

# The 2011 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:



agility  
made possible™



Produced by:



# Introduction

While ensuring acceptable application delivery has always been important, it historically was not a top of mind issue for the majority of IT organizations. That changed several years ago when IT organizations began to develop a concerted focus on it. Throughout the **2011 Application and Service Delivery Handbook**, this era that began several years ago will be referred to as the Application Delivery 1.0 era.

At the same time that many IT organizations are still in the process of implementing solutions that respond to the challenges of the Application Delivery 1.0 era, a new generation of challenges is emerging. These challenges are driven in large part by the:

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

Throughout this handbook, the emerging generation of application delivery challenges and solutions will be referred to as the Application Delivery 2.0 challenges and solutions.

As we enter the Application Delivery 2.0 era, leading edge IT organizations must focus on improving their ability to ensure acceptable application and service delivery. Throughout the 2011 Application and Service Delivery 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

***The goal of the 2011 Application and Service Delivery Handbook is to help IT organizations ensure acceptable application delivery when faced with both the first and the second generation of application delivery challenges.***

# Application and Service Delivery Challenges

The following challenges are associated with the Application Delivery 1.0 era:

- Limited focus on application performance during application development
- Chatty protocols and applications
- The Webification of applications
- Security vulnerabilities
- Server consolidation
- Data center consolidation and single hosting
- Distributed employees
- Distributed applications
- Complexity

All of the challenges listed above continue to impact IT organizations. In the Application Delivery 2.0 era, however, some of these challenges have either become notably more difficult or have morphed to become a different challenge. An example of a challenge that has become notably more difficult is the webification of applications. Current research shows that the number of hosts for a given web based user transaction varies around the world, but it typically is in the range of six to ten. An example of a challenge that has changed to where it has spawned a new challenge is that the task of supporting distributed employees has morphed to where a large percentage of those employees are now mobile.

Some of the other challenges that are associated with the Application Delivery 2.0 era include:

- Service Oriented Architectures (SOA) with Web Services
- Web 2.0 and Rich Internet Applications
- The increased focus on services
- Internal Service Level Agreements (SLAs)

Many of the challenges listed above have become top of mind issues for the majority of IT organizations. For example, recent market research indicates that two thirds of IT organizations believe that getting better at managing internal SLAs is either very or extremely important.

## Virtualization

Relative to the Application Delivery 2.0 era, virtualization is a double-edged sword as it both presents challenges and solutions. Of all of the myriad forms of virtualization, server virtualization receives the most attention. In early 2010, 20% of IT organizations had virtualized the majority of their data center servers. Today, 32% of IT organizations have virtualized the majority of their data centers servers. In addition, market research that was recently conducted predicts that within a year, that 40% of IT organizations will have virtualized the majority of their data center servers.

One of the challenges associated with server virtualization is supporting the dynamic movement of VMs between data centers. Recent market research indicates that fifty percent of IT organizations find that it is either very or extremely important for them to get better at supporting the dynamic movement of VMs. Some of the other challenges associated with server virtualization include:

- Contentious management of the vSwitch
- Breakdown of network design and management tools
- Limited visibility into VM-to-VM traffic
- Poor management scalability
- Multiple hypervisors
- Inconsistent network policy enforcement
- Over subscription of server resources
- Complex troubleshooting on a per-VM basis

At the present time, there is no overarching solution for the comprehensive management of a computing environment composed of virtualized servers, storage, and networks. Listed below are some the key developments that can help IT departments meet the challenges of virtualization.

- Dynamic infrastructure management
- Virtualized performance and fault management
- Distributed virtual switching
- Edge virtual bridges
- Orchestration and provisioning

Desktop virtualization doesn't receive as much attention as server virtualization does in part because today the majority of IT organizations haven't virtualized any of their desktops. However, twenty five percent of IT organizations that haven't yet virtualized any desktops plan to do so in limited fashion in the next year.

The two fundamental forms of desktop virtualization are:

- Server-side application/desktop virtualization
- Client-side application/desktop virtualization

With server-side virtualization, the client device plays the familiar role of a terminal accessing an application or desktop hosted on a central presentation server. There are two primary approaches to server-side application/desktop virtualization. They are:

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

Client-side application virtualization is based on a model in which applications are streamed on-demand from central servers to client devices. On the client-side, streamed applications are isolated from the rest of the client system by an abstraction layer inserted between the application and the local operating system.

One of the primary challenges that are associated with implementing desktop virtualization is achieving an acceptable user experience for client-to-server connections over a WAN. For example, VDI requires at least 200 Kbps of bandwidth per simultaneous user and the minimum peak bandwidth required for a PCoIP connection is one Mbps. In most cases, the successful deployment of desktop virtualization requires the broad deployment of WAN optimization techniques that focus on the particular characteristics of the various traffic streams that are associated with desktop virtualization. Current market research indicates that by the end of the year, the vast majority of virtualized desktops will be utilizing server side virtualization.

While not widely deployed currently, IT organizations are showing a significant interest in implementing virtualized appliances. This interest is driven in part by the fact that a virtual

appliance can help IT organizations respond to some of the challenges created by server and desktop virtualization. **A Virtual Appliance** is based on the appropriate software running in a VM. Virtual appliances can include WOCs, ADCs, firewalls, and performance monitoring solutions among others.

One of the compelling advantages of a virtualized appliance is that the acquisition cost of a software-based appliance can be notably less than the cost of a hardware-based appliance with same functionality. In addition, a software-based client can potentially leverage the functionality provided by the hypervisor's management system to provide a highly available system without having to pay for a second appliance.

One of the potential downsides of a virtual appliance is performance. The conventional wisdom in the IT industry is that a solution based on dedicated, purpose-built hardware performs better than a solution in which software is ported to a generic piece of hardware, particularly if that hardware is supporting multiple applications. However, conventional wisdom is often wrong and IT organizations that are considering deploying a virtual appliance need to test the performance of that appliance in their production environment.

A critical factor that must be considered when evaluating the deployment of virtual appliances in a dynamic, on-demand fashion is the degree of integration that the virtual appliance has with the virtual server management system. Ideally this management system would recognize the virtual appliances as another type of VM and understand the associations between the appliance VM and the application VMs in order to allow a coordinated migration whenever this is desirable.

## Cloud Computing

Cloud computing is based on a number of familiar concepts including time-sharing, automation, virtualization and the rental of applications. What is new about cloud computing is the synthesis of these concepts combined with the dynamic creation and movement of IT resources.

The goal of cloud computing is to enable IT organizations to achieve a dramatic improvement in the cost effective, elastic provisioning of IT services that are good enough. The phrase **good enough** refers in part to the fact that the SLAs that are associated with public cloud computing services such as Salesforce.com or Amazon's Simple Storage System are generally very weak.

Some of the primary characteristics of a cloud computing solution are:

- The centralization of applications, servers and storage resources
- The extensive virtualization of every component of IT
- The standardization of the IT infrastructure.
- The simplification of the applications and services provided by IT
- Technology convergence
- Service orchestration
- Automation
- Self-service
- Usage sensitive chargeback
- The dynamic movement of resources

There are three primary classes of cloud computing solutions. They are:

- Public
- Private
- Hybrid

Most public cloud based solutions are delivered over the Internet and no vendor will provide an end-to-end performance SLA for a service delivered over the Internet. As a result, many of the approaches to providing public cloud-based solutions will not be acceptable for the applications, nor for the infrastructure that supports the applications, for which enterprise IT organizations need to provide an SLA.

One way that an IT organization can improve the performance and availability of cloud computing solutions is to implement cloud balancing. Cloud balancing refers to routing service requests across multiple data centers based on myriad criteria. One way to think about cloud balancing is that it is the logical extension of global server load balancing (GSLB).

As is true with any new technology or way to deliver technology based services, there are risks associated with the adoption of all three classes of cloud computing. While the security risks get the most attention, cloud computing also presents significant management and performance challenges. However, the biggest risk accrues to those companies that don't implement any form of cloud computing.

