# Application Delivery

## Reality Report

***By Dr. Jim Metzler***

Sponsored by:




Produced by

# Table of Contents

## 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. As part of this focus, many IT organizations began to deploy a first generation of solutions that were intended to protect the organization from a growing number of security attacks, to mitigate the impact of chatty protocols such as CIFS (Common Internet File System), to offload computationally intensive processing (e.g., TCP termination and multiplexing) from servers and to manage application performance on an end-to-end basis. Throughout this white paper, the application delivery challenges and solutions of this era will be referred to as Application Delivery 1.0.

However, 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
- Shifting emphasis and growing sophistication of cyber crime
- Adoption of varying forms of virtualization
- Adoption of cloud computing

Throughout this white paper, this new generation of application delivery challenges and solutions will be referred to as Application Delivery 2.0.

The goal of this white paper is to provide a reality check on the status of application delivery with a focus on optimization and management. Towards that end, this paper will identify which optimization and management challenges are currently top of mind for IT organizations and will also identify what steps IT organizations plan to take to respond to these challenges. One of the advantages of identifying the steps that IT organizations are taking to respond to a set of challenges is that this information enables an IT organization to see how their approach compares to that of their peers. That is important because most IT organizations do not want to be too far outside of the mainstream. That is particularly important at this point in time given the fundamental changes being driven by the challenges listed above.

## The Changing Environment

Throughout this report, the phrase *ensuring acceptable application delivery* will refer to ensuring that the applications that an enterprise uses:

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

This section of the report will discuss some of the factors that make ensuring acceptable application delivery complex today and will also discuss some additional factors which have begun to add dramatically to the complexity of ensuring acceptable application delivery.

### The Traditional Complex Application Delivery Environment

One of the factors that makes the traditional IT environment complex is the broad use of distributed applications. For example, most IT organizations have deployed a form of distributed computing often referred to as *n-tier applications*. The browser on the user's device is typically one component of an n-tier application. The typical 4-tier application (**Figure 1**) is also comprised of a Web tier, an application tier and a data base tier which are implemented on a Web server(s), an application server(s) and a database server(s). In the traditional IT environment, none of the servers are virtualized.

**Figure 1: A 4-Tier Application**

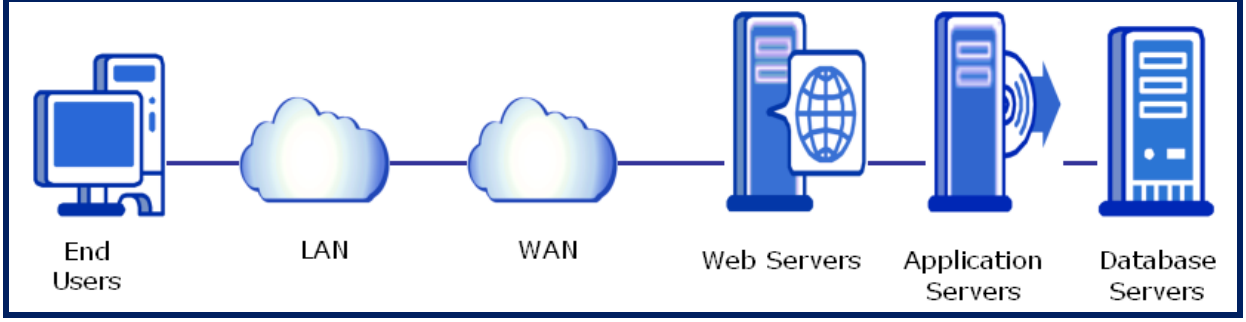


The complexity of the traditional n-tier application comes in part from the fact that each tier of the application is typically implemented on a separate system from which management data must be gathered. It also comes from the fact that the networks that support these applications are comprised of a variety of switches, routers, access points, WAN optimization controllers, application delivery controllers, firewalls, intrusion detection systems and intrusion protection systems from which management data must also be gathered.

The complexity of the traditional n-tier application also comes in part from the fact that if any of the application components depicted in Figure 1 are not available, or are not performing well, the performance of the overall application or service is impacted. In some instances, each component of the application architecture is performing well, but due to the sheer number of components, the overall delay builds up to a point where some function, such as a database query, fails.

Due to the complexity of the current environment, when the performance of an application or service is degrading, that degradation is typically noticed first by the end users and not by the IT organization. In addition, when the IT organization is made aware of the fact that the performance of an application or service has degraded, it typically is unaware of the cause of the degradation. Given the sheer number of service components, the time to troubleshoot the problem can be lengthy. The combination of these facts tends to diminish the reputation of the IT organization. As the IT environment becomes more complex due to the factors discussed below, the typical IT organization is going to become even less aware of performance issues and take even more time to troubleshoot problems if the organization does not fundamentally change its approach to management.

## Emergence of a Sophisticated Mobile Workforce

In many cases, the challenges that are associated with the Application Delivery 2.0 era are the logical extension of the challenges that were associated with the Application Delivery 1.0 era. In some cases, the challenges have not changed over the last five years. For example, optimizing the performance of TCP was a challenge in the Application Delivery 1.0 era and, as will be discussed below, it remains a challenge in the Application Delivery 2.0 era.

In other cases, however, the challenges associated with The Application Delivery 1.0 era have evolved to become new challenges in The Application Delivery 2.0 era. For example, one of the changes that is associated with the Application 1.0 era is the fact that many employees who had at one time worked in a headquarters facility now work someplace other than a headquarters facility; i.e., a regional, branch or home office. As a result of this decentralization of employees combined with the ongoing initiative that most IT organizations have to consolidate servers into centralized data centers, most employees now access applications over a relatively low-speed, high-latency WAN. One of the challenges associated with the Application Delivery 1.0 era was that running a chatty protocol[1] such as CIFS (Common Internet File System) over a WAN usually resulted in unacceptable performance. As a result, optimizing the performance of a chatty protocol was one of the most important tasks of The Application Delivery 1.0 era. As will be discussed below, it is not one of the most important tasks of The Application Delivery 2.0 era.

[1] Chatty protocols typically require hundreds of round trips to complete a single transaction.

In The Application Delivery 2.0 era the challenge of supporting decentralized employees has evolved because many of these employees are now mobile[2]. What has also evolved is the type of applications that these workers access. At one time, mobile workers tended to primarily access either recreational applications or applications that are not delay sensitive; e.g., email. However, in the current environment mobile workers also need to access a wide range of business critical applications, many of which are delay sensitive. One of the challenges associated with supporting mobile workers access to delay sensitive, business critical applications is that because of the way that TCP functions, even the small amount of packet loss that is often associated with wireless networks results in a dramatic reduction in throughput.

## Shifting Emphasis and Growing Sophistication of Cyber Crime

Preventing security breaches was a key component of the Application Delivery 1.0 era. One key thing that has evolved over the last couple of years, however, is the level of sophistication of cyber crime. For example, McAfee recently published a report[3] based on a survey of 800 CIOs that was performed by Purdue University's Center for Education and Research in Information Assurance and Security. The McAfee Report discussed some of the ways that cyber crime was maturing and becoming more sophisticated. The report stated that, "Credit card fraud and identity theft have moved into the so-called 'cash cow' phase of criminal strategy. In other words, it's a source of revenue, but there's not much room for growth, so criminals are looking for the new stars of their portfolios. And intellectual property has emerged as a favorite." Also included in the report was the observation that many malware writers now have R&D and test departments.

