

The 2012 Cloud Networking Report

Part 5: Management & Security

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:



Management & Security	1
Executive Summary.....	1
Management.....	2
A New Set of Management Challenges	3
Management Challenges Associated with Server Virtualization.....	3
Management Challenges Associated with Cloud Computing	5
Importance of Managing Cloud Computing.....	6
The Traditional Management Environment	8
Network Performance Management Systems.....	8
Application Performance Management.....	8
Synthetic Transactions	9
Internal SLAs.....	9
Delay Sensitive Traffic.....	11
The Emerging Management Environment	12
The Evolving Focus on Services	12
Service Delivery Management.....	14
Dynamic Infrastructure Management.....	15
Virtualized Performance and Fault Management.....	16
Converged Infrastructure Management	16
Orchestration and Provisioning.....	18
Application Performance Management	20
Impediments.....	20
A Top Down Approach	21
Root Cause Analysis	21
Designing for Application Performance.....	22
Application Performance Engineering.....	23
Application Performance Management Tools	25
Management as a Cloud Provided Service.....	26
Security	27
The Current Environment for Security Breaches.....	27
The Current Environment for Implementing Security	28
Security as a Cloud Provided Service	32
Web Application Firewall Services.....	34
The Role of a Traditional Firewall	34
The Role of a Web Application Firewall Service	35

Management & Security

Executive Summary

The **2012 Cloud Networking Report** (The Report) will be published both in its entirety and in a serial fashion. This is the fifth of the serial publications. The first publication in the series described the changes that are occurring in terms of how cloud computing is being adopted, with a focus on how those changes are impacting networking. The second publication in the series focused on data center LANs. The third publication discussed Software Defined Networks (SDNs) and included the results of a survey that was done in conjunction with Information Week. The fourth publication focused on Wide Area Networking.

The focus of this publication of The Report is security and management. The Report will also be published in its entirety and there will be a separate executive summary that covers the totality of The Report.

This section of The Report includes the results of surveys that were recently given to the subscribers of Webtorials.com. Throughout this report, the IT professionals who responded to those surveys will be referred to as the ***Survey Respondents***.

Management

One of the questions that were administered to the **Survey Respondents** was “Please indicate how important it is to your organization to get better at each of the following tasks over the next year.” The question included twenty wide-ranging management tasks. The possible answers were to the question were:

- Extremely important
- Very important
- Moderately important
- Slightly important
- Not at all important

In order to avoid restating that question each time it is referenced in this section of The Report, it will be referred to as The Question.

A New Set of Management Challenges

Management Challenges Associated with Server Virtualization

As discussed in the section of The Report entitled *The Emergence of Cloud Computing and Cloud Networking*, one of the key characteristics of a cloud computing solution is virtualization. Server virtualization is the most commonly implemented form of virtualization and it creates a number of management challenges. For example, until recently, IT management was based on the assumption that IT organizations performed tasks such as monitoring, baselining and troubleshooting on a server-by-server basis. Now, given the widespread adoption of server virtualization, the traditional approach to IT management must change to enable management tasks to be performed on a virtual machine (VM)-by-VM basis. Another assumption that underpinned the traditional approach to IT management was that the data center environment was static. For example, it was commonly assumed that an application resided on a given server, or set of servers, for very long periods of time. However, part of the value proposition that is associated with server virtualization is that it is possible to migrate VMs between physical servers, both within the same data center and between disparate data centers. This ability to migrate VMs between physical servers is just one example of the fact that

IT organizations need to adopt an approach to management that is based on the assumption that the components of a service, and the location of those components, can and will change frequently.

Additional management challenges that are associated with server virtualization include:

- **Breakdown of Network Design and Management Tools**
The workload for the operational staff can spiral out of control due to the constant stream of configuration changes that must be made to the static data center network devices in order to support the dynamic provisioning and movement of VMs.
- **Limited VM-to-VM Traffic Visibility**
The first generation of vSwitches doesn't have the same traffic monitoring features as does physical access switches. This limits the IT organization's ability to do security filtering, performance monitoring and troubleshooting within virtualized server domains.
- **Poor Management Scalability**
Many IT organizations have experienced VM proliferation sometimes called VM sprawl. In addition, the normal best practices for virtual server configuration call for creating separate VLANs for the different types of traffic to and from the VMs. The combined proliferation of VMs and VLANs places a significant strain on the manual processes that are traditionally used to manage servers and the supporting infrastructure.
- **Contentious Management of the vSwitch**
Each virtualized server includes at least one software-based vSwitch. This adds yet another layer to the existing data center LAN architecture. It also creates organizational stress and leads to inconsistent policy implementation.