## Optimizing and Securing the Use of the Internet

The applications that are typically identified with public cloud computing are well known enterprise applications including CRM, SCM and ERP. While those applications will continue to be closely associated with cloud computing, it is becoming increasingly common for organizations to acquire a different category of application from a cloud computing service provider (CCSP). That class of applications is traditional network and infrastructure services such as VoIP, unified communications, management, optimization and security. Such applications will be referred to as a Cloud Networking Service (CNS) and recent market research indicates that over the next year that many IT organizations intend to make significant use of CNSs

The traditional application delivery solutions based on WAN optimization controllers (WOCs) and application delivery controllers (ADCs) were designed to address application performance issues at both the client and server endpoints. These solutions make the assumption that performance characteristics within the WAN itself can't be optimized. This assumption is reasonable in the case of WAN services such as ATM or 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 based on the use of a CNS. Such a CNS would have to 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.



## Planning

Many planning functions are critical to the success of application delivery. One such function is identifying the company's key applications and services and establishing SLAs for them. Another key planning activity that will be elaborated on in the following chapter is Application Performance Engineering (APE). One of the characteristics of APE is that it is a life cycle approach to planning and managing application performance. Addressing performance issues throughout the application lifecycle is greatly simplified if there are tight linkages between the IT personnel responsible for the planning and operational functions.

For those organizations that run a large, complex network there often is a significant gap between network planning and network operations. A class of management tool that can facilitate the integration of planning and operations is typified by an IP route analytics solution. The goal of route analytics is to provide visibility, analysis and diagnosis of the issues that occur at the routing layer in complex, meshed networks.

Route analytics is gaining in popularity because the only alternative for resolving logical issues involves a very time-consuming investigation of the configuration and log files of numerous individual devices. Route analytics is also valuable because it can be used to eliminate problems stemming from human errors in a router's configuration by allowing the effect of a configuration change to be previewed before the change is actually implemented.

Most IT organizations that have already implemented either public or private cloud computing have not done so in a highly systematic fashion. In order to maximize the benefit of cloud computing, IT organizations need to develop a plan (The Cloud Computing Plan) that they update on a regular basis. The Cloud Computing Plan should identify the opportunities and risks associated with both public and private cloud computing. The Cloud Computing Plan must identify a roadmap of what steps the IT organization will take on a quarter-by-quarter basis for the next two to three years and ensure that the steps are in line with the corporate culture.

The Cloud Computing Plan should look systematically across multiple technologies because of the interconnected nature of the technologies. As part of creating this plan, IT organizations need to understand the cloud computing strategy of their existing and potential suppliers, including the partnerships that the suppliers are establishing between and amongst themselves.

## Application Performance Management

Successful APM requires a holistic approach based on integrated management of both the application and/or service itself as well as the end-to-end IT infrastructure. However, only a small percentage of IT organizations take such an approach.

A holistic approach to APM must focus on the experience of the end user of the application or service. Monitoring actual user transactions in production environments provides valuable insight into the end-user experience and provides the basis for an IT organization to be able to quickly identify, prioritize, triage and resolve problems that can affect business processes.

A holistic approach to APM must also address the following aspects of management:

- The adoption of a system of service level agreements (SLAs) at levels that ensure effective business processes and user satisfaction for at least a handful of key applications.
- Automatic discovery of all the elements in the IT infrastructure that support each service. This functionality provides the basis for an IT organization to be 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.

Some of the challenges that make APM more difficult include:

- Port hopping
- Instant messaging
- Peer-to-peer networks
- The port 80 black hole
- Server virtualization
- Mobility
- Cloud computing

A concept that is closely related to APM is APE. APE is the practice of first designing for acceptable application performance and then testing, measuring and tuning performance throughout the application lifecycle. During the operational, or production phase of the lifecycle, APM is used to monitor, diagnose, and report on application performance. APM and APE are therefore highly complementary disciplines. For example, once an APM solution has identified that an application in production is experiencing systemic performance problems, an APE solution can be used to identify the root cause of the problem and to evaluate alternative solutions.

The key components of APE are:

- Setting Performance Objectives
- Discovery of the Network Topology
- Performance Modeling
- Performance Testing and Analysis
- Optimization

Another concept that is closely related to APM is route analytics. As noted, the goal of route analytics is to provide visibility, analysis and diagnosis of the issues that occur at the routing layer in complex, meshed networks. These issues are often referred to as logical issues in contrast to a physical issue such as an outage. One of the reasons why route analytics is a key component of APM is that recent market research has shown that in the vast majority of cases, logical factors cause as much or more business disruption than do physical factors. That same research also showed that in the vast majority of instances, logical errors take either somewhat more or notably more time to troubleshoot and repair than do physical errors.



# Network and Application Optimization

The phrase **network and application optimization** refers to an extensive set of techniques the goal of which is to optimize the performance of networks and applications as part of assuring acceptable application performance. The primary role that these techniques play is to:

- Reduce the amount of data sent over the WAN;
- Ensure that the WAN link is never idle if there is data to send;
- Reduce the number of round trips (a.k.a., transport layer or application turns) necessary for a given transaction;
- Overcome the packet delivery issues that are common in shared (i.e., over-subscribed) networks;
- Mitigate the inefficiencies of certain protocols and/or applications;
- Offload computationally intensive tasks from client systems and servers;

There are two principal categories of network and application optimization products. One category focuses on mitigating the negative effect that WAN services such as MPLS have on application and service performance. This category of products has historically included WAN optimization controllers (WOCs). However, due to some of the second generation of application and service delivery challenges, this category of products now also contains an emerging class of WAN optimization device - the Data Mobility Controller (DMC). As described in detail later in this section of the handbook, WOCs are focused primarily on accelerating end user traffic between remote branch offices and central data centers. In contrast, DMCs are focused on accelerating the movement of bulk data between data centers. This includes virtual machine (VM) migrations, storage replication, access to remote storage or cloud storage, and large file transfers. WOCs and DMCs are often referred to as **symmetric solutions** because they typically require complementary functionality at both ends of the connection. However, as is explained later in this section of the handbook, one way that IT organizations can accelerate access to a public cloud computing solution is to deploy WOCs in branch offices. The WOCs accelerate access by caching the content that a user obtains from the public cloud solution and making that content available to other users in the branch office. Since in this example there is not a WOC at the CCSP's site, this is an example of a case in which a WOC is an asymmetric solution. Roughly fifty percent of IT organizations have already deployed WOCs, although relatively few IT organizations have deployed them broadly within the organization. Over the next year, IT organizations plan to make a moderate increase in their deployment of WOCs.

The second category of optimization products is often referred to as an Application Delivery Controller (ADC). This solution is typically referred to as being an **asymmetric solution** because an appliance is only required in the data center and not on the remote end. The genesis of this category of solution dates back to the IBM mainframe-computing model of the late 1960s and early 1970s. Part of that computing model was to have a Front End Processor (FEP) reside in front of the IBM mainframe. The primary role of the FEP was to free up processing power on the general purpose mainframe computer by performing communications processing tasks, such as terminating the 9600 baud multi-point private lines, in a device that was designed specifically for these tasks. The role of the ADC is somewhat similar to that of the FEP in that it performs computationally intensive tasks, such as the processing of Secure Sockets Layer (SSL) traffic, hence freeing up server resources. However, another role of the ADC that the FEP did not provide is that of Server Load Balancer (SLB) which, as the name implies, balances traffic over multiple servers.

## 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](mailto:jim@webtorials.com) or Steven Taylor at [taylor@webtorials.com](mailto:taylor@webtorials.com).

**Published by  
Webtorials  
Editorial/Analyst  
Division**  
[www.Webtorials.com](http://www.Webtorials.com)

**Division Cofounders:**

Jim Metzler  
[jim@webtorials.com](mailto:jim@webtorials.com)  
Steven Taylor  
[taylor@webtorials.com](mailto:taylor@webtorials.com)

**Professional Opinions Disclaimer**

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

**Copyright © 2011 Webtorials**

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



# The Fastest Growing Application Networking Company



## 64-bit AX Series

### Application Delivery

- Advanced Application Delivery Controller (ADC)
- New Generation Server Load Balancer (SLB)

### IPv6 Migration

- Large Scale NAT
- Dual-Stack Lite
- NAT64 & DNS64
- IPv6 ↔ IPv4 (SLB-PT)

### Cloud Computing & Virtualization

- SoftAX & AX-V
- AX Virtual Chassis
- AX Virtualization (Application Delivery Partitions)

**Advanced Core Operating System (ACOS)**

### AX Series Advantage

- All inclusive pricing for hardware appliances, no performance or feature licenses
- Most scalable appliances in the market with unique modern 64-bit ACOS, solid-state drives (SSD) and multiple hardware acceleration ASICs
- Faster application inspection with aFlex TCL rules
- aXAPI for custom management

### Application Solutions

The AX Series increases scalability, availability and security for enterprise applications. Visit A10's web site for deployment guides, customer usage scenarios and to participate in the Application Delivery Community.