The Center for Strategic and International Studies (CSIS) recently released a report[4] entitled "In the Crossfire – Critical Infrastructure in the Age of Cyber-War" that provided further evidence that cyber crime is becoming more sophisticated and costlier. That report included the results of a survey that was completed by six hundred IT and security executives in 14 countries. According to the report, "More than half of the executives surveyed (54 percent) said they had experienced 'Large-scale denial of service attacks by high level adversary like organized crime, terrorists or nation-state (e.g. like in Estonia and Georgia).' The same proportion said they had been subject to 'stealthy infiltration" of their network by such a high-level adversary "e.g. like GhostNet'—a large-scale spy ring featuring individualized malware attacks that enabled hackers to infiltrate, control, and download large amounts of data from computer networks belonging to non-profits, government departments and international organizations in dozens of countries."

---

[2] One analyst firm has predicted that there will be one billion mobile workers worldwide by 2011 - http://findarticles.com/p/articles/mi_m0EIN/is_2008_Jan_15/ai_n24230213/

[3] http://www.mcafee.com/us/about/press/corporate/2009/20090129_063500_j.html

[4] http://csis.org/event/crossfire-critical-infrastructure-age-cyber-war

Another aspect of the changing nature of cyber crime is the sheer scale of the attacks. In January 2010 Arbor networks published their Fifth Annual Infrastructure Security Report[5]. According to that report, the peak rate of DDoS attacks has grown from 400 Mbps in 2001 to 49 Gbps in 2009. The report also stated that, "We expect DDoS attack rates to continue to grow, but given that most enterprises are still connected to the Internet at speeds of one gigabit per second (Gbps) or less, any attack over one Gbps will be typically effective, and often trigger collateral damage to adjacent network or customer service elements as well."

## Adoption of Varying Forms of Virtualization

Most of the current interest in virtualization revolves around the virtualization of servers. In a recent survey, three hundred and thirty nine IT professionals were asked to indicate the percentage of their company's data center servers that have either already been virtualized or that they expected would be virtualized within the next year. Their responses are shown in **Table 1**.

**Table 1: Deployment of Virtualized Servers[6]**

| | None | 1% to 25% | 26% to 50% | 51% to 75% | 76% to 100% |
|---|---|---|---|---|---|
| **Have already been virtualized** | 21.6% | 33.0% | 18.9% | 15.1% | 11.3% |
| **Expect to be virtualized within a year** | 12.4% | 25.6% | 21.9% | 21.9% | 18.2% |

The data in Table 1 shows the ongoing interest that IT organizations have relative to deploying virtualized servers. Two observations that can be drawn from Table 1 are that within the next year:

- The number of IT organizations that have not implemented server virtualization will be cut almost in half.
- The number of IT organizations that have virtualized the majority of their servers will grow by sixty percent.

Some of the challenges associated with server virtualization include being able to:

- Manage the traffic that goes between virtual machines (VMs) on a single physical server

[5] http://www.marketwire.com/press-release/Arbor-Networks-Fifth-Annual-Infrastructure-Security-Report-Finds-Service-Application-1103590.htm
[6] http://www.webtorials.com/content/2009/12/cloud-computing-a-reality-check-guide-to-risk-mitigation.html

- Keep track of VMs as they are moved between physical servers
- Discover VMs
- Dynamically move VMs, and all of the supporting management functionality, between physical servers
- Perform traditional management tasks such as troubleshooting and performance management, on a per VM basis

While server virtualization receives far more attention than does any other form of virtualization, there is growing interest in other forms of virtualization such as desktop virtualization. For example, the results (**Table 2**) of the same survey as previously referenced indicate that within the next year:

- The number of IT organizations that have virtualized the majority of their desktops will almost double.
- The number of IT organizations that have not implemented desktop virtualization will be cut in half.

**Table 2: The Percentage of Desktops that Already Have or Will be Virtualized**

| | None | 1% to 25% | 26% to 50% | 51% to 75% | 76% to 100% |
|---|---|---|---|---|---|
| Have already been virtualized | 49.5% | 34.7% | 8.9% | 1.0% | 5.9% |
| Expect to be virtualized within a year | 22.0% | 46.3% | 18.3% | 7.3% | 6.1% |

The primary factor that limits the deployment of virtualized desktops is the concern that IT organizations have over the performance of the solution. The performance problems that are associated with virtualized desktops are a result of the fact that all of the data that supports the applications and/or the desktops has to transit the WAN from one or more data centers to the user. Protocols optimized for hosted application virtualization such as ICA, RDP and PCoIP are more efficient in their use of the WAN than are chatty protocols such as CIFS. However, network latency can still cause the performance of these protocols to degrade. This can have a very negative impact on tasks such as a user moving his/her mouse around a computer screen.

## Adoption of Cloud Computing

Cloud Computing Service Providers (CCSPs), such as Akamai and Google, provide services either over the public Internet or other WAN services. They are offering a class of solution that is often referred to as the *Public Cloud*. Public cloud is what many IT

organizations think of when they hear the phrase *cloud computing*. The primary types of services provided by CCSPs are:

- Software-as-a-Service (SaaS)
- Infrastructure-as-a-Service (IaaS)
- Platform-as-a-Service (PaaS)

As identified in a recent report[7], there is a well-defined set of characteristics that are typically associated with public cloud computing solutions; e.g., virtualization, automation, orchestration. Many IT organizations have decided to implement services themselves within their own environment that reflect these same characteristics. This is what is meant today by the phrases *Private Cloud* or *Private Cloud Computing*. Private Clouds have the advantages of not being burdened by many of the potential security exposures, data confidentiality and control issues that are associated with public clouds. A third class of cloud computing solution is a *hybrid* solution. As is illustrated below, a hybrid solution is comprised of a combination of both private and public cloud services.

Because virtualization is often associated with cloud computing, the challenges associated with virtualization are typically also challenges for cloud computing. One of the additional challenges associated with cloud computing is being able to ensure acceptable performance for the services that are acquired from a CCSP and to manage the service on an end-to-end basis. Part of managing on an end-to-end basis means being able to troubleshoot performance problems.

To put the challenge of troubleshooting a cloud computing service in perspective, it is worth repeating the fact that in the traditional IT environment of The Application Delivery 1.0 era, it is the end user, and not the IT organization that first notices application degradation. In addition, once the IT organization is aware of the fact that there is a performance problem, the time to identify the root cause can be quite lengthy.

**Figure 2** depicts the typical 4-tier application that was previously discussed. However, this time the application is delivered in one of the many ways that IT organizations are

**Figure 2: The New 4-Tier Application: Virtualized Hybrid Cloud**

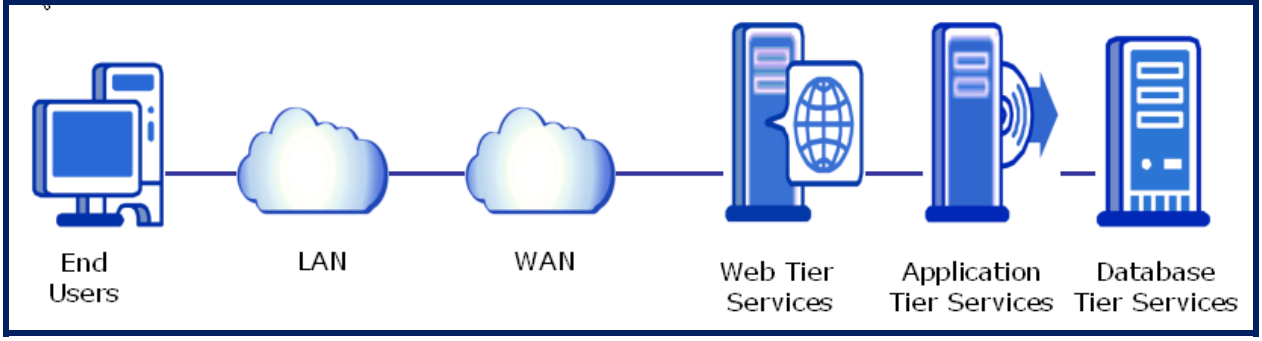


[7] http://www.webtorials.com/content/2009/11/a-guide-for-understanding-cloud-computing.html

implementing hybrid cloud solutions. In this example:

- The enterprise IT organization is responsible for managing the wired and wireless LANs that the company's employees use to access myriad WAN services.
- The WAN is comprised of an MPLS service from one service provider and Internet access services from multiple providers. Some of the Internet access is wired and some is wireless.
- A CCSP is hosting the Web access tier in virtualized servers in multiple data centers.
- The enterprise IT organization hosts the application and database tiers of the application in virtualized servers in their primary data center.

