WAN.

One way to minimize the degradation in application performance is to not backhaul the traffic but hand it off locally to the Internet. For this approach to be successful, IT organizations must be able to find another way to implement the security and control that it has when it backhauls Internet traffic. One way that this can be done is to use an IOS to carry traffic directly from the branch office to the SaaS provider. With this approach, in addition to providing optimization functionality, the IOS can provide the security functionality that was previously provided in the corporate data center.

Another approach to optimizing Internet traffic is to implement a form of WAN optimization that enables IT organizations to keep its current approach to backhauling traffic, but which eliminates the performance issues surrounding the fact that once the traffic is handed off to the Internet, the traffic is typically no longer optimized. For this approach to work, the optimization that is in place for enterprise WANs must be integrated with the optimization that is provided by the IOS.

Hybrid WANs

The key concept behind Hybrid WANs is to create end-to-end WAN connections either based on multiple WAN services (e.g., MPLS, Frame Relay and the Internet) or based just on multiple Internet connections. Part of the value proposition of a hybrid WAN is that traffic is allocated across alternative paths based on real-time traffic analytics, including:

- The instantaneous end-to-end performance of each available network: This allows the solution to choose the optimal network path for differing traffic types.

- The instantaneous load for each end-to-end path: Typically the load is weighted based on the business criticality of the application flows.

- The characteristics of each application: This includes the type of traffic (e.g., real time, file transfer); the performance objectives for delay, jitter and packet loss; as well as the business criticality and information sensitivity.

One of the primary reasons why IT organizations backhaul their Internet traffic to a central site over an enterprise WAN service is because of security concerns. In order to mitigate those concerns when using a hybrid WAN for direct Internet access, the hybrid WAN should support security functionality such as encryption.

Management and Security

Management

The Handbook identifies the challenges that are forcing IT organizations to change how they manage applications and services. In this context, a *service* is comprised of the following four components:

- Either a multi-tier application or multiple applications
- Supporting protocols
- Enabling network services; e.g., DNS, DHCP
- The end-to-end network

The key challenges discussed in The Handbook include:

- Server Virtualization
- Cloud Balancing
- Delay Sensitive Traffic
- Converged Infrastructure

As pointed out in The Handbook, since any component of a complex service 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; i.e., Cisco's UCS and VCE's Vblock platforms. 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 NSPs and one or more 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 supply chain management and relate these business level KPIs to the performance of the IT services that support the business processes.

In addition to discussing how application performance management can be performed within an enterprise's IT environment, The Handbook also identifies 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 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.

Security

The security landscape has changed dramatically in the last few years. In the not too distant past, the typical security hacker worked alone, relied on un-sophisticated techniques such as dumpster diving, and was typically motivated primarily 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 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.

To help IT organizations understand how their approach to security compares to what others are doing, The Handbook identifies the approach to security being taken by the majority of IT organizations. In particular, The Handbook uses input from The Survey Respondents to identify:

- The network security systems currently in place
- The approaches that IT organizations are taking to comprehensive IT security
- Alternative ways to support employee owned devices
- Techniques for implementing full disc encryption
- Approaches for implementing Identity and Access Management
- Security governance models

The Handbook also discusses the use of cloud-based security services (CBSSs). As discussed in The Handbook, one part of the value proposition of a CBSS is the same as the value proposition of any public cloud service. For example, a CBSS reduces the investment in security that an organization would have to make. In addition, a 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[2]. 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 POPs as well as functionality in a software agent that is deployed on each mobile device.

Also included in The Handbook is a discussion of Web application firewalls. As pointed out in The Handbook, whereas network firewalls are focused on parameters such as IP address and port numbers, 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. They 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.

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

[2] http://en.wikipedia.org/wiki/Zero-day_attack

to an approach to security that is typically referred to as *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. This means that Web application functionality is close to the source of security attacks and hence can prevent many security attacks from reaching the organization.