**Microsoft**



# Transforming the Internet into a Business-Ready Application Delivery Platform



## Ensuring applications perform to support your business goals

As organizations expand globally, they need to make a variety of business-critical applications available to employees, partners and customers across the globe. Application delivery strategies are increasingly leveraging Cloud based options for hosting enterprise applications on Cloud infrastructure and outsourcing applications via SaaS vendors. Organizations must also be sensitive to the economic pressures driving IT consolidation and centralization initiatives.

Whether delivering applications from behind the firewall, hosting in the cloud, or using a hybrid model, the Internet remains an integral part of application delivery strategy. Though global delivery of enterprise applications over the Internet can provide remote users with essential business capabilities, poor application performance can quickly sour user experience. Business applications must perform quickly, securely, and reliably at all times, or adoption and intended benefits will suffer.

## Key Challenges in Delivering Applications

IT organizations often use the public Internet to support globalization efforts because of its lower cost, quick time to deploy, and expansive reach. However, when delivering applications via the Internet to global users, business can face many challenges, including:

- Poor performance due to high latency and chatty protocols (like HTTP & XML)
- Spotty application availability caused by unplanned internet disruptions
- Inadequate application scalability and spiky peak usage
- Growing security threats, including distributed denial of service, cross-site scripting, and SQL injections

These problems can severely undermine application effectiveness and ROI and do not disappear by moving to the Cloud.

## Akamai's Application Performance Solutions

Today, thousands of businesses trust Akamai to distribute and accelerate their content, applications, and business processes. Akamai Application Performance Solutions are a portfolio of fully managed services designed to accelerate performance and improve reliability of any application delivered over the Internet, hosted behind the firewall or in the Cloud, with no significant IT infrastructure investment.

Akamai leverages a highly distributed intelligent Internet platform, comprised of tens of thousands of servers, within a single network hop of 90% of the world's Internet users. The Akamai Protocol optimizes application delivery at the routing, transport, and application layers, not only caching content at the Internet's edge, close to end users, for fast delivery, but accelerating dynamic content from the origin to global users. This intelligent Internet platform also extends the security perimeter to the edge of the Internet with modules providing a cloud based Web Application Firewall and DDoS defense.

Application Performance Solutions drive greater adoption through improved performance, higher availability, and an enhanced user experience, ensuring consistent application performance, regardless of user location, and delivers capacity on demand, where and when it's needed. This helps reduce infrastructure costs and support data center consolidation. Examples of applications delivered by Application Performance Solutions include Web-based enterprise applications, Software as a Service (SaaS), applications deployed on IaaS and PaaS, Web services, client/server or virtualized applications, live chat, productivity, and administration functions, such as secure file transfers.

To learn more about Akamai Application Performance Solutions, visit [www.akamai.com/aps](http://www.akamai.com/aps).



# Next-Generation Application Acceleration



Organizations everywhere face tough challenges in optimizing business application performance. For today's distributed enterprises, centralization and server consolidation can create user response and network capacity problems; business applications are often slow or unpredictable; and bandwidth costs are out of control. Now, IT is expected to deliver even more — including corporate communication videos and cloud delivered software-as-a-service (SaaS) applications — all while containing costs.

To solve these and other application delivery problems, you have to understand how application performance requirements have changed, know which technologies can meet your business demands today and prepare for capacity needs down the road.

## The Foundation: Optimizing Traditional Applications

Rapid growth of files, email, storage and backup systems put an incredible burden on WAN connections and create significant end-user performance issues — unless you can accelerate traffic. Blue Coat's protocol optimization, byte caching, compression and QoS are the technologies required to accelerate remote and branch office access to centralized files, email and backup systems. These technologies offer significant performance benefits by mitigating the latency caused by chatty file protocols, caching data and expanding bandwidth for high-volume transfers. Besides data applications, however, you need specialized technologies to optimize performance of key emerging applications.

## Next Generation WAN Optimization Requirements

Many of the latest applications are changing the way we collaborate, educate, and communicate. Video, for instance, is increasingly used for training and live communications, and Cloud delivered SaaS applications are enabling new business processes. However, the traditional acceleration technologies cannot address these newer types of applications.

### Streaming video and rich media

Delivering high-quality, on-demand or live streaming video requires massive amounts of bandwidth on specialized protocols. For example, a single live stream can be 200KB to 1.5MB and large on-demand files can reach 25MB, 100MB and even 1GB in size. In addition, bandwidth-hungry rich media applications can dominate the entire network and still fail due to insufficient resources.

### Cloud Delivered SaaS applications

SaaS applications, such as Salesforce.com, or SaaS-hosted SAP and SharePoint applications have unique management challenges due to their location and the encryption used to secure them. Because SaaS offerings are located outside of your network they are outside of your control, but still need to be accelerated. They are also encrypted with SSL and use certificates and keys controlled by the SaaS provider and the Web browser — not your organization.

Traditional WAN Optimization technologies would require you to place an appliance on the SaaS provider's network, which is simply not possible. Because SaaS applications rely on HTTP and SSL delivery, you need optimization technologies that can asymmetrically accelerate HTTP and SSL, as well as secure client-side certificate handling so you can decrypt and accelerate the sessions.

### Next-Generation Acceleration

The good news is next generation acceleration technologies available today can help you optimize your most critical applications and reclaim bandwidth from non-essential traffic. These new optimization technologies include:

- Video caching, stream splitting and Content Delivery Network (CDN) to enable optimized delivery of business video and minimize the impact of recreational video over the WAN.
- Asymmetric optimizations technologies and external SSL certificate handling that don't require changes to the SaaS infrastructure, like Blue Coat CloudCaching engine.
- URL classification and content filtering with usage and QoS policies to identify and contain recreational content and traffic.
- Integration with web security service to protect Internet-connected branch offices from malware and enable faster SaaS, 100% recreational offload and high availability networking.

**Figure 1: Performance gains by technology type**



#### Video Optimization

- Scale internal Video 10x – 100x – 1000x
- Reduce Recreational Video by 30-80% across the distributed enterprise



#### Cloud Acceleration

- Accelerate SaaS applications directly to branch offices by 7 – 93x
- Eliminate back-hauling SaaS/Internet applications over WAN



#### Traditional WAN Optimization

- Accelerate applications by 3x-300x from data center to branch office
- Reduce storage replication and backup bandwidth by up to 90%

### Get the right acceleration strategy

Acceleration requirements have rapidly moved beyond CIFS and MAPI acceleration. Video and SaaS application delivery are IT's challenges today. With the right acceleration strategy, you can gain superior business value from internal and external infrastructure. Find out how Blue Coat can help at [www.bluecoat.com](http://www.bluecoat.com)

### About Blue Coat

Blue Coat Systems secures and optimizes the flow of information to any user, on any network with leading web security and WAN Optimization solutions. Blue Coat enables the enterprise to tightly align network investments with business objectives, speed decision making and secure business applications for a long-term competitive advantage.



## Application Performance: Your Window to Service Delivery

Virtually all organizations depend on online services to transact business. For online brokerage, retail companies and others, online services are their business. For insurance companies and manufacturers, online services enable their business. Regardless of your business, you want to deliver a positive customer experience. Satisfied customers come back and customer retention is the foundation of your bottom line. Pressure is mounting on IT departments to deliver on this requirement and, as a result, continually increasing amounts of IT budgets are spent on tools and processes to assure that services are performing.

Keeping customers happy is nothing new. For years, organizations focused on the domains – the network, the databases, the servers – assumed that if all the domains worked so would the service. But that strategy exposed an IT paradox: IT services are more than the sum of their parts. Managing each domain for peak performance is no guarantee of success. The information essential to assuring services include the service delivery pathway - the route through the infrastructure the service takes to reach the customer - and the components in that pathway - the network links, databases and servers that are essential to delivering the service. Tools and teams, dedicated to supporting individual domains, often have a hazy view of which components actually impact specific services. For example, a single server outage may have nothing to do with your critical business service...or it may have everything to do with it. Managing each domain for peak performance without a clear asset-to-service view is not a guarantee that you will stay ahead of calls to the help desk.