In this example, effective end-to-end management of the performance of the application requires gathering multiple levels of management data from disparate organizational domains; i.e., the enterprise IT organization, multiple network service providers and one cloud computing service provider. Within a given organizational domain (e.g., the enterprise IT organization), management data must be gathered on multiple technology domains; e.g., networks, servers, security devices. However, just being able to gather this data is not sufficient. The data needs to also be effectively processed to provide the information necessary to perform tasks such as identifying the root cause of poor application performance. Identifying the root cause of poor application performance in a public cloud computing solution such as the one described above is an order of magnitude more difficult than it is in the traditional IT environment.

## Optimization

This section of the report will focus on the optimization component of application delivery. In particular, this section will discuss both the optimization challenges that IT organizations find to be the most important as well as what they intend to do in the next year to respond to those challenges. This section, and the subsequent section on management, is based in large part on two surveys that were administered in the first quarter of calendar year 2010 to the subscribers of Webtorials. Throughout this report, the IT professionals who completed those surveys will be referred to as The Survey Respondents.

### Optimization Challenges

The Survey Respondents were asked to indicate how important it is to their organization to get better at seventeen optimization tasks over the next year. They were given the following five-point scale:

1. Not at all important
2. Slightly important
3. Moderately important
4. Very important
5. Extremely important

The Survey Respondents were also asked to indicate how difficult it would be for their organization to get better at each of the seventeen optimization tasks over the next year. They were given the following five-point scale:

1. Not difficult
2. Slightly difficult
3. Moderately difficult
4. Significantly difficult
5. Very significantly difficult

**Table 3** shows the ten application delivery related optimization tasks that are the most important for IT organizations to improve on in the next year. Included in Table 3 are the task and the percentage of The Survey Respondents who indicated that the task was either very or extremely important to get better at over the next year. Also included is the percentage of The Survey Respondents who indicated that it would either be significantly or very significantly difficult for their IT organization to get better at the task in the next year.

| Table 3: Top Ten Optimization Tasks | | |
|---|---|---|
| **Optimization Task** | **Importance: Very or Extremely** | **Difficulty: Significant or Very Significant** |
| Relating the performance of applications to their impact on the business | 70% | 39% |
| Ensuring acceptable performance for VoIP traffic | 68% | 29% |
| Improving the performance of applications used by mobile workers | 60% | 38% |
| Ensuring acceptable performance for video or telepresence traffic | 57% | 31% |
| Ensuring acceptable performance for the applications that you acquire from a Software-as-a-Service provider | 56% | 32% |
| Optimizing the performance of TCP | 54% | 19% |
| Controlling the cost of the WAN by reducing the amount of traffic that transits the WAN | 50% | 31% |
| Optimizing the Web tier of a multi-tiered application for peak utilization | 50% | 23% |
| Optimizing the performance of specific applications such as SharePoint | 49% | 19% |
| Optimizing the performance of protocols other than TCP; e.g., HTTP and MAPI | 49% | 24% |

An analysis of the right hand column in Table 3 (the column that indicates the difficulty of getting better at the task) shows that relative difficulty of the tasks as indicated by The Survey Respondents seems reasonable. For example, The Survey Respondents indicated that of the ten tasks listed in Table 3, "Relating the performance of applications to their impact on the business" was the most difficult and "Optimizing the performance of TCP" was the least difficult. This makes sense as the former task involves both implementing sophisticated tools and changing organizational culture while the subsequent task involves implementing well understood, somewhat narrowly focused technology.

Another measure of the creditability of the data in Table 3 is that 68% of The Survey Respondents indicated that ensuring acceptable performance for VoIP traffic was an important task for them to get better at next year and 29% indicated that it would be significantly or very significantly difficult to do so. This is very consistent with the data in **Table 4** (below) in which 67% of The Survey Respondents indicated that effectively

managing QoS was an important task for them to get better at next year and 29% indicated that it would be significantly or very significantly difficult to do so.

One obvious conclusion that can be drawn from the data in Table 3 is that there are a large number of optimization tasks that are important for IT organizations to get better at in the next year. The data also indicates that IT organizations don't think that getting better at those tasks is very difficult. One possible reason for that is the *can do* attitude of most IT professionals. However, even if getting better at a single task, such as optimizing the performance of TCP, is not that difficult unto itself, marshalling the resources to get better at multiple tasks will be very difficult for most IT organizations.

The data in Table 3 does raise some interesting questions. One such question is that given that VoIP has been widely deployed for many years, why is it still so important to get better at ensuring acceptable performance for VoIP traffic if the task is not that difficult? Also, most Software-as-a-Service (SaaS) providers do not provide any form of a service level agreement (SLA) relative to the performance of the applications that they provide. In addition, in the majority of instances users access applications from a SaaS provider by using the Internet, and there is no SLA associated with the use of the Internet. Because of factors such as these, it is reasonable to suggest that The Survey Respondents underestimated how difficult it will be to get better at certain tasks such as ensuring acceptable performance for the applications that they acquire from a SaaS provider.

One of the tasks that did not make the list of top ten optimization tasks was optimizing the performance of chatty protocols such as CIFS. This is interesting in part because as was previously alluded to, this is one of the tasks that ignited the initial deployment of WAN optimization controllers in The Application Delivery 1.0 era. Thirty seven percent of The Survey Respondents indicated that this task was either very or extremely important. While not enough to make it into Table 3, that is still a large percentage. However, the fact that another task that didn't make the top ten list, optimizing the performance of virtualized desktops, was deemed to be more important indicates that somewhat of a shift is occurring in terms of the types of optimization tasks that are most important to IT organizations. In particular, some of the emerging challenges associated with the Application Delivery 2.0 era are becoming more important than some of the traditional challenges associated with the Application Delivery 1.0 era.

## Optimization Solutions

The Survey Respondents were asked to indicate their organization's approach to optimizing network and application performance. They were given four choices:

- We implement very little if any functionality specifically to optimize network and application performance.

- We implement optimization functionality primarily on a case-by-case basis in response to high visibility problems.
- We have broadly implemented optimization functionality throughout our organization.
- Other

Their responses are shown in **Figure 3**.