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

Customer Driven Innovation

Cost Effective Cloud Computing

Reduce Your Application Delivery Costs

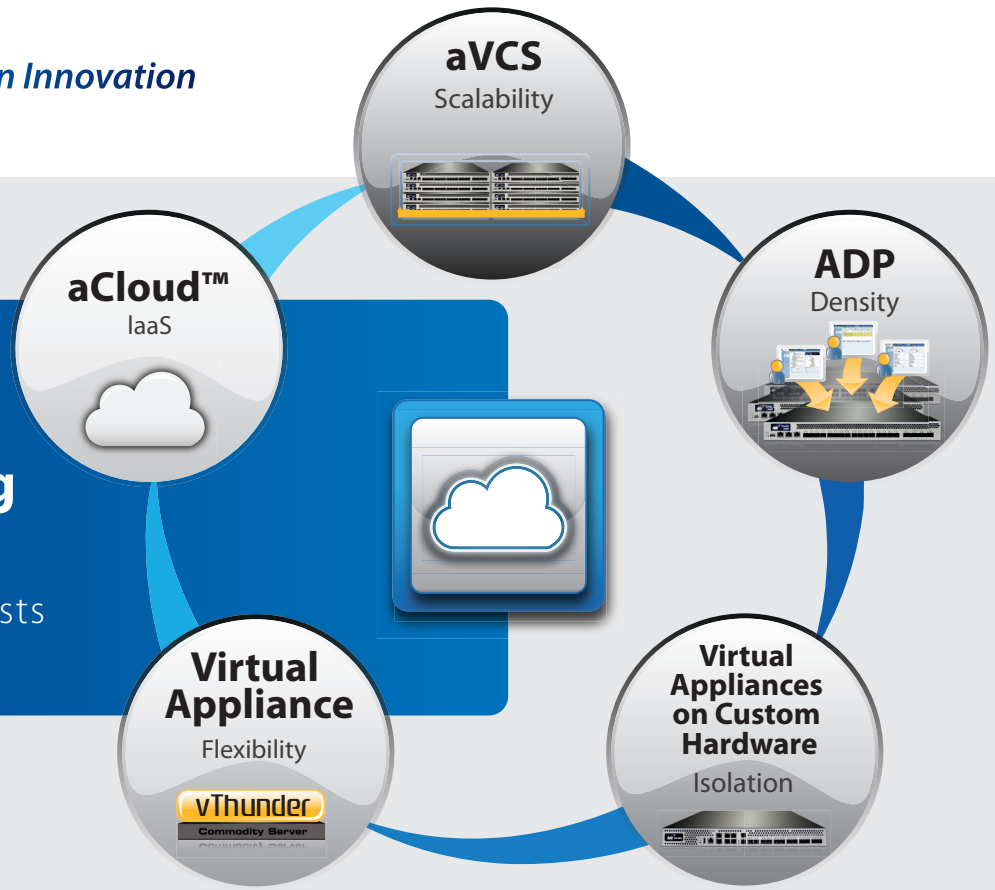

A10 Thunder™ and AX Series Products & Solutions

Based on A10's award-winning Application Delivery Controllers (ADCs) and Advanced Core Operating System (ACOS™) architecture, enterprises and service providers will have the flexibility to choose the following scale-as-you-grow virtualization options.

Virtual Appliances

- Virtual machine (VM) on commodity hardware
- Rapidly scale with commodity hardware
- Reduce hardware costs and upload to compatible cloud providers

Virtual Chassis System

- Cluster multiple ADCs to operate as a unified single device
- Scale while maintaining single IP management
- Reduce costs and simplify management while adding devices as you grow

Virtual Appliances on Custom Hardware

- Hardware appliances with embedded hypervisor running virtual appliances
- Flexibility with hardware performance and reliability

Application Delivery Partitions

- Partition the ADC platform resources for individual applications
- Enable quality multi-tenancy with granular resource allocation
- Reduce the number of appliances to host multiple applications