### Applications as Bellwethers

Are applications another domain or a service? Actually, it depends. Some applications, such as online trading, are the end-user service. An email application is often an end-user service but can also be an enabling part of an online retail service, thereby putting the application into the role of a domain in a service delivery pathway. What can be said with certainty is that applications are an essential part of any service, and applications, like services, rely on the other domains to function. Consider how Web-based applications rely on the full range of IT infrastructure components to be operational. So, whether your application is the service itself or a service enabler, its performance is linked tightly with your business service delivery and that makes your application performance an open window into your service performance.

The CA Technologies [Service Assurance](#) portfolio is built around proactive performance management. On a foundation of [CA eHealth® Performance Manager](#) for client-server applications, CA Technologies added and integrated CA Introscope and CA Customer Experience Manager, the [CA Application Performance Management](#) (CA APM) solution, to detect, triage and diagnose performance problems in your complex, composite and Web application environments. CA APM supports both Java and .NET applications and provides end-to-end visibility to online transactions. To complete the picture, CA Technologies acquired NetQoS, bringing products like [CA NetQoS SuperAgent®](#) and [CA NetQoS ReporterAnalyzer™](#) into the fold. CA NetQoS SuperAgent tracks every TCP application packet traversing the network between clients and servers, providing metrics such as network, server and application latency for all applications. CA NetQoS ReporterAnalyzer provides historical, real-time and predictive behavioral views through traffic composition metrics that show how applications tax and compete for network resources. With these detailed application performance metrics, application delivery bottlenecks are quickly





pinpointed, root cause established and performance issues corrected, often before user impact. And the ripple effect on business services is all positive.

### **Service Assurance: Application Performance Plus**

The CA Technologies Service Assurance portfolio provides a layer of intelligence that leverages data from your existing infrastructure and application performance management tools used to directly manage your IT assets, including [Infrastructure Management](#) products like CA Spectrum® Infrastructure Manager, CA eHealth Performance Manager and CA NetQoS Performance Center, and CA [Application Performance Management](#) products like CA Introscope® and CA Customer Experience Manager. Consolidating information from these performance managers, [CA Service Operations Insight \(formerly CA Spectrum® Service Assurance\)](#) provides the business service analytics, uniquely linking applications to infrastructure to calculate key performance indicators (KPIs) for service delivery and risk.

CA Spectrum Service Assurance creates a single service model, leveraging information from the domain managers, that is updated dynamically as things change, so you know what components – infrastructure or application – are in the pathway of your critical business service and you know if there is a problem that will impact service delivery, even as configurations and virtual machines change. With the CA Technologies Service Assurance portfolio, you can prioritize your efforts, have confidence in the information you have and fix the important things first to minimize customer and business impact. Even better, CA Technologies can show you where a potential problem is chipping away at performance, for example, telling you when a server farm is losing machine power even if it is not yet impacting service. This puts you where you want to be - two steps ahead of your customer.

### **Integration Works at Rooms To Go**

Customers that have benefitted from the tight integration in the CA Technologies Service Assurance portfolio have compelling stories to tell. Putting it all together was the key for Rooms To Go. To enhance the customer experience at its 150 showrooms across the U.S., [Rooms To Go](#) added CA Technologies software for network, application and virtual system performance management to its existing Service Assurance products to maintain service availability and improve support of its retail and distribution outlets.

Rooms To Go is using the [CA NetQoS Performance Center](#), a key component of the Service Assurance portfolio, and [CA Virtual Assurance for Infrastructure Managers](#) to improve the performance of its most business-critical, networked applications and their supporting infrastructure. For example, Rooms To Go uses the two CA Technologies solutions to monitor and manage its point-of-sale (POS) application that provides immediate purchase-related information and fast credit application processing and approvals.

“The CA NetQoS Performance Center and CA Virtual Assurance for Infrastructure Managers will help Rooms To Go be more proactive in ensuring a high level of service across our stores and improving the customer experience as a result,” said Jason Hall, Director of IT systems for Rooms To Go. “Combined with our other products from CA Technologies, the CA NetQoS and CA Virtualization Management solutions will give us a more complete understanding of what is happening across our network and virtualized infrastructure and where we need to direct our attention to solve problems faster, prepare for future capacity needs, and optimize application performance.”

In addition to monitoring how well the network delivers the POS application to the Rooms To Go showrooms, the CA NetQoS solution will help Rooms To Go understand how application traffic affects network performance, with views into the composition of traffic on every network link, and which applications and users consume bandwidth. Before installing NetQoS, Hall had no visibility into how end users were experiencing application and service performance across the WAN or LAN. “It was purely the end user,” he said. “We waited for someone to call. Operationally, that gives the end user the perspective that the systems are slow ... and that we’re not doing



anything about it.” Hall said that adding NetQoS's performance management capabilities to his suite of tools has also helped him solve some service delivery mysteries, particularly with his company's intranet. You can read more on this story on SearchNetworking.com in their June 16, 2010 article by Shamus McGillicuddy titled, [“Service delivery management: Integrating IT management tools.”](#)

### **Jack Henry & Associates Put Service First**

No one doubts the importance of accuracy and high performance when it comes to financial applications. [Jack Henry & Associates](#) processes transactions, automates business processes, and manages mission-critical information for more than 8,700 financial institutions and corporate entities, serving around six million end-users who depend on Jack Henry to run business-critical applications and financial processes. Initially, the company had no consistent means of monitoring end-to-end performance across its network and applications, which made it difficult to safeguard service levels and manage capacity.

“We have to prove every single day that our performance is meeting customer requirements, which, without end-to-end monitoring, was challenging,” said Josh Bovee, Senior Network Engineer, Jack Henry & Associates. “We needed to focus on application performance from the end-user perspective and create a baseline of how well we were serving those customers so we could understand when performance degraded and what impact things like infrastructure changes might have. We were reliant on getting all the IT groups in the same room, and then putting our heads together until we located the source of the issue. With limited insight into network and application performance metrics, this would often take days.”

Realizing they needed to take a more proactive approach to managing its business critical banking applications, Jack Henry looked for a solution that would address its performance management challenges. After struggling for several months with a competitive product, they arranged with CA Technologies for a Proof of Concept with the NetQoS Performance Center, starting with the CA NetQoS SuperAgent. “We started the POC at 8 a.m. and by 1 p.m. we were capturing more meaningful data with SuperAgent than after six months working with the competitive product. SuperAgent was also easier to implement. We didn’t need to install an agent on the server or re-architect our infrastructure, which was something we very much wanted to avoid,” notes Bovee. Having made the decision to deploy CA NetQoS SuperAgent, the company decided to implement additional modules of the CA NetQoS Performance Center.

Jack Henry now has a finger firmly on the pulse of its customers’ business-critical applications, furthering its commitment to industry-leading client satisfaction and retention rates. As a result of their investment in CA Service Assurance solutions, the company is already benefiting from improved service, more cost-effective support and greater business agility. “We now have a great foundation on which to continue to improve our service levels and customer satisfaction,” concludes Bovee.

### **CA Technologies Manages Risk to Assure Application and Service Delivery**

Service Assurance and risk management is achieved through new, advanced technology that can model the IT assets that comprise services, track service quality (end-user experience), the status of each IT asset (network devices, systems, databases and applications) and calculate each asset’s risk to each service dynamically. With this information, you’ll know how to proactively fix problems before they impact users.

These capabilities also factor dimensions of risk beyond typical KPIs to include compliance, answering questions such as: “Are my business services at risk because configurations do not meet the gold standard? Do we have the latest security patches deployed on every device?”

Identifying and measuring risk to business services benefits both IT executives and the technical staff who manage the IT environment “hands-on.” By understanding risk, IT executives can make more informed decisions about



capital and operational investments. Technical staff benefit because they can see the root cause of trends that will impact services in the future and can proactively prevent impact to quality.

CA Technologies Service Assurance is a mature, integrated portfolio that provides end-to-end visibility into business services, applications and transactions linked with top-to-bottom insight over the entire infrastructure. Providing great service in a consistent manner, meeting SLAs and having the agility in your infrastructure to roll out new services quickly and efficiently is just table stakes in today's complex IT environment. No matter what business you are in, service assurance is critical to your success, and CA Technologies can work with you to help you deliver the service your customers demand.