**Figure 3: Current Deployment of Optimization Functionality**

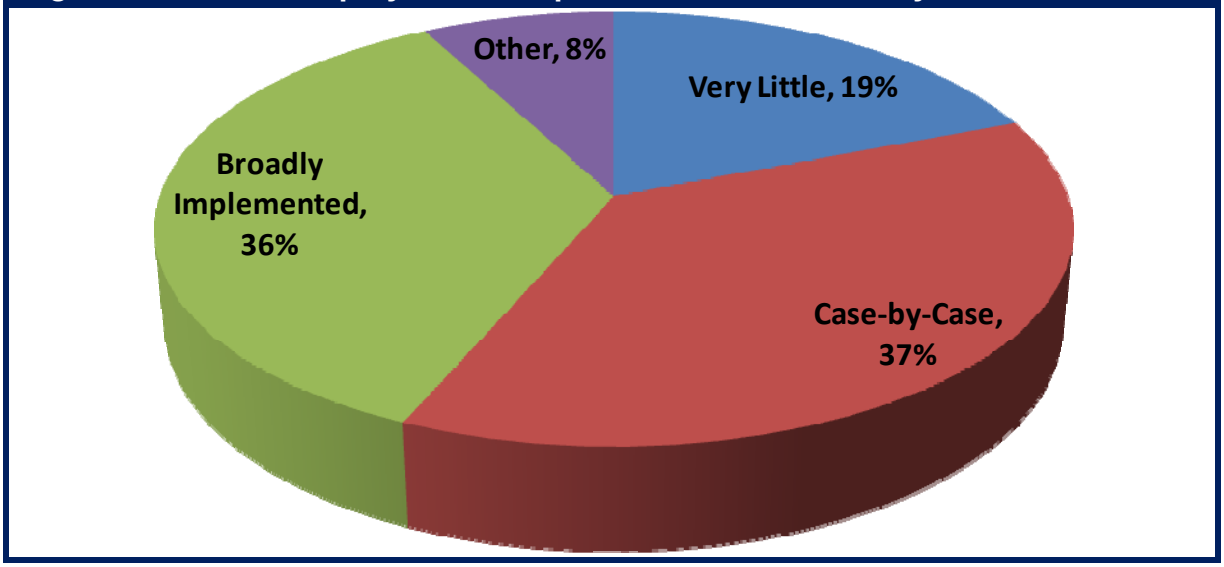


The most common response in the *other category* is that the IT organization doesn't deploy any optimization functionality itself in part because the organization relies on functionality provided as part of the MPLS service that it uses.

The data in Figure 3 is a classic good news – bad news situation. The good news is that just over a third of IT organizations have broadly implemented optimization functionality throughout their organization. The bad news is that two thirds of IT organizations have not. In order to quantify the current and planned deployment of WAN Optimization Controllers (WOCs) The Survey Respondents were asked to indicate the percentage of locations (e.g., branch offices and end users' computers) to which their organization has either already deployed WOC functionality or will deploy it within the next year. The Survey Respondents were given three possible form factors:

- WOCs that are appliance based
- A soft or virtualized WOC
- A software client running on a user's desktop

Roughly two thirds of IT organizations have already deployed WOCs that are appliance based, although only slightly more than a quarter of IT organizations have deployed WOCs to the majority of their locations. The Survey Respondents indicated that they would make only a modest increase in the deployment of appliance based WOCs over the next year.

Less than half of The Survey Respondents indicated that their organization had either already deployed a soft WOC or a software client running on users' desktops. However, The Survey Respondents indicated that their IT organization would increase their deployment of both forms of optimization tools in the next year. In particular, twenty percent of The Survey Respondents indicated that within a year that the majority of the desktops that they support will have a software client running on it the purpose of which is to optimize network and application performance. This result is highly consistent with the data in Table 3 in which sixty percent of The Survey Respondents indicated that improving the performance of applications used by mobile workers was either very or extremely important to their organization in the coming year.

In order to quantify the current and planned deployment of Application Delivery Controllers (ADCs) The Survey Respondents were asked to indicate the percentage of their application traffic that is front-ended by an ADC either currently or within the next year. The Survey Respondents were given two possible form factors:

- Appliance based ADCs
- Software based ADCs

Over half of IT organizations already front-end at least some of their application traffic with an appliance based ADC, and roughly a quarter of IT organizations front-end the majority of their application traffic with an appliance based ADC. Both of those percentages will increase slightly over the next year. Less than half of IT organizations already front-end at least some of their application traffic with an software based ADC, and that percentage will increase slightly over the next year.

The Survey Respondents were asked to indicate how their organization's budget in 2010 for WOC and ADC functionality compared to what they budget was in 2009. Their answers are shown in **Table 4**.

**Table 4: 2010 Budget vs. 2009 Budget**

| | Reduced by 10% or more | Reduced by less than 10% | About the same | Increased by less than 10% | Increased by more than 10% |
|---|---|---|---|---|---|
| **WOC** | 11.3% | 6.3% | 53.8% | 12.5% | 16.3% |
| **ADC** | 9.7% | 8.3% | 54.2% | 13.9% | 13.9% |

The data in Table 4 shows that in 2010 the budget that the vast majority of IT organizations (80+%) have for WOC and ADC functionality will either stay the same as it was in 2009 or it will increase.

## Management

This section of the report will focus on the management component of application delivery. In particular, this section will discuss both the management challenges that IT organizations find to be the most important as well as what they intend to do to respond to those challenges.

### Management Challenges

In addition to evaluating seventeen optimization tasks, The Survey Respondents were also asked to evaluate twenty management tasks. **Table 5** shows the ten management tasks that are the most important for IT organizations to improve on in the next year.

**Table 5: Top Ten Management Tasks**

| Management Task | Importance: Very or Extremely | Difficulty: Significant or Very Significant |
|---|---|---|
| Rapidly identify the root cause of degraded application performance | 76% | 47% |
| Identify malicious traffic and eliminate it | 71% | 31% |
| Effectively manage QoS | 67% | 29% |
| Prevent large scale DDOS attacks | 66% | 32% |
| Identify the components of the IT infrastructure that support the company's critical business applications | 66% | 20% |
| Obtain performance indicator metrics and granular data that can be used to detect and eliminate impending problems | 64% | 38% |
| Effectively mange services, where services are comprised of multiple, inter-related applications | 61% | 39% |
| Effectively manage SLAs for one or more business critical applications | 61% | 35% |
| Obtain real-time, or nearly real-time, insight into how specific applications and end user sessions are performing | 59% | 43% |
| Track end user experience and relate it to factors such as Web response time | 51% | 40% |

Similar to the format in the preceding section, included in Table 5 are the tasks and the percentage of The Survey Respondents who indicated that each task was either very or extremely important for their organization to get better at over the next year. Also included is the percentage of The Survey Respondents who indicated that it would either be significantly or very significantly difficult for their IT organization to get better at the task in the next year.

As was the case with optimization, the data in Table 5 indicates that there is a large number of management tasks that are important for IT organizations to get better at in the next year. Comparing the data in Table 4 with the data in Table 5 indicates that The Survey Respondents believe that getting better at the top ten management tasks will be more difficult than getting better at the top ten optimization tasks.

Given the previously discussed shifting emphasis and growing sophistication of cyber crime, it was not surprising to see that two of the management tasks that are most important to The Survey Respondents are identifying malicious traffic and eliminating it, and preventing large scale DDoS attacks. It was also not surprising that the management task that is the most important for The Survey Respondents to get better at during the next year is rapidly identifying the root cause of degraded application performance. That follows because as previously noted, in the current environment the end user typically notices application degradation before the IT organization does and when the IT organization is made aware of the fact that the performance of an application is degrading, it typically is unaware of the cause of the degradation. The fact that 47% of The Survey Respondents stated that getting better at identifying the root cause of degraded application performance would be either significantly difficult or very significantly difficult is in line with the fact that as previously discussed, the movement to adopt virtualization and cloud computing will make troubleshooting an order of magnitude more difficult that it is currently.

The data in Table 5 exhibits a lot of consistency. For example, given that rapidly identifying the root cause of degraded application performance was so important to The Survey Respondents, it makes sense that obtaining performance indicator metrics and granular data that can be used to detect and eliminate impending problems was also regarded by The Survey Respondents as being important. In addition, given that effectively managing SLAs for one or more business critical applications is so important to The Survey Respondents, it makes sense that identifying the components of the IT infrastructure that support the company's critical business applications is also important to The Survey Respondents.

One of the interesting results of the survey is that effectively managing services, where services are comprised of multiple, inter-related applications is very important to IT organizations. That is interesting in large part because over the last few years that has been considerable discussion about the fact that IT organizations need to move away from focusing on managing individual technology domains and need to focus more on

managing the end-to-end performance of applications. There hasn't been anywhere near as much discussion about the need to focus on managing services where a service is comprised of multiple, inter-related applications.