Mission-critical app? You need CA Application Performance Management

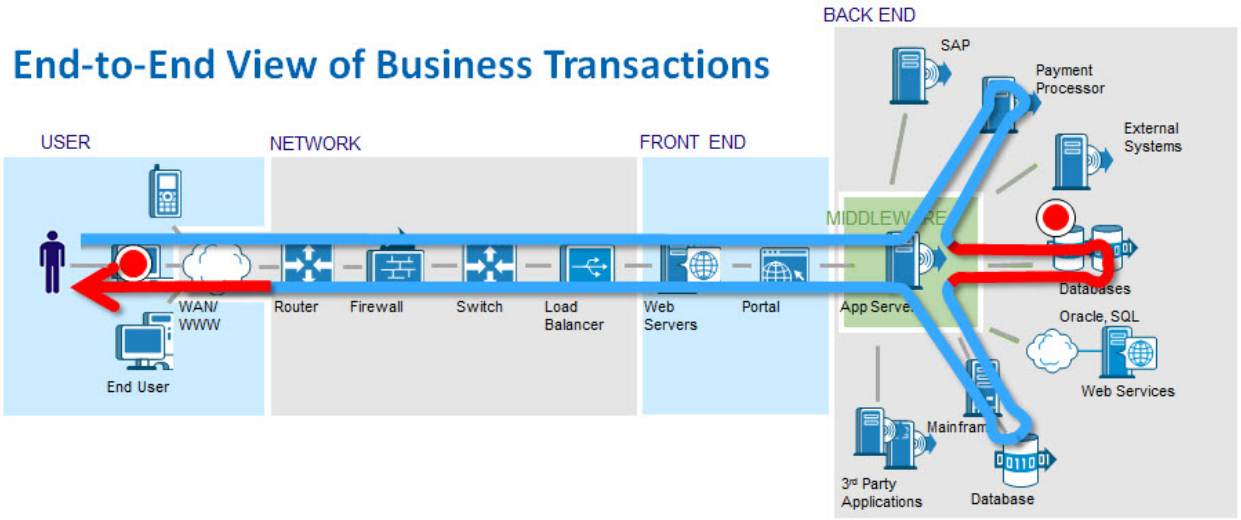

What does an hour of an application outage or downtime cost your business? For internal systems, it's lost productivity. For external systems it's lost customers and lost revenue. Let's face it: neither is good for the bottom line. Applications are the lifeblood of virtually all organizations today, meaning they require a robust application performance management solution that helps ensure end users get the experience they expect and business services are reliably delivered. Today's complex, business-critical applications require CA Application Performance Management (CA APM) to ensure business success.

CA APM delivers 360-degree visibility into and analysis of all user transactions across the hybrid-cloud infrastructure – physical, virtual, cloud and mainframe – to understand the health, availability, business impact and end-user experience of critical enterprise, mobile and cloud applications. Advanced Application Behavior Analytics add deeper visibility into the wealth of metric performance collected by CA APM, giving IT operators another set of eyes to look for potential trouble spots. With CA APM, organizations can proactively identify, diagnose and resolve problems throughout the application lifecycle to put organizations firmly in control of the end-user experience and optimize the performance of critical, revenue-generating services.

CA APM is uniquely designed for today's large, complex and heterogeneous production and pre-production environments. CA APM deploys rapidly with little overhead, scales to manage billions of transactions and is easy for both business and IT executives to use, accelerating time to value while increasing return on investment. CA APM is trusted by more than 2,500 large enterprises and service providers to manage the performance and availability of their critical and revenue-generating services. *Learn more at ca.com/apm*.

agility
made possible™

ca technologies

Simplify and Accelerate Virtual Application Deployments with Cisco Cloud Network Services

Cisco and a Multi-vendor Ecosystem Provide Cloud-Ready Application Services

ROLE OF THE NETWORK FOR THE CLOUD
Access to Critical Data, Services, Resources and People • Core fabric connects resources within the data center and across data centers to each other. • Pervasive connectivity links users and devices to resources and each other. • The network provides identity- and context-based access to data, services, resources, and people.
Granular Control of Risk, Performance, and Cost • Manage and enforce policies to help ensure security, control, reliability, and compliance. • Manage and enforce service-level agreements (SLAs) and consistent quality of Service (QoS) within and between clouds, enabling hybrid models and workload portability. • Meter resources and use to provide transparency for cost and performance.
Robustness and Resilience • Supports self-healing, automatic redirection of workload and transparent rollover. • Provide scalability, enabling on-demand, elastic computing power through dynamic configuration.
Innovation in Cloud-Specific Services • Context-aware services understand the identity, location, proximity, presence, and device. • Resource-aware services discover, allocate, and pre-position services and resources. • Comprehensive insight accesses and reports on all data that flows in the cloud.

Overview

Cisco has announced the evolution of its network services strategy for virtual and cloud networks, Cisco® Cloud Network Services, a complete portfolio of application networking and security services built on top of the Nexus® 1000V virtual networking portfolio and part of the Cisco Unified Data Center architecture. Cloud Network Services simplifies and accelerates cloud network deployments without compromising the critical security and application delivery services that critical data center applications require.