*Advertorial*



# Software WAN Optimization



“... application performance ... one of the top three inhibitors of cloud adoption” <sup>1</sup>

## aCelera™

Secure Automated Optimized

- Any deployment model:
  - Enterprise
  - Hosted
  - Cloud
  - Or - Any combination
- Any hypervisor & Windows Server 2008 R2
- Any number of instances
- Any throughput capacity
- Any security requirement
- Any routing mode
- Any Failover mode
- Automated Management
- Meet cost savings objectives
- Match footprint limitations
- Bundle best of breed applications:
  - Video streaming
  - Directory services
  - Security



Certeon Inc.  
4 Van de Graaff Drive  
Burlington, MA 01803  
781 425 5200

<http://www.certeon.com>

## APPLICATION DELIVERY PERFORMANCE

### From datacenter and cloud to any user anywhere

#### Maximum Value & Maximum Performance

The **business value** of any application must be measured by its ability to increase business agility, decrease cost through on-demand provisioning and teardown of infrastructure and services, accelerated development, and improved reliability. Solutions must be utility-based, self-service, secure and most importantly, have levels of application performance that improve productivity.

Maximizing the business value of any networked application requires full featured, secure, scalable, high performance WAN Optimization software that allows applications to perform as expected, and can be part of any on demand architecture. Tactical hardware or virtual appliances with limited performance don't measure up.

#### aCelera: Built for Global Performance

aCelera software exceeds the scalability and performance of purpose-built hardware appliances. aCelera WAN Optimization software can support hundreds of thousands of connections and gigabits of throughput. It is built to support global enterprise scalability requirements and is ready for the Internet scale demands of managed services and cloud computing providers.

aCelera software and virtual appliances deliver these performance benefits and advantages without the costs or the friction of hardware appliances or limited scope virtualization. aCelera can easily be scaled on any existing hardware platform or migrated to more powerful platforms and processors when business conditions dictate, leveraging any industry standard management tool.

#### aCelera: Built for Global Deployment

Enterprise and cloud infrastructures are not uniform. aCelera software can be deployed in any heterogeneous mix of hardware, virtualization platforms, storage technologies, networking equipment and service providers supporting any custom or off the shelf application.

Hardware WAN optimization products require more planning and are more labor intensive to install. aCelera software packages are delivered over a network and installed in a data centers, remote sites, or end user PCs in less than 30 minutes. aCelera creates a high performance WAN infrastructure that can span the globe and scale to meet your application and user performance needs.

aCelera can be deployed in any private, public, and hybrid cloud computing environment and is poised to meet ANY future performance, scale and connection demands imposed by any enterprise IT environment, private network, private cloud, public cloud or a hybrid of them all.

**aCelera software WAN optimization:**  
**60% better 3 year TCO & 50% better scalability**


<sup>1</sup>, Clouds and Beyond: Positioning for the Next 20 Years of Enterprise IT, Frank Gens, IDC

# Cisco Unified Network Services



Highly virtualized data center and cloud environments impose enormous complexity on the deployment and management of network services. Provisioning dynamic services and accommodating mobile workloads present challenges for layered services, such as security and application controllers, that traditionally have required in-line deployment and static network topologies. Cisco® Unified Network Services meets these challenges with integrated application delivery and security solutions for highly scalable, virtualized data center and cloud environments.

**Any Service:** Cisco Unified Network Services is a critical component of the Cisco Data Center Business Advantage architecture. It consists of Cisco Application Control Engine (ACE) application controllers, Cisco Wide Area Application Services (WAAS) WAN acceleration products, Cisco Adaptive Security Appliances (ASA) data center security solutions, Cisco Virtual Security Gateway (VSG), Cisco Network



Analysis Module (NAM), and associated management and orchestration solutions.

**Any Form Factor:** Cisco Unified Network Services provides consistency across physical and virtual services for greater scalability and flexibility. One element of the Cisco Unified Network Services approach is the concept of a virtual service node (VSN), a virtual form factor of a network service running in a virtual machine. Cisco VSG for

Cisco Nexus® 1000V Series Switches and Cisco Virtual WAAS (vWAAS) are examples of VSNs that enable service policy creation and management for individual virtual machines and individual applications.

**Outstanding Scalability:** In addition to virtualization-aware policies and services, Cisco Unified Network Services supports greater data center scalability and cloud deployments, with the services themselves being virtualized. The application and security services can be provisioned and scaled on demand and can be easily configured to support the needs of dynamically deployed and scalable virtual applications.

**Integrated Management Model:** Cisco Unified Network Services enables consistency of management across different services and across physical and virtual form factors. Cisco Unified Network Services is thus a critical component of a fabric-centered data center architecture that is well integrated with the virtual servers and applications to readily enable scalable public and private cloud environments.

## Application Delivery Controllers


Enhanced web application performance, availability, and server scalability



Cisco ACE module and appliance, Cisco GSS

## WAN Optimization

Reduce branch IT costs and enhanced application performance for the distributed enterprise



Cisco WAAS appliances and modules

Cisco vWAAS

## Network Analysis and Monitoring

Simplifies application performance monitoring



Cisco NAM appliances, modules, and virtual blades

## Data Center Security

Physical and virtual solutions remove multi-tenant security risks and external threats



Cisco ASA 5585-x

Cisco VSG



### 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](http://www.cisco.com/go/offices).

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at [www.cisco.com/go/trademarks](http://www.cisco.com/go/trademarks). Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)



---

# UNIFIED PERFORMANCE MANAGEMENT

---

VISIBILITY | CONTROL | OPTIMIZATION

## COMPLETE WAN OPTIMIZATION

Increase the speed and efficiency of your wide area network.

Exinda's Unified Performance Management (UPM) solution delivers everything you need to manage your application performance and ensure the highest quality user experience.

Point solutions lack inter-communication between the functions of visibility, control, and optimization. This creates contention between these independent solutions, as each function is unaware of the effect its actions has on the other.

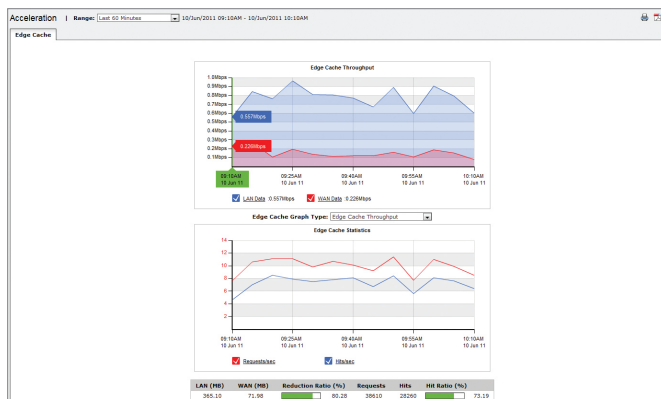
Exinda's unique, holistic approach to WAN Optimization eliminates the communication barriers and contention of point solutions, by integrating visibility, control and optimization, into a single, unified solution.



## LATEST ADVANCES IN UNIFIED PERFORMANCE MANAGEMENT

Exinda's development team is continually adding new features and functionality into our unified performance management solution. It is because of our agile development cycle and constant push to add innovation to our product line that Exinda has become the fastest growing WAN optimization vendor in the world. The following are some of the latest advances in our UPM solution.

### EDGE CACHE



Exinda Edge Cache will allow you to reduce bandwidth usage, decrease network costs, and accelerate content delivery, improving user experience and productivity.

### Edge Cache

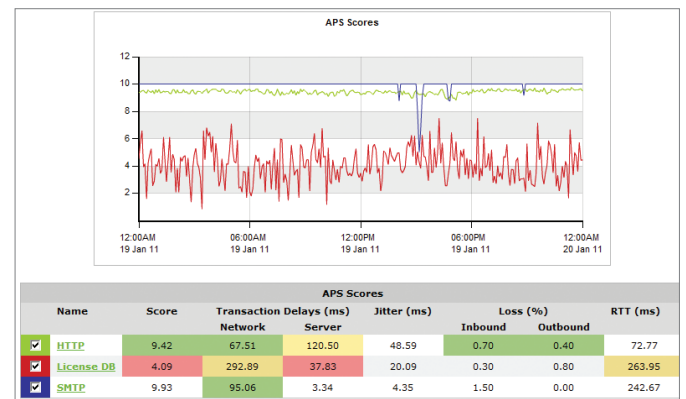
The Exinda Edge Cache™ enables single-sided caching of Internet-based content at the network edge, including web objects, videos and software updates, delivering a superior user experience and reducing WAN resource utilization.