The preceding section of this report entitled *Adoption of Varying Forms of Virtualization* identified a number of the challenges associated with server virtualization. None of those challenges unto themselves were important enough to make it into Table 5. The importance and the difficulty of these challenges as seen by The Survey Respondents are shown in **Table 6**.

**Table 6: Importance of Server Virtualization Tasks**

| Server Virtualization Management Task | Importance: Very or Extremely | Difficulty: Significant or Very Significant |
|---|---|---|
| Perform traditional management tasks such as troubleshooting and performance management, on a per VM basis | 49% | 20% |
| Keep track of VMs as they are moved between physical servers | 38% | 17% |
| Dynamically move VMs, and all of the supporting management functionality, between physical servers | 37% | 20% |
| Discover VMs | 33% | 12% |
| Manage the traffic that goes between virtual machines (VMs) on a single physical server | 31% | 23% |

As shown in Table 6, The Survey Respondents indicated that getting better at many of the individual challenges associated with server virtualization is important to their organization. In addition, it is reasonable to look at the five challenges contained in Table 6 as being a single challenge - managing server virtualization. When looked at that way, getting better at server virtualization is extremely important to The Survey Respondents. However, The Survey Respondents don't find any of the individual challenges associated with server virtualization as being very difficult. The reality is that some tasks such as dynamically moving VMs, and all of the supporting management functionality, between physical servers, is extremely difficult. In particular, while the task of moving the VM itself is a relatively simple function of the virtual server management system, there can be significant challenges in assuring that the VM's network configuration state (including QoS settings, ACLs, and firewall settings) is also transferred to the new location.

## Management Solutions

The Survey Respondents were asked to indicate the approach their organization takes to management. They were given the following choices and allowed to choose all that applied to their environment.

- We have a focus primarily on individual technology domains such as LAN, WAN and servers
- We have a focus on managing the performance of applications as seen by the end user
- We have a focus on managing the performance of services as seen by the end user, where service refers to multiple, inter-related applications
- Other

Their responses are shown in **Figure 4.**

**Figure 4: Focus of Management**

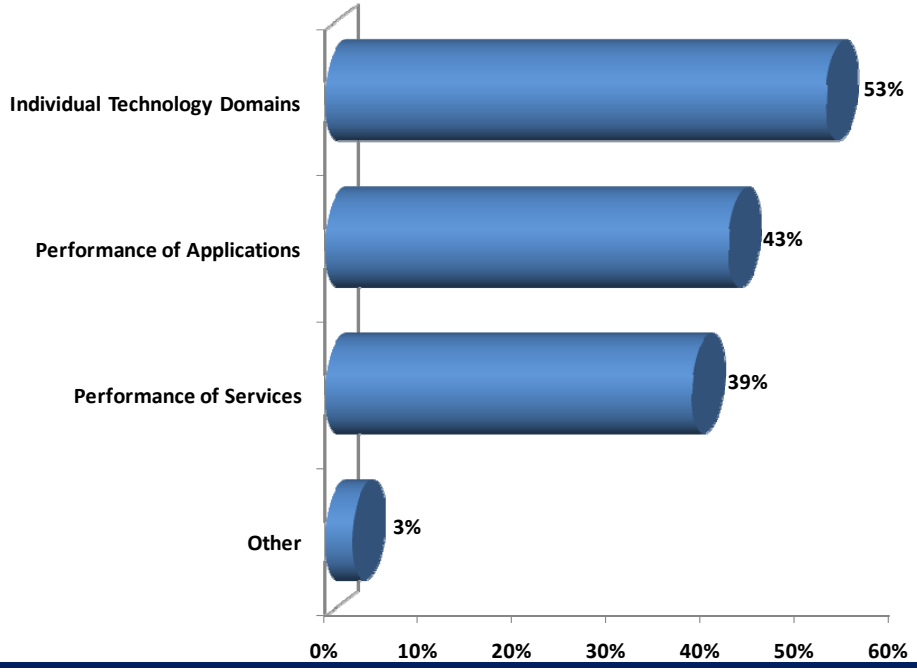


The most common response in the *other category* was that the IT organization had outsourced the management of the network to a third party. The data in Figure 4 indicates that the most frequent approach that IT organizations take to management is to focus on individual technology domains, but that a significant percentage of IT organizations focus their management activities on the performance of applications and/or services.

Roughly half of The Survey Respondents indicated that their organization has a well-understood, effective strategy that describes how they will manage each component of IT. Roughly half of The Survey Respondents also indicated that the effective management of IT is a very important topic for their senior IT managers. It was somewhat of a surprise to see that whether or not an organization's senior IT managers regard effective management as being very important didn't significantly impact whether or not that IT organization had a well-understood, effective management strategy.

Some of the additional insight into the state of network management that was provided by The Survey Respondents is:

- Almost fifty percent of The Survey Respondents indicated that they offer SLAs to their customers for at least some key applications. However, far fewer indicated that they do an effective job of managing those SLAs.
- Just over a third of The Survey Respondents indicated that their company acquires software from a SaaS provider. However, less than half of the companies that utilize a SaaS provider indicated that they do an effective job of managing the software that they acquire.
- Almost three quarters of The Survey Respondents indicated that when acquiring new management tools over the next year that they would buy best of breed solutions vs. buying a large, integrated solution. Companies with 10,000 or more employees were slightly more interested in buying large, integrated solutions than was the overall survey base.

The Survey Respondents were asked to indicate how their organization's 2010 management budget compares to their 2009 management budget. Their answers are shown in **Table 7**.

**Table 7: 2010 Management Budget vs. 2009 Management Budget**

| Reduced by 10% or more | Reduced by less than 10% | About the same | Increased by less than 10% | Increased by more than 10% |
|---|---|---|---|---|
| 10.5% | 9.5% | 51.6% | 16.8% | 11.6% |

The data in Table 7 shows that in 2010 the budget that the vast majority of IT organizations (80%) have for management functionality will either stay the same as it was in 2009 or it will increase.

.

## Summary

In The Application Delivery 1.0 era IT organizations began to develop a concerted effort to protect their company from security attacks, ensure the performance of key applications, offload computationally intensive activities off of servers, and gain some insight into the end-to-end performance of applications. However, 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
- Shifting emphasis and growing sophistication of cyber crime
- Adoption of varying forms of virtualization
- Adoption of cloud computing

The data in tables 3 and 5 drive home the fact that there are a large number of optimization and management tasks that are important for IT organizations to get better at in the next year and that these tasks represent a combination of Application Delivery 1.0 and Application Delivery 2.0 era challenges, including:

- Ensuring acceptable performance for VoIP traffic
- Improving the performance of applications used by mobile workers
- Preventing large scale DDoS attacks
- Effectively managing services, where services are comprised of multiple, inter-related applications

While IT organizations realize that there are a number of optimization and management tasks that they must get better at in the next year, they don't tend to think that getting better at these tasks is terribly challenging. In some cases, such as improving the performance of TCP, they are correct. In other cases, however, such as:

- Ensuring acceptable performance for the applications that are acquired from a SaaS provider
- Dynamically moving VMs, and all of the supporting management functionality, between physical servers

The Survey Respondents seem to be underestimating the difficulty of the task.

Fortunately the budget that IT organizations have in 2010 for implementing both optimization and management functionality is up somewhat over the corresponding budget that IT organizations had in 2009. Relative to the current state of optimization, the vast majority of IT organizations have implemented at least some optimization functionality, but only one third of IT organizations have broadly implemented

optimization functionality within their company. On a going forward basis, IT organizations intend to continue to deploy both appliance-based and software-based WOCs and ADCs. The largest incremental increase in the use of optimization functionality over the next year, however, will be a result of a significant increase in the deployment of client-based software to improve the delivery of applications to mobile workers.