Introducing Cisco Cloud Network Services

Advances in cloud computing, data center consolidation, mobility, and big data are imposing new demands on the network, along with demands for greater network simplification and automation.

As virtual networking and programmable overlay networks evolve to meet these challenges, a similar evolution needs to take place in Layer 4 through 7 application networking services to support widespread virtualization, application mobility, cloud architectures and network automation.

Cisco's solution to this challenge is Cisco Cloud Network Services, a portfolio of integrated, application-aware network services offerings designed for virtual and cloud environments. The Cloud Network Services framework eliminates the obstacles of physical service appliances to accommodate the requirements of virtual applications and cloud deployments, such as:

- Limited scalability of physical services in fixed locations

CISCO VIRTUAL NETWORK PORTFOLIO

Routing and Switching

- Cisco Nexus 1000V virtual switch
- Cisco Cloud Services Router (CSR) 1000V

Security and VPN

- Cisco Virtual Security Gateway for Nexus 1000V (included in Nexus 1000V Advanced Edition)
- Cisco Adaptive Security Appliance (ASA) 1000V Cloud Firewall
- Imperva SecureSphere Web Application Firewall

WAN Optimization

- Cisco Virtual Wide Area Application Services (vWAAS)

Network Analysis and Monitoring

- Cisco Prime Virtual Network Analysis Module (NAM)

Application Delivery Controllers

- Citrix NetScaler VPX virtual application delivery controller

Virtual Services Deployment Platform

- Cisco Nexus 1100 Series Cloud Services Platform

Cloud Orchestration and Management

- Cisco Intelligent Automation for Cloud (IAC)
- Cisco Prime Network Controller
- OpenStack

To learn more about Cisco's complete virtual networking portfolio, see http://cisco.com/go/1000v

- Inconsistent application performance based on workload location relative to services
- Difficulty in inserting security and network services into virtual networks
- Lack of control over services and policies for applications deployed at cloud service providers

The Cisco Cloud Network Services portfolio includes the Cisco Adaptive Security Appliance (ASA) 1000V Cloud Firewall, Cisco Virtual Security Gateway (VSG) virtual firewall, Cisco Virtual Wide Area Application Services (vWAAS) WAN optimization solution, and Cisco Prime™ Virtual Network Analysis Module (vNAM). This new architecture provides a complete services portfolio while delivering scale-out architecture, elastic instantiation, and multi-tenant operation, all with a common approach to service provisioning and management.

Cloud Network Services also includes best-in-class third-party virtual service offerings that integrate transparently into the framework. It now includes the Citrix NetScaler VPX virtual application delivery controller (ADC), and the Imperva SecureSphere Web Application Firewall (WAF).

Cisco Cloud Network Services form the virtual network services strategy of the larger Cisco Unified Data Center framework, which brings together a seamless architecture of virtualization and cloud-ready compute servers (Unified Compute Servers), network fabric (Unified Fabric) and automation platform (Unified Management).

Cloud Network Services, based on the Nexus 1000V virtual switch, are designed to run across major hypervisors, including VMware vSphere, Microsoft Hyper-V, and Linux Kernel-based Virtual Machine (KVM). It is also designed to support multiple cloud orchestration and virtualization management platforms, such as VMware vCenter and Microsoft Systems Center Virtual Machine Manager (SCVMM), giving customers greater flexibility.

Cisco Cloud Services Platform

With the evolution to Cisco Cloud Network Services as the Layer 4 through 7 framework for virtual and cloud networks, organizations are increasingly looking for a flexible platform on which to deploy virtual service nodes rather than use existing application servers. The Nexus Cloud Services Platform was created to address this need.

The Cisco Nexus 1100 Cloud Services Platform is a group of Cisco Unified Computing System™ (Cisco UCS) appliances dedicated to running Cloud Network Service nodes. In addition to the virtual services listed earlier, the Nexus 1100 series runs the management platforms for the virtual network, the virtual security module (VSM), and the Cisco Data Center Network Manager (DCNM) application. The Cloud Services Platform can be dynamically configured to allocate its virtual CPUs to each service as needed based on current application and performance requirements. Current models of the Nexus 1100 Cloud Services Platform include the Nexus 1110-S and 1110-X.

vPath: Enabling Services in Virtual and Cloud Networks

vPath is a component of the Cisco Nexus 1000V virtual switch that directs traffic to appropriate virtual service nodes, such as firewalls and ADCs, in the correct order for each application, independent of the topology of the network or the location of the network services. This feature allows greater application mobility and more reliable service delivery (Figure 1).