Web objects are cached at the network edge when they are first downloaded from the Internet or across WAN links. These objects can then be delivered to the users on subsequent requests over the corporate local area network much faster without needing to download the data over the WAN again, providing a better user experience and increased productivity to the workforce. By caching web objects in the local office, organizations can drive down the network traffic consumed by each office, which directly reduces network costs.

The Exinda Edge Cache enables caching of web objects, video, software update and other content on the WAN. It also offers cache statistics, which provide insight into the amount of repetitive data being off-loaded from the WAN link, how cacheable the network data is, how frequently the cache is being accessed, and by how many hosts, helping organizations to understand the nature of their network traffic over time.

The Exinda Edge Cache can also be aligned with an organization's optimization policies, allowing the administrator to only cache specific content for specific users or groups of users, and to maintain very precise controls over how much WAN bandwidth should be made available for each application traversing the network.

### APPLICATION PERFORMANCE SCORE



Gain proactive reports on users perception of application performance & responsiveness.

### Application Performance Score

A significant feature of Exinda's WAN Optimization solutions is its ability to provide Application Performance Scores (APS). Exinda's APS provides a single data point to monitor and report on the overall health and performance of an application on your network. With APS, you can set performance thresholds for the applications on your network, and easily monitor if and when the thresholds are met or exceeded. When WAN application performance issues arise, the APS allows you to quickly troubleshoot the problems, by drilling down into individual metrics for the application, including network delay, server delay, jitter and loss, and round trip time, helping you to pinpoint and address the source of the performance issue.

Exinda also allows you to monitor and report on TCP efficiency and health. With Exinda, TCP efficiency reports let you examine how efficiently packets flow through the network, based on the number of dropped packets and retransmitted packets for the application. When combined with Exinda's TCP health monitoring, TCP efficiency reporting gives you a more in-depth view of network and application performance. TCP Health monitoring displays the health of TCP Connections by showing the total number of TCP connections, and how many were aborted, ignored, or refused by the server. With Exinda, you get a simple graphical view of the TCP health of the network, allowing rapid drill down for troubleshooting network and application performance issues.



# Unified Performance Management

Network Visibility, Control and Optimization - All in a Single Appliance

**“Unified Performance Management is driven by improving the quality of a user’s experience.”**

**- Ed Ryan, Exinda Vice President of Products**

## The Best Solution For You.

### Identify and Improve Application Performance

- Application Performance Measurement technology measures user experience objectively.
- Identify the source of application performance issues - Network, Server or Application.
- Apply application performance scoring to more than 2,000 applications.

### Offer a Superior User Experience

- Dramatically increases user download speeds for internet applications, videos, and software updates.
- Accelerate delivery of content to users at LAN speeds from a web cache with a single appliance.
- Optimize and accelerate mission critical applications.

### Real-time and Historical Reporting

- Real time reporting showing all traffic on the network over the last 10-60 seconds.
- Up to 2 years of historical reporting on applications, hosts, conversations, URL's, and performance scores “on appliance”.
- Microsoft Active Directory Integration allows you to report on users or groups regardless of IP Address.
- Netflow v9 export, providing in-depth layer 7 details of your network usage and application performance.

### Conserve WAN Resources

- Guarantee bandwidth for critical applications while controlling recreational traffic.
- Byte and Object level caching with dual or single appliances reduces the footprint of traffic on the WAN serving files, software updates, and video to users at LAN speeds.
- Reclaim up to 90% of the bandwidth on your WAN circuits to deliver data more efficiently.

### Leverage Your Investment

- Exinda is fully scalable supporting WAN circuits from 256k to 10Gbps, and includes mobile client support.
- Exinda auto-discovery limits the operational burden and cost of managing large scale multi-site deployments.
- Exinda's Service Delivery Platform (SDP) is available as an appliance or on a cloud-based management platform, offers a flexible and cost-saving option to manage your network.
- A single appliance delivering visibility, control, and optimization makes it easier and more cost-effective to manage and expand over time.

## Features & Benefits

### Visibility

Provides insight into network activity, usage and performance. Gives you the information you need to keep your network operating at peak performance

- Layer 7 Classification
- Heuristic Classification
- URL Classification
- Drill Down Capabilities
- Real Time Monitoring
- Top Talkers/Top Conversations
- Active Directory User ID
- Anonymous Proxy Detection
- Application Performance Score
- Service Level Agreements
- Network Health
- Citrix Published Applications
- Automated PDF Reporting

### Control

Maximize network resources to the needs of your organization through comprehensive control over network traffic without placing heavy-handed restrictions on users.

- QoS / Dynamic per IP User
- Bandwidth Management
- Traffic-shaping
- Prioritization
- Active Directory Integration

### Optimization

Rapidly, turn understanding into action that drives network performance, improves the user experience, and optimizes productivity.

- Layer 4 TCP Optimization
- Layer 7 Application Acceleration
- Universal Caching
- Compression
- Intelligent Acceleration
- Peer Auto-Discovery
- SSL Acceleration

**Americas**  
+1 877 219 0603  
[info.americas@exinda.com](mailto:info.americas@exinda.com)

**UK / Europe**  
+44 808 120 1996  
[info.emea@exinda.com](mailto:info.emea@exinda.com)

**Asia Pacific**  
+61 3 9415 8332  
[info.apac@exinda.com](mailto:info.apac@exinda.com)

**IMEA**  
+971 4 295 5049  
[info.imea@exinda.com](mailto:info.imea@exinda.com)



## EXPAND ENABLES SERVER CONSOLIDATION, THIN-CLIENT COMPUTING AND BANDWIDTH OPTIMIZATION AT RIDLEY INC – DELIVERS SAVINGS OF \$250,000 PER ANNUM

Having initially deployed Expand Networks' Accelerators as part of a bandwidth consolidation project in 2006, Ridley Inc – the leading animal nutrition company - was already aware of the benefits that WAN optimization technology could bring; this initial \$200,000 investment paid for itself through efficiency savings in just over six months.

However, with many of its 42 locations being extremely harsh and dusty environments, Ridley recently embarked on a thin-computing strategy, removing servers and computers from branches and delivering server based computing from a central location in Minnesota.

In order to meet renewed bandwidth requirements and ensure the company's new thin-computing IT initiatives were to succeed, Ridley Inc. re-assessed the company's WAN environment.

Chad Gillick, the IT Manager that led the project at Ridley Inc, explained, "Moving to a thin computing environment could help us streamline processes, increase productivity and reduce costs. However, I knew WAN optimization would be essential to the success of these projects, to ensure the user experience and productivity wouldn't suffer across our distributed network environment."

By investing further in new Expand technology, Ridley has been able to remove expensive desktop and laptop computers at the remote sites and replace with thin client terminals, without costly bandwidth upgrades.

The company chose Expand because of its superior capabilities in accelerating Citrix and web based traffic, and the Accelerators have been deployed in 31 key sites.

Combining compression, byte-level caching, layer 7 QoS and small packets mitigation techniques, Expand's technology enables available bandwidth and real-time interactive TCP traffic to be maximized, extending Ridley existing network infrastructure investments and providing 'virtual bandwidth' capacity to its users.


With substantially faster data transfer speeds over WAN links, Ridley is gaining an estimated 45 minutes of productivity per person, per day. Furthermore, Expand's Wide Area File Services (WAFS) capabilities with QoS have enabled the IT team to tailor traffic flows across the managed network and dynamically manage bandwidth requirements 'on the fly'.

"Without the Expand solution we would have needed a 45mbps connection at the central site that would have cost in the region of \$26,000 per month. With Expand we were able to reduce this to a 9mbps link costing \$4,500, an annual saving of over \$250,000," said Gillick.

He concluded, "On top of this, using Expand as an enabler of server consolidation and thin client computing, we have managed to reduce our technical refresh costs which were running at \$400,000 annually down to \$220,000. We believe we will be reaping the benefits of the Expand solution for many years to come."