Relative to the current state of management, the most frequent approach that IT organizations take is to focus on individual technology domains. However, a significant percentage of IT organizations focus their management activities on the performance of applications and/or services. IT organizations have been relatively aggressive at implementing internal SLAs for at least some key applications and at acquiring software from a SaaS provider. However, most IT organizations don't do an effective job of managing either their internal SLAs or the software they acquire from a SaaS provider.

On a going forward basis, IT organizations are more likely to acquire best of breed management tools than they are to acquire large, integrated solutions. The most important management task facing IT organizations this year is improving their ability to identify the root cause of degraded application performance. Unfortunately, due to the complications brought about by virtualization and cloud computing, this task is quickly becoming an order of magnitude more difficult than it is in the traditional IT environment.

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

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

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

# Unlocking the WAN Architecture for On-demand Cloud Computing


## Overview

Communications service providers are being called upon by their enterprise customers to help with two seemingly opposing objectives: the expansion of mission-critical WAN services to support global business objectives and cost rationalization. The Alcatel-Lucent Application-Assured (AA) VPN architecture is ideally suited to helping service providers address this requirement. The Alcatel-Lucent next-generation architecture enables service providers to deliver two application-assured business services over the same network infrastructure: VPNs and on-demand cloud computing services, such as Software as a Service (SaaS). This flexible architecture enables service providers to tailor service offerings to each enterprise's specific needs and to grow service offerings as the enterprise's needs change. Leveraging a single architecture to deliver two different services ensures the service provider gets maximum return on its infrastructure investment. Equally important, both services can be delivered using a common management system, which reduces operational overhead and ensures rapid service delivery.

## Two for the price of one

The Alcatel-Lucent AA VPN architecture enables communications service providers to:

- Deliver enhanced application-assured VPN services to enterprises with application-reporting and policy control capabilities. These value-added capabilities enable service providers to differentiate their VPN service offering.
- Deliver application-assured cloud-computing. Cloud computing is of benefit to enterprises who do not wish to take on the high capital investment of software licenses and servers, maintaining 24x7 resources, and so on. They would rather pay a monthly fee for hosted application services to the service/application provider.

Adding on-demand cloud computing into the service mix enables the service provider to strengthen its relationship with enterprise customers and evolve from a connectivity provider to a trusted ICT partner.

## Enterprise challenge equals opportunity

While enterprises are increasingly reliant on their business applications for successful day-to-day operations, they have little visibility of how applications are performing when they are being accessed over a WAN VPN. This limited visibility can be a deterrent to adopting or expanding use of cloud computing offerings.

Service providers have an opportunity to capitalize on this gap by enhancing their existing VPN and offering new on-demand cloud-computing services with the ability to monitor and address application-performance issues. The Alcatel-Lucent AA VPN architecture provides extensive service assurance capabilities, enabling the service provider to expand the reach and value of its business VPN service and to introduce a cloud computing applications portfolio.

## The Solution

The Alcatel-Lucent AA architecture enables service providers to offer application-assured VPNs and cloud computing applications. The architecture is an extension of the Alcatel-Lucent Business VPN Services (BVS) solution, which supports the convergence of IP voice, data and video over Layer 2 and/or Layer 3 business VPNs. Enhancing the BVS solution with application assurance allows service providers to be highly

competitive, leveraging their existing IP/MPLS network and service management infrastructure to offer application-level visibility and policy control with minimal incremental investment.

VPN and on-demand cloud computing services can be delivered over a single, cost-effective Alcatel-Lucent AA architecture.

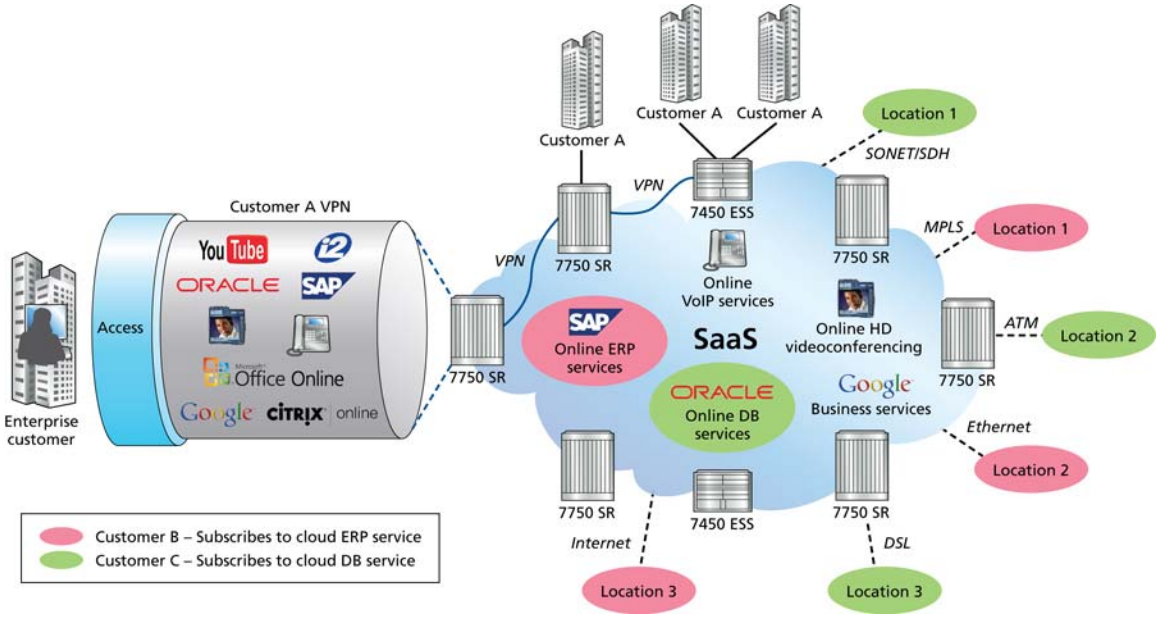


The Alcatel-Lucent AA architecture provides the ability to recognize applications and application flows throughout the network. This in turn enables the service provider to report on the applications and apply application-level QoS controls. The Alcatel-Lucent AA architecture supports multiple small, medium and large enterprise customers cost effectively, providing dramatic cost savings.

## How it works

The Alcatel-Lucent AA architecture relies on the application-assurance feature set of the Alcatel-Lucent Service Router (SR) Operating System (OS), including the purpose-built Alcatel-Lucent Multiservice Integrated Service Adapter (MS-ISA). The Alcatel-Lucent MS-ISA enables stateful application traffic-flow inspection and application assurance for IP/MPLS-based Layer 2 and Layer 3 services. By integrating the application recognition and assurance functions in the Alcatel-Lucent 7750 SR and 7450 Ethernet Service Switch (ESS) situated at the edge of the network, the number of managed network elements and overall power consumption can be dramatically reduced while the services themselves can scale to meet demand.

The Alcatel-Lucent AA architecture is fully managed with the Alcatel-Lucent 5620 Service Aware Manager suite, which includes the Alcatel-Lucent 5670 Reporting and Analysis Manager. This fully integrated network- and service-management solution simplifies operational integration and enables service providers to deliver new application reporting and assurance capabilities to their business services customers quickly.

### *Addressing the application visibility issue*

Service providers can enable web-based service portals to allow enterprises to view application reports for their VPN and cloud computing services on a highly granular basis. Enterprises can use this information to understand their service usage patterns and troubleshoot application-related issues. The solution also allows customers to manage their policies to control applications traversing the VPN, ensuring that resources are

used in accordance with business priorities. Web-based portal provides enterprises critical insight into their applications' performance.