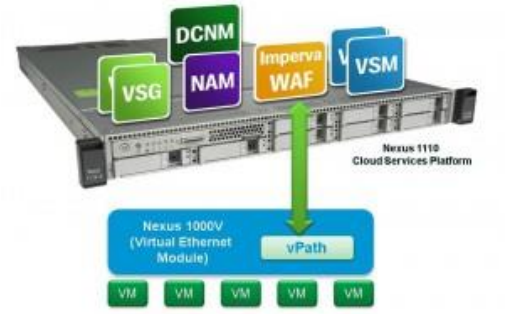

Figure 1 – vPath Connects Virtual Applications to Services Running on the Cisco Cloud Services Platform

Nexus 1000V InterCloud: Enable Hybrid Cloud Connectivity and Cloud-Based Virtual Services

As virtual networks extend from the data center to cloud service providers, organizations are concerned about the consistency of security and application delivery policies and about how these policies are enforced in the cloud. Applications that migrate from the data center to cloud providers can expect different behavior, and organizations may struggle to address compliance issues.

Cisco Nexus 1000V InterCloud complements Cisco Cloud Network Services, allowing seamless hybrid cloud connectivity between data centers and cloud providers and creating one extended network for application and Cloud Network Service deployments.

By deploying Cloud Network Services in all cloud locations, public and private, organizations help ensure consistent policy enforcement, quality of service (QoS) and compliance independent of the location of the virtual applications. Because Cloud Network Services are virtual machines themselves, they are easily deployed within public cloud providers regardless of the infrastructure they provide, and they provide the service consistency required for mission-critical applications.

For More Information

Learn more about Cisco virtual networking portfolio: http://cisco.com/go/1000v

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

About Ipanema Technologies

Application Performance Guarantee

Go Beyond WAN Optimization

About Ipanema

- Selected by worldwide enterprises across all industry sectors.
- One of the largest customer bases (over 150,000 managed sites).
- Visionary in Gartner's WOC Magic Quadrant 2013.
- Leader for Application-Aware Network services (BT, Colt, Vodafone, KDDI, KPN, OBS, Swisscom, Telecom Italia, Telefónica, Easynet…).

79%*

of organizations suffers application performance problems while increasing their IT budget.

*Ipanema Killer Apps survey 2012

Losing 5 minutes

per day for poor application performance means 1% of productivity drop which can turn down profitability by 10%.

"Thanks to Ipanema, our network is totally aligned with our business requirements. With the flexible application-based managed service delivered by e-Qual, we can guarantee the performance of our business critical applications including our ERP and MS Lync, anytime anywhere while reducing our IT costs".

Philippe Faure, Chief Information Officer, Gemalto

Go Beyond WAN Optimization to guarantee your applications performance

Ipanema provides enterprises with a direct connection between application performance and their business requirements.

With Ipanema Technologies, enterprises automatically understand which applications use the network and deliver guaranteed performance to each user. Enterprises can support their strategic IT transformations (like cloud computing and Unified Communications) and control Internet growth while reducing their IT expenses.

Using Ipanema, **enterprises**:

- Guarantee their business application performance;
- Protect unified communications;
- Enable hybrid cloud applications;
- Deploy hybrid networks;
- Control Internet, social media and video traffic;
- Save on IT costs.

Ipanema: the only solution that integrates all the features to guarantee application performance

Ipanema's self-learning and self-optimizing Autonomic Networking System™ (ANS) tightly integrates all the features to guarantee the best application performance:

- Application Visibility,
- Application Control,
- WAN Optimization,
- Dynamic WAN Selection
- And Network Rightsizing.

Do You Have Best-in-Class Application Delivery?

Today's data center challenges extend beyond the traditional needs for application availability, performance and security – challenges well-served previously with classic load balancers / application delivery controllers (ADCs). Nowadays, the adoption of data center virtualization, synchronization with dynamic data center changes, true "awareness" of deployed business applications, and the need for end-to-end visibility– all require a new class of advanced (yet cost-effective) ADCs.