### Enabling Strategic Initiatives

- **Virtualization** - The foundation infrastructure for delivering on all strategic IT initiatives, Expand's technology is unique in its combined ability to be deployed within a virtualized infrastructure and to accelerate and control virtualized traffic out of it. The software can be effectively integrated into virtual server environments, such as VMWare, Citrix XenServer and Microsoft HyperV, and as a truly virtualised solution Expand can also be deployed under extreme conditions such as on aircraft, mobile environments and remote and unattended locations..
- **VDI and Thin Computing** - Expand accelerates within Virtual Desktop Infrastructure (VDI) and thin computing environments optimizing protocols including Microsoft Terminal Services (RDP), Citrix XenDesktop (ICA) and Sun Sunray (ALP). Unlike competitive offerings, Expand works on the IP layer, this enables Expand to accelerate all IP & uniquely UDP applications over the WAN, applying advanced compression, byte level caching, layer 7 QoS and small packet mitigation techniques.
- **Server Consolidation** - Expand's integrated 'virtual server' technology enables complete server consolidation by replacing the need for an additional branch office file server. Expand's unique "Virtual Branch Server" feature sets also enable to customer to replace features that used to be delivered by a remote server, such as DHCP, DNS and Printing, all within the AOS and not via third party plug-ins like other vendors.
- **Satellite** - With integrated Space Communication Protocol Specifications (SCPS) Standard technology, Expand helps distributed organizations overcome the traditional limited bandwidth, high latency obstacles that impede the speed and performance of applications and services over satellite links. Communication Protocol Standard technology, helps distributed organizations overcome the traditional low bandwidth, high latency obstacles that impede the speed and performance of applications and services over satellite links.



## Software as a Service (SaaS)

### A Cloud-Ready Network ensures rollout success

[www.ipanematech.com](http://www.ipanematech.com)

Cloud adoption adds complexity to network management. Cloud applications such as SaaS collaboration bring many of the same issues as licensed software, but each IT implementation project can have a larger impact because of its reliance on your WAN. By aligning your network with business and Application Performance Objectives, WAN Governance puts you in control of this complexity and network impact.

WAN Governance improves the IT Governance you already have in place by providing:

- A holistic approach to global visibility, control and optimization of application performance, as opposed to conventional solutions operating as independent agents
- Business continuity and control as SaaS applications are adopted
- Guaranteed application performance for any network architecture
- Network capabilities to absorb enterprise requirements for agility, flexibility and growth
- Next-generation solutions for implementing and managing a cloud-ready network

Using WAN Governance, your organization can:

- Understand the nature of application traffic
- Control and optimize this traffic
- Guarantee application performance
- Improve users' Quality of Experience
- Simplify network operations
- Control network costs and leverage savings


IT infrastructure directors today find themselves in one of two situations: the business side of their organization is planning for SaaS applications that the VPN will need to support, or existing SaaS applications are underperforming or impacting the performance of other business applications.

VPNs and the tools used to manage them are optimized for traditional private applications residing in data centers, not those stored in the cloud. For example, SaaS collaboration applications, such as Google Apps, Microsoft BPOS/Office 365 and IBM LotusLive, consume much more network bandwidth than many traditional applications. Moving from traditional on-premise collaboration to a SaaS counterpart dramatically changes the way traffic flows across the WAN.

In order to avoid application performance issues and ensure optimal end-user experience, infrastructure directors need to make their VPN "cloud ready." A cloud-ready network (CRN) is a network that provides full application performance visibility and total control of both SaaS and on-premise applications. Ideally, the best time to prepare is prior to your first SaaS implementation, so that the impact of SaaS on your VPN can be mastered from the pilot phase through full enterprise rollout.

With Ipanema for a fraction of the cost per user of your SaaS you can:

- Guarantee the performance of SaaS across the WAN
- Ensure peaceful co-existence of SaaS and existing applications (ERP, CRM...)
- Obtain a dashboard of application performance for all critical applications including SaaS
- Take full advantage of hybrid MPLS + Internet networks
- Shift to WAN governance, plan and grow your network according to business needs







www.ipanematech.com



### Valeo Embraces the Cloud and Maximizes Value

Valeo, one of the world's leading suppliers of components, integrated systems and mod-ules for automotive CO2 emissions reduction, rolled out a hybrid network with MPLS + Internet for its migration from conventional email and collaboration applications to Google Apps.

Valeo's network supports approximately 160 sites worldwide, 52,000 users, and the delivery of applications such as ERP and CATIA.

Using Ipanema's ANS to dynamically manage application performance over their hybrid network, Valeo successfully deployed Google Apps with full Applications Visibility, QoS & Control, and Dynamic WAN Selection.

*"With Ipanema, we divided by three the transfer cost of each Gbyte of band-width over our global network," says Alain Meurou, Infrastructure and Network Manager at Valeo.*

### Return on investment

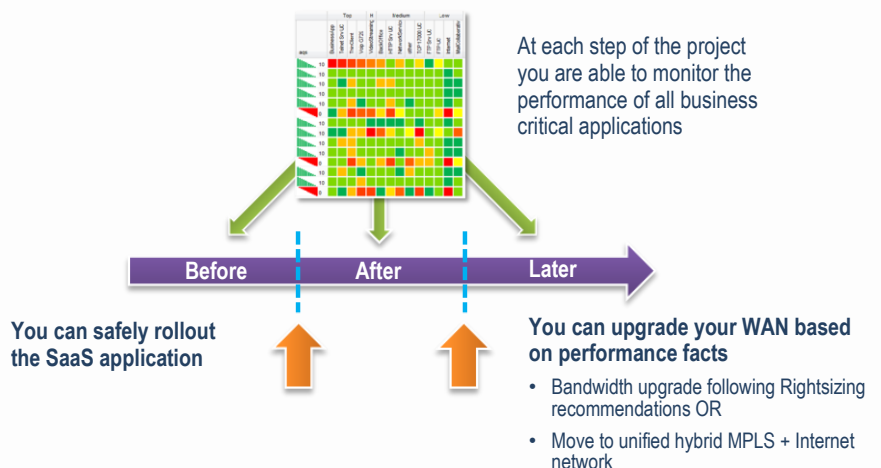
Enterprises that have chosen to move to a unified hybrid net-work controlled by ANS typically chose not to upgrade their MPLS bandwidth in favor of the less-expensive Internet bandwidth. Including the price of the deployed ANS solution, typically 1 to 2 € per user per month, most enter-prises were able to obtain a 20% decrease in overall network costs, upgrade available band-width by a factor of three, and adequately pre-prepare for traffic increases over the next three to five years.

### All-in-One Solution for Guaranteeing Application Performance

Ipanema's Autonomic Networking System (ANS) tightly couples into a single, all-in-one solution.

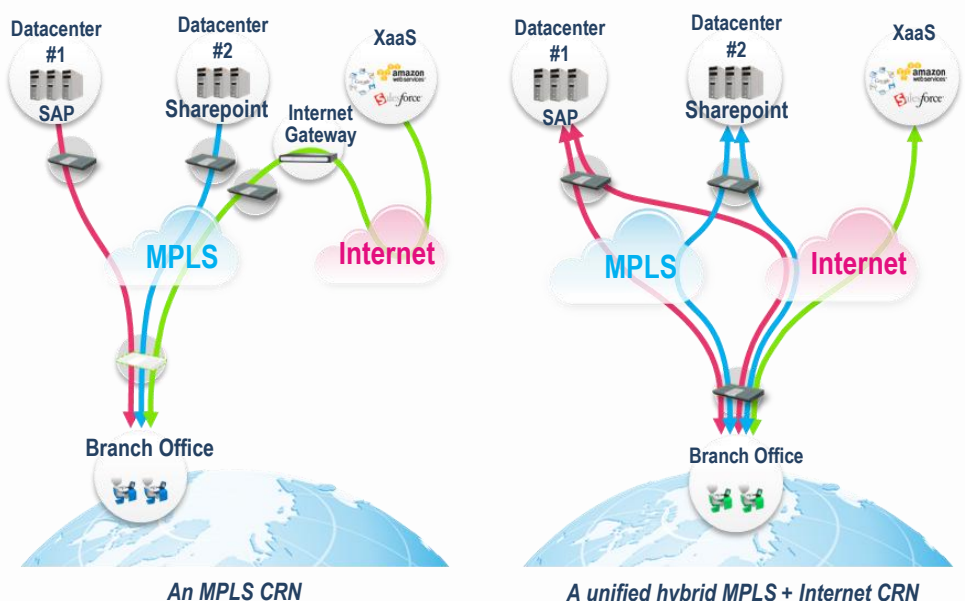
- QoS & Control
- Application Visibility
- WAN Optimization
- Dynamic WAN Selection (hybrid network unification)

With ANS, all application performance challenges can be managed with a holistic approach over the global network. The autonomic networking solution automates tasks that IT organizations cannot perform with traditional approaches. Orchestrating network traffic in real-time, ANS manages the complexity of the hybrid cloud and guarantees application performance for public and private applications. ANS not only helps to guarantee the performance of SaaS during and after implementation, but the end-user experience for all applications over your WAN, and much more cost effectively.