## Reaping application assurance rewards

The Alcatel-Lucent AA architecture allows service providers to introduce a variety of tiered business VPN service plans as well as cloud computing services. Both types of service can be delivered with stringent SLAs for more direct support of enterprises' business objectives and strengthened customer relationships. These enhanced relationships provide increased opportunities for up-selling services to enterprises.

The Alcatel-Lucent AA architecture's cloud-based approach for application assurance delivers many benefits:

- Ability to identify business applications and prioritize them according to their performance behavior to ensure that an enterprise experiences consistent end-to-end application performance
- Better understanding of how an enterprise customer's applications are traversing the WAN, and deeper insight into how to baseline business-critical application traffic to improve performance
- Support for all business VPN and cloud computing applications, regardless of the access type connecting the enterprise site, allowing providers to overcome the limitation of using costly access-specific devices
- Matching of the VPN and cloud computing applications to site-specific application needs, such as volume of use by time and application
- Scalable performance, depending on the service provider's specific infrastructure. For example, a provider with Alcatel-Lucent 7750 SRs and 7450 ESSs can scale up to 70 Gb/s of application processing per node to support hundreds and thousands of enterprise sites.
- Reduced time and investment to operationalize the solution. Integrating application intelligence in the network nodes reduces complexity and eliminates the need for truck rolls to CPE locations. An integrated service management capability also minimizes installation and provisioning times.
- Ability to offer service portals to enterprise customers, allowing them to monitor application statistics, download customized application performance reports, and manage their VPNs and cloud computing applications themselves using application-based policy control.

## Conclusion

The Alcatel-Lucent cloud-based application assurance solution is deployed once by the service provider and can be offered to all enterprise customer sites at minimal cost. Moreover, the application assurance service can be remotely activated in minutes without requiring onsite installation.

These new service capabilities translate directly to increased revenue opportunities, with a faster time-to-market. By entrenching a competitive differentiator in its business service offerings, the service provider is safeguarding future revenue streams from price erosion and reinforcing customer stickiness — something that is extremely valuable in today's highly competitive market.

For the enterprise, these new service capabilities ensure that they are able to expand their mission-critical WAN application services effectively, to support global business objectives and cost rationalization.

**For more information on Alcatel-Lucent's Business VPN Services, please visit**
**http://www.alcatel-lucent.com/businessvpns**

# Managing End-user Experience, Application, and Network Performance to Deliver Business Services

FLUKE network

## Overview

The fundamental challenge of any IT organization today is aligning its technology with the business goals. In order to achieve alignment, IT organizations need to have visibility into how the performance and changes in the infrastructure impact application and business service delivery. This actionable insight is integral to bridging the gap between business goals, customer experience and IT technology.

But, within many enterprises, IT organizations tend to narrowly focus on technology issues and specific project demand without understanding the implications across the entire organization ranging from the end-user to the CEO.

A more comprehensive approach is needed to demonstrate the value of IT in a business context by delivering insight into individual user experience, application and network performance. IT can improve end user productivity through better application and business services delivery by having visibility into how the infrastructure impacts the business. IT equipped with the proper intelligence, can reduce operating costs, make better decisions that impact business services delivery.

A successful strategy for managing the impact of infrastructure on application and business services delivery includes:

- Measuring the impact of infrastructure performance and changes on the business and users
- Identifying performance degradation incidents
- Determining the root cause of degradation
- Resolving the problem

## Impact of the infrastructure on the business

For many organizations, there is a distinct separation between business challenges and operational challenges. While extremely different in their approach, a successful strategy can help align the business and operational requirements to provide a holistic view for managing end user experience and application and network performance.

Let's use, as an example, a national supplier of automotive parts with approximately 40 distribution centers across the country. Since the primary business is the timely delivery of needed parts, the business challenge is supporting a networked inventory management and delivery status system that works across two very different business models: the wholesale distribution to commercial customers, and retail sales to consumers through 3,400 store-fronts.

The operational challenge is identifying the source of application slowdowns and restoring normal service as quickly as possible. The IT team needed to integrate network and applications performance data to assess and validate end-user experiences. This involved both voice components and data components with specific Service Level Agreements (SLAs) for each. To provide an integrated view of infrastructure performance, it

requires actionable intelligence provided by correlating information collected from a variety of instrumentation options across the enterprise LAN, WAN and data center.

## Identify Incidents That Impacts Performance and Business

While the goal of managing application and network performance is to deliver an enterprise-wide solution, individuals in different IT roles have distinct information requirements to do their jobs properly.

What makes application and network performance framework so important is that even though these different groups have different requirements, they're all intertwined so one group can impact the success of all others To illustrate this point, let's revisit the automotive parts supplier example from a different viewpoint.

A problem first revealed itself by a complaint from a retail store-front. The issue might be that the store cannot get inventory information quickly enough to meet the customer's requirements and they are starting to lose business. The IT department first needs to determine the extent of the problem; is it one store or is it company-wide? They need a high level view showing response times to all locations, matched against SLAs and compared to this one store.

Maybe the problem is limited to this single location. If so, what is different about this one location? Was something changed in how this store is communicating with the data center, making it a network problem? Is some backup process for this store incorrectly scheduled and running during business hours, making it a problem for the application group? What if the problem is companywide, because all traffic is routed to one overloaded server? Or did a change to an application work fine in tests but fail when running over the WAN? The entire IT team needs the ability to drill into these types of issues, because most performance problems cross organizational boundaries.

**What is Enterprise Service Intelligence (ESI)? Why is it Important?**

The term Enterprise Service Intelligence is the Fluke Networks Systems vision to help IT professionals and business stake-holders understand the true impact of the IT infrastructure on mission critical applications and business services. The implementation of ESI demonstrates the value of IT in the business context by delivering insight into individual user experience, application and network performance. Find out more, and how you can begin to put ESI into action in your environment, at **www.flukenetworks.com/ESI.**

## Identifying the True Root Cause of Performance Problems

A third aspect of managing application and business service delivery infrastructure performance is having the flexibility and visibility for managing a wide range of deployment requirements. Using the previous example, we can see how necessary it is to be able to have a true sense of application performance as opposed to just looking at the availability of a network infrastructure. Operational requirements can have a substantial impact for an organization implementing a guaranteed-uptime application. Potential impacts include the guaranteed availability service level agreement and troubleshooting capabilities dealing with so many remote locations.

If a holistic application and network performance framework is not followed, organizations will have difficulties when different groups are only focused on their individual goals and don't have the visibility to understand the impact on business-critical resources and to truly identify the root cause of performance problems.

## Increasing the business value of IT

For many CIOs, one key concern is not only increasing the business value of IT, but also quantifying the positive impact. While many organizations view IT as a cost center with a focus on reducing the expenses,

many leading companies view IT as a strategic asset. These organizations focus on how IT can improve overall business value to the organization.

A comprehensive strategy for managing the impact of infrastructure on application and business services delivery allows organizations to focus on the strategic asset and align technology with business goals. With a complete understanding of the impact of infrastructure performance and changes on the business and users, an enterprise can reduce the risk of downtime and degradation, reduce the cost of operations and troubleshooting, and optimize IT support staff.

IT organizations have begun implementing business service dashboards and automating service desk workflow. A business service dashboard provides the line of business owners a clear view of the availability and performance of critical application and services that impact the bottom line. For IT, this dashboard increases the visibility of impact of infrastructure performance and changes on the business and users. An automated service desk work flow improves incident management operational efficiency and improve the accuracy of IT incident impact reporting.

The ideal underlying unified performance management system, supplying the intelligence to the business service dashboard and service desk, needs to be built on an application-aware architecture with the ability to correlate data gathered from a range of instrumentation options covering end user experience monitoring, application, network and VoIP performance across the enterprise LAN, WAN and data center. Identifying and leveraging the right solution will minimize the challenges and limitations presented with the traditional, siloed approach to IT and ultimately help to align them to overall business goals.