Radware **Alteon® 5224** is an advanced ADC specifically targeted to address all of these challenges. Offering the very latest in next generation application delivery technology with ease of operations, it's simply the best-in-class application delivery choice. Here are four reasons why, we know you'll appreciate:

Reason 1: ADC Virtualization and Consolidation

ADC-VX™, part of Radware's Virtual Application Delivery Infrastructure (VADI)™ strategy, is the industry's first ADC virtualization hypervisor, allowing for the most cost-effective ADC consolidation capabilities. ADC-VX is built on a unique architecture that virtualizes the resources of Radware's ADC including CPU, memory, network and acceleration resources. This specialized hypervisor runs virtual ADC instances (vADC) where each delivers full ADC functionality. Each virtual ADC instance contains a complete and separated environment of resources, OS, configurations and management.

In turn, this allows allocation of a separate, fully-isolated vADC instance for each application. Companies can then maximize application availability and meet application SLA requirement with a resource reservation mechanism. Moreover, this deployment model simplifies operations, reduces the ADC infrastructure footprint, and increases business agility with faster roll out of new vADCs and applications. With vADC per application, application lifecycle management is streamlined and its associated cost is significantly reduced compared to traditional ADC deployment models.

Reason 2: Result-Driven Application Acceleration

Radware's FastView™ result-driven acceleration technology adds Web Performance Optimization (WPO) capabilities on top of standard ADC application acceleration features (e.g., caching, compression, SSL acceleration, etc.), to deliver the fastest Web application response time and ensure best application SLA while offloading server processing. This results in increased revenues, higher conversion rates, higher customer loyalty as well as improved employee productivity when using enterprise web applications. It applies to all browsers, all end-user device types and all users, located anywhere. Radware's leading WPO capabilities include:

- Reducing the # of server requests per page
- Accelerate entire web transaction
- Custom optimization templates for each browser
- Static and dynamic, browser-side caching
- Dedicated, mobile caching based on HTML 5 local storage
- Content minification

Reason 3: End-to-End Application QoE & Performance Visibility

Ensuring applications deliver the best quality of experience requires IT administrators to gain maximum visibility on all application delivery chain components, throughout the life cycle of the application. Radware's multilayer approach for monitoring the application delivery infrastructure, coupled with its integrated application

performance monitoring (APM) module, provide a powerful tool to guarantee continuous high application SLA throughout the entire application life cycle, by displaying actual user transactions and errors. The only APM-integrated ADC on the market, the solution enables easy detection and resolution of SLA degradations, while eliminating the need to manually script synthetic transactions. Cross-ADC infrastructure historical reports on resource utilization provide a holistic view enabling better capacity planning when rolling out new applications. In addition, drilldown-able real-time dashboards, that span multiple ADCs, enable instant visibility for spotting problems and a powerful tool for fast and accurate troubleshooting.

Reason 4: Application Awareness with AppShape™ & AppShape™++

Radware's AppShape technology transforms the ADC into a 'smart' device to accelerate, ease and optimize application deployment on the ADC. With Radware's AppShape, each ADC service is tailored to and *aware* of a specific business application (such as SAP, Microsoft, Oracle, IBM and more). In this way, the ADC can be managed from an application-oriented perspective via application specific configuration templates and wizards – resulting in fast application roll-out and simplified application management. Plus, AppShape offers logs and reports for: compliance, per application trends analysis and resources utilization.

Radware's also provides ADC policy scripting capabilities with its AppShape++ technology to further enable the customization of the ADC service per specific application flows and scenarios. By leveraging scripts, examples in Radware's library and dev-community, customers can easily use AppShape++ to refine various layer 4-7 policies including HTTP, HTTPS, TCP, UDP, SSL and more – with no application modifications to further reduce cost and risk.

Simply the Best-in-Class ADC Choice

The combination of these advantages – along with an industry unique 5-year longevity guarantee, "pay-as-you-grow" approach in throughput, # of vADCs and services, plus performance leadership in all layer 4-7 metrics – makes Alteon 5224 simply your best application delivery choice. Want to see for yourself? We invite you to download our Radware ADC Solution white paper here or contact us at: info@radware.com.