Since every enterprise is different, IT strategy on whether or not to change network architecture for SaaS collaboration varies from one company to another. You do not necessarily need to change your architecture to make your network "cloud-ready".

All companies, however, must implement a minimum set of capabilities in order to avoid application performance issues during and after SaaS implementation, or to fix issues resulting from a prior SaaS deployment. Companies that use or plan to use a hybrid (MPLS + Internet) network architecture will also want to consider additional capabilities to further optimize their "cloud-ready network" (CRN).



## Packet Design

# Network-Wide IP Routing and Netflow Monitoring, History, Modeling & Planning

### Optimize IP Networks with Traffic Explorer

- Monitor and analyze critical traffic dynamics across all IP network links and routes by Class of Service (CoS)
- Strengthen change management with operationally accurate network modeling based on real-time, network-wide routing and traffic state
- Reduce Internet transit costs with IGP/BGP-aware peering and transit routing and traffic analysis
- Analyze network-wide traffic usage, even per MPLS VPN
- Improve network continuity with easy traffic trending
- Perform network-wide traffic capacity planning

### Overview of Traffic Explorer

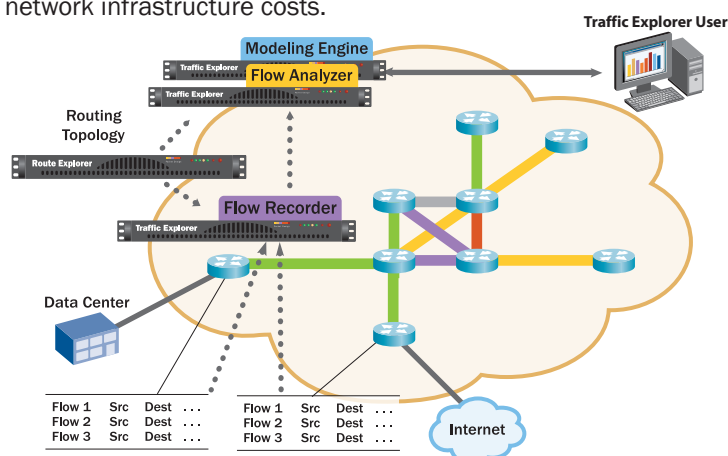
Traffic Explorer is the first solution to combine real-time, integrated routing and Netflow traffic monitoring and analysis, with “what-if” modeling and capacity planning capabilities. Unlike previous traffic analysis tools that only provide localized, link by link traffic visibility, Traffic Explorer’s knowledge of IP routing enables visibility into network-wide routing and traffic behavior. Powerful “what-if” modeling capabilities empower network managers with new visibility to strengthen change management processes and optimize network infrastructure costs.

Traffic Explorer delivers the industry’s only integrated analysis of network-wide routing and traffic dynamics. Standard reports and threshold-based alerts help engineers track significant routing and utilization changes in the network. An interactive topology map and deep, drill-down tabular views allow engineers to quickly perform root cause analysis of important network changes, including the routed path for any flow, network-wide traffic impact of any routing changes or failures, and the number of flows and hops affected. This information helps operators prioritize their response to those situations with the greatest impact on services or applications.

Traffic Explorer provides extensive “what-if” planning features to enhance ongoing network operations best practices. Traffic Explorer lets engineers model changes on the “as running” network, using the actual routed topology and traffic loads. Engineers can simulate a broad range of changes, such as adding or failing routers, interfaces and peerings; moving or changing prefixes, BGP policy configurations, link capacities or traffic loads; even adding new MPLS VPNs. Simulating the effect of these changes on the actual network results in faster, more accurate network operations and optimal use of existing assets, leading to reduced capital and operational costs and enhanced service delivery.

### Packet Design Overview:

- Founded in 2003, Packet Design pioneered and is the market-leading provider of routing and traffic analysis solutions
- 500+ global enterprises, Service Providers, and government and military agencies utilize Packet Design solutions to manage their complex IP networks.
- Packet Design solutions offer IT departments significant operational cost savings by increasing the accuracy and efficiency of key IT business processes.



**Proven, Market Leading Solutions:** Based in Palo Alto, Packet Design Inc. is the pioneer and market leader in routing-aware network management solutions. Packet Design is a member of the Cisco Technology Developer Partner program. Find out more at [www.packetdesign.com](http://www.packetdesign.com)



## Why an Application Delivery Fabric is Essential for Agile & Scalable Virtualization



### Data Center Virtualization and Application Delivery


Network infrastructure virtualization/consolidation has had a major impact on the Application Delivery Controller (ADC) role, position and deployment models. For instance, ADCs which were previously tightly coupled with a single application must now be able to service a layer of virtualized applications sharing a common server infrastructure.

### A New Paradigm: Virtual Application Delivery Infrastructure (VADI)

Radware's VADI strategy allows for the transformation of computing resources and ADC devices and software into an integrated, agile and scalable set of application delivery services that can be dynamically provisioned, decommissioned, and migrated.

Radware's VADI delivers the following business benefits:

- ✓ Significant cost reduction via ADC consolidation
- ✓ Simpler path to data center virtualization
- ✓ Improved business agility
- ✓ Greater IT efficiency via data center workflow automation
- ✓ Full application delivery resource elasticity
- ✓ On demand scalability in throughput, advanced services and virtual ADCs
- ✓ Full investment protection, increased asset ROI, and CAPEX savings



### VADI Key Components

#### Virtual ADC Instances

A Virtual ADC (vADC) instance is a service providing a consistent and complete set of application delivery features such as load balancing, global server load balancing, application acceleration, integrated security, bandwidth management and more. A vADC runs on top of specialized and general purpose computing resources, thus transforming ADCs into services.

#### ADC Computing Resources via Three Form-Factors

- Dedicated ADC – a dedicated, physical ADC device running a single vADC, which is designed to provide application delivery services for siloed data center architectures, hybrid (virtualized and physical) data centers, and applications requiring high SLA and performance predictability.
- ADC-VX™ - the industry's first ADC hypervisor that runs multiple vADCs on a dedicated ADC hardware, Radware's OnDemand Switch platform.
- Alteon VA™ - Radware's Soft ADC is a vADC deployed on a general server virtualization infrastructure, running as a virtual appliance, providing the full functionality of a physical ADC.

#### Virtual Data Center Ecosystem Integration

Radware's vDirect™ is the industry's first ADC management orchestration plug-in, designed specifically for virtual data centers. It provides all the building blocks and management interfaces required for an orchestration system to provision, decommission, configure and monitor Radware's vADCs and computing resources within a virtual data center.

#### Advanced VADI Services

VADI services, such as ADC service provisioning, decommissioning, and migration of virtual ADC instances across form factors, enables business agility goals while delivering the matching resilience and SLA per application. Radware provides various VADI services such as:

- Provisioning and decommissioning - vADCs are instantly provisioned and/or decommissioned through the ADC management system or orchestration systems' API
- vADC migration - Easily move a vADC instance between different form factors, allowing scheduling ADC maintenance with zero downtime, thus reducing the potential loss of business and revenue
- Dynamic elasticity - Dynamically instruct the orchestration system to allocate additional resources for an application when the existing computing resources are completely utilized
- Cloud burst - Dynamically instruct the orchestration system to allocate additional resources in the cloud or in a second data center when the resources are completely utilized in the main data center

For additional information on ADC-VX, Alteon VA and vDirect please refer to <http://www.radware.com/Solutions/Enterprise/Virtualization/DataCenterVirtualization.aspx> or for customer case examples please visit our press release section: <http://www.radware.com/NewsEvents/PressReleases.aspx>.