## About the author

**Belinda Yung-Rubke**
**Director of Marketing**
**Fluke Networks Systems**

Belinda Yung-Rubke is the Director of Marketing for Fluke Networks Systems. She has over 20 years of experience in the communications and networking industry in product marketing, planning, business development and manufacturing. Prior to joining Fluke Networks Systems, Belinda held a number of management positions at Hewlett-Packard, Agilent Technologies and several start-up companies. Belinda holds a Bachelors of Science Electrical Engineering degree from the University of California at Davis and is based in Colorado Springs, Colorado.

# Meeting the Challenges of Today's Distributed Enterprise by Juniper Networks

## Overview

This document describes some of the most important reasons that more than 4,000 high-performance IT organizations globally have chosen the Juniper Networks® Application Acceleration solution to dramatically improve application performance.

## Challenges

The business goal of today's high-performance enterprises—deliver new, differentiated products and services that wow customers—cannot be met without fundamentally changing the way they interact with customers, contractors, and partners. Therefore, these innovative enterprises have embraced new methods of connecting with customers and collaborating with partners—and seek to enhance their employees' productivity by enabling them to choose their workplace to suit the task.

## Meeting the Challenges

Application acceleration solutions help businesses make more efficient use of their WAN resources and connect all employees, regardless of their location, by delivering LAN-like response times everywhere.

Juniper Networks takes this a step further by integrating all these capabilities—and more —in a single solution to address the bandwidth, latency, contention, and manageability. Working together, these attributes—detailed in **Table 1**—help customers achieve the desired application performance.

**Table 1: Application Acceleration Problems and Requirements**

| Problem | Application Acceleration Requirements |
|---|---|
| Limited WAN bandwidth | Reduce the amount of traffic traversing WAN links through a combination of compression and caching features that make the most efficient use of existing resources |
| WAN latency | Accelerate a broad set of applications that make up the majority of WAN traffic and delivery technologies that compensate for inefficiencies of certain chatty protocols (TCP, MAPI, CIFS, and HTTP) over the WAN |
| Application contention | Deliver application control capabilities such as quality of service (QoS) that allow users to prioritize traffic flows to ensure sufficient bandwidth for critical business operations |
| Management | Provide visibility into application performance so users can understand, anticipate, and predict application behavior to make informed decisions |

## The Juniper Networks WXC Series Application Acceleration Solution

Juniper Networks WXC Series Application Acceleration Platforms are based on a unique framework, which outlines the attributes defined previously and describes how specific features of the WXC Series solutions meet those requirements.

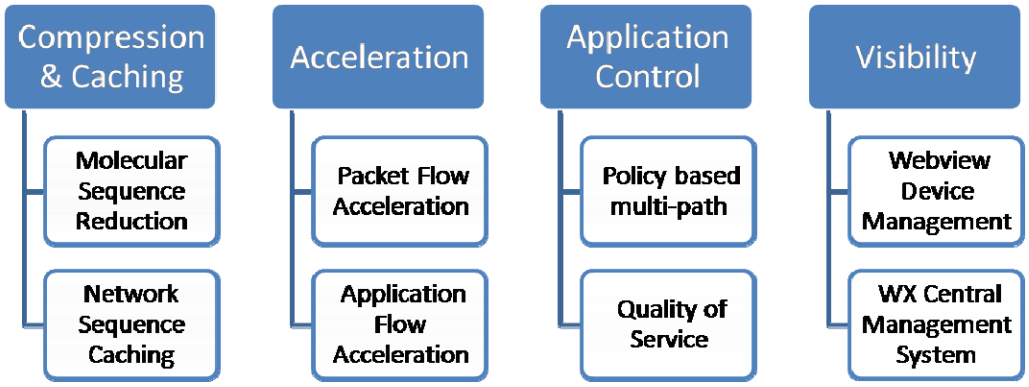


**Compression and Caching:** To satisfy compression and caching components, Juniper Networks integrates memory-based Molecular Sequence Reduction (MRS) compression technology, which increases WAN throughput up to 10 times by eliminating repeated data patterns from traffic flows traversing the WAN. MSR compression is complemented with Network Sequence Caching (NSC), which uses hard disks to recognize and store larger repeated patterns last seen days or even weeks earlier to increase throughput by up to 50 times.

**Acceleration:** Acceleration is delivered in the form of Packet Flow Acceleration (PFA) and Application Flow Acceleration (AppFlow) technologies. The PFA techniques—including Fast Connection Setup and Active Flow Pipelining—combat the effects of TCP latency by accelerating connection setup and substituting a more efficient client transport across the WAN. For lossy networks, an additional PFA feature—forward error correction (FEC)—makes use of recovery packets to reconstruct lost data, eliminating the need for retransmissions.

**Application Control: WXC Series** delivers application control capabilities via QoS and Policy-Based Multipath capabilities. QoS is combined with bandwidth management tools to ensure bandwidth is always available for delay-sensitive applications such as VoIP. Policy-based Multipath keeps designated flows to a specific WAN link when multiple options are available.

**Visibility:** The WXC Series solutions integrate visibility and reporting functions that arm IT with a set of monitoring and management tools for controlling application performance over the WAN. Juniper Networks WX Central Management System provides system-wide visibility into WAN performance while the WebView device management software enables IT to configure and manage individual appliances from a central location.

## Going Beyond Acceleration

### *Junos Pulse*

Junos Pulse is a dynamic, integrated network client. A core component of Junos Platform, Pulse delivers integrated, anytime/anywhere connectivity, acceleration, and security, while drastically simplifying user experience. With Junos Pulse, users no longer need to interact with network access and security software. From any location, users simply supply their credentials and Junos Pulse takes care of the rest.

Junos Pulse enables fast, easy, secure access to corporate networked and cloud-based data and applications from mobile devices and smartphones. Enterprises and service providers can deploy granular role and device-based security policies when provisioning mobile handset and device access. Plus Pulse can increase

smartphone sales, per unit revenues, retention rates, and customer satisfaction for service providers and mobile operators.

### Scalability

What flows from the rigorous and integrated approach is unmatched scalability of Juniper Application Acceleration solutions. Juniper operates arguably the largest WAN optimized networks in the world[8]. This is made possible because Juniper Networks WXC Series Application Acceleration Platforms are built to support mission-critical services that should be deployed pervasively and not limited to select remote sites with low bandwidth, high-latency links.

### Diversity of Product Portfolio

Juniper's diverse portfolio is cited as one of the main reasons independent analysts recommend Juniper Networks should be on your short list. The WXC Series solutions are available as a software client (see Junos Pulse above) for mobile employees, small and medium appliances for branch offices, and high-end systems for data centers. For branch offices with pre-existing Juniper Networks J Series Services Routers, the WXC Series ISM blade is ideal. Additionally, the WX client software is fully compatible with the award-winning Juniper Secure Access (SSL VPN) solutions.

## Conclusion

The optimization speed, ease of deployment, and security integration of Juniper Networks WXC Series Application Acceleration Platforms—along with Juniper's award-winning technical support capabilities—can effectively enhance performance and collaboration of your enterprise.

The ROI of a typical Juniper Application Acceleration deployment is from six months to a year. When making your selection, focus on scalability and stability of the product feature set and choose a vendor with a strong track record of support and operational maturity.

**For further information about the WXC Series Application Acceleration Platforms, technical features, multimedia demos, case studies, and ROI tools, visit http://www.juniper.net/us/en/products-services/application-acceleration/ or call Juniper at 866.298.6428 (in the U.S.) or 978.589.0500 (outside the U.S.).**

[8] Please contact a Juniper sales representative for specifics of the customer case study.