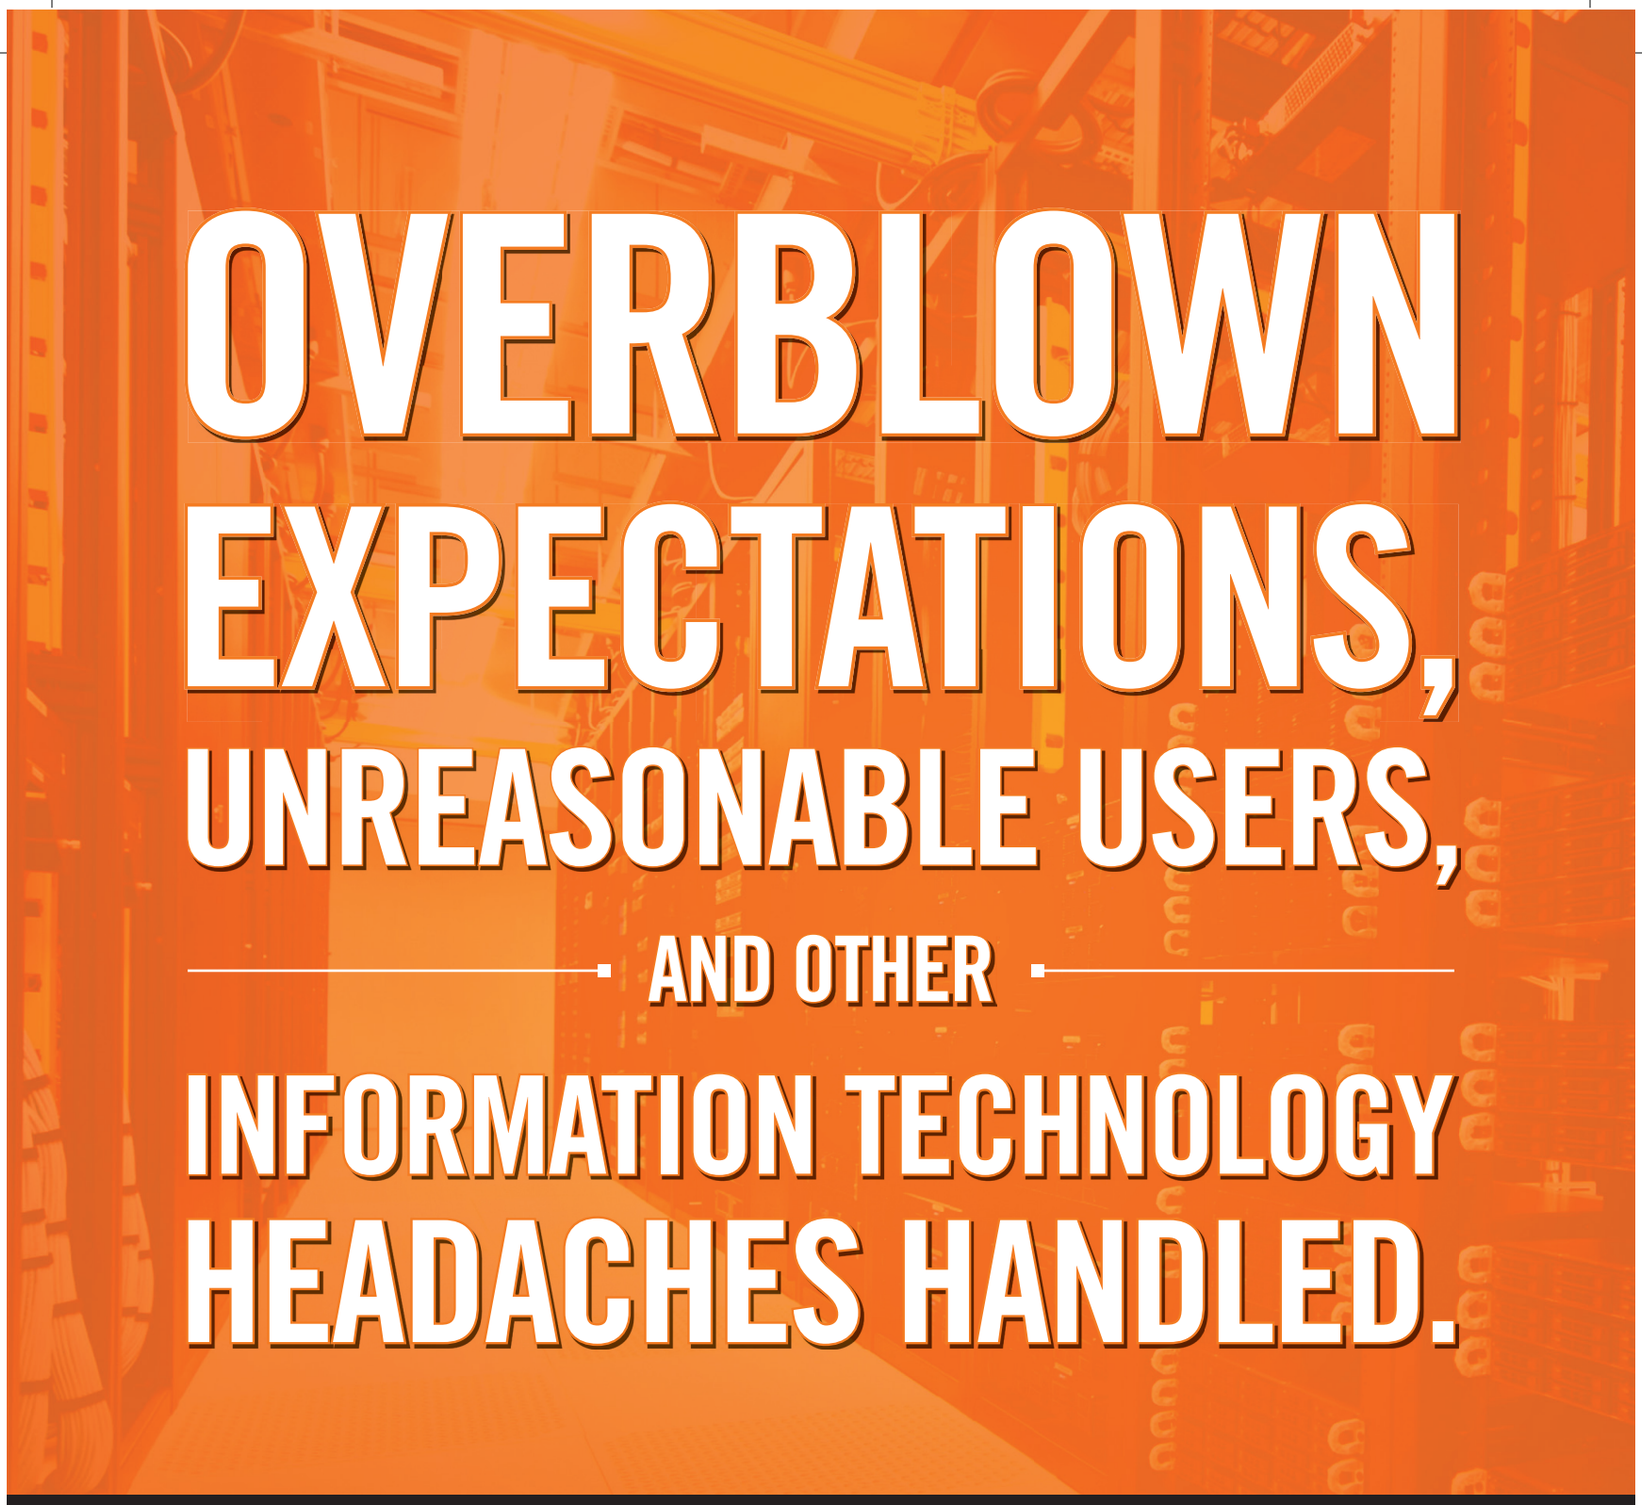

OVERBLOWN EXPECTATIONS, UNREASONABLE USERS, AND OTHER INFORMATION TECHNOLOGY HEADACHES HANDLED.

There can be no doubt that increasing user demands, computing trends and new technology are placing incredible expectations on IT professionals. Thanks to Riverbed Technology, you can now "right-place" your IT infrastructure in a way that meets growing application and data needs — without sacrificing performance or blowing your entire budget.

Get started now at **www.datacenter.riverbed.com**

riverbed®
Think fast.®

Silver Peak

Visit the Silver Peak Marketplace Today!

WANop Anywhere. Anytime.

Data acceleration has never been easier!

www.silver-peak.com/marketplace