- **Inconsistent Network Policy Enforcement**

Traditional vSwitches lack some of the advanced features that are required to provide a high degree of traffic control and isolation. Even when vSwitches support some of these features, they may not be fully compatible with similar features that are offered by physical access switches. This situation leads to the implementation of inconsistent end-to-end network policies.

- **Multiple Hypervisors**

It is becoming common to find IT organizations using multiple hypervisors, each of which comes with their own management system and their own management interface. In addition, the management functionality provided by each hypervisor varies as does the degree to which each hypervisor management system is integrated with other management systems.

- **Management on a per-VM Basis**

IT organizations typically perform management tasks such as discovery, capacity planning and troubleshooting on a per server basis. While that is still required, IT organizations must also perform those tasks on a per-VM basis.

In order to quantify the interest that IT organizations have in responding to the management challenges that are created by server virtualization, three of the twenty tasks that were included in The Question were:

- Manage the traffic that goes between virtual machines (VMs) on a single physical server.
- Support the movement of VMs between servers in different data centers.
- Perform traditional management tasks such as troubleshooting and performance management on a per VM basis.

The responses of the **Survey Respondents** are summarized in **Table 1**.

Table 1: Importance of Managing Server Virtualization			
	Traffic Between VMs	Move VMs Between Servers	Manage on a per VM Basis
Extremely	6%	10%	11%
Very	27%	28%	34%
Moderately	37%	36%	33%
Slightly	21%	14%	19%
Not at All	9%	13%	3%

One conclusion that can be drawn from the data in **Table 1** is that:

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

Management Challenges Associated with Cloud Computing

Even in the traditional IT environment¹ when the performance of an application is degrading the degradation is typically noticed first by the end user and not by the IT organization. In addition, when IT is made aware of the fact that application performance has degraded, the process to identify the source of the degradation can be lengthy.

Unfortunately:

The adoption of cloud computing makes troubleshooting application performance an order of magnitude more difficult than it is in a traditional environment.

In order to illustrate some of the challenges of managing a cloud computing environment, assume that a hypothetical company called SmartCompany has started down the path of implementing private cloud computing by virtualizing their data center servers. Further assume that one of SmartCompany's most important applications is called BusApp and that the users of the application complain of sporadic poor performance and that BusApp is implemented in a manner such that the web server, the application server and the database server are each running on VMs on separate physical servers which have been virtualized using different hypervisors.

In order to manage BusApp in the type of virtualized environment described above, an IT organization needs detailed information on each of the three VMs that support the application and the communications amongst them. For the sake of example, assume that the IT organization has deployed the tools and processes that are necessary to gather this information and has been able to determine that the reason that BusApp sporadically exhibits poor performance is that the application server occasionally exhibits poor performance. However, just determining that it is the application server that is causing the application to perform badly is not enough. The IT organization also needs to understand why the application server is experiencing sporadic performance problems. The answer to that question might be that other VMs on the same physical server as the application server are sporadically consuming resources needed by the application server and that as a result, the application server occasionally performs poorly.

Part of the challenge associated with troubleshooting this scenario is that as previously noted, in most cases once an IT organization has virtualized its servers it loses insight into the inter-VM traffic that occurs within a physical server. Another part of the challenge is that as was also previously noted, each of the hypervisors comes with their own management system.

Staying with this example, now assume that SmartCompany has decided to evaluate the viability of deploying BusApp using either a public or hybrid cloud computing solution. For the sake of this example, consider two alternative approaches that SmartCompany might implement. Those approaches are:

- 1. Public Cloud Computing**

SmartCompany acquires BusApp functionality from a SaaS provider. The employees of SmartCompany that work in branch and regional offices use an MPLS service from a

¹ This refers to an IT environment prior to the current wave of virtualization and cloud computing.

network service provider (NSP) to access the application, while home office workers and mobile workers use the Internet.

2. Hybrid Cloud Computing

SmartCompany hosts the application and data base servers in one of their data centers and the web servers are provided by a cloud computing service provider. All of the users access the web servers over the Internet and the connectivity between the web server layer and the application server layer is provided by an MPLS service.

In order to monitor and manage either deployment, consistent and extensive management data needs to be gathered from the cloud computing service provider(s), the MPLS provider(s) and the provider(s) of Internet access. In the case of the first option (public cloud computing) similar management data also needs to be gathered on the components of the on-site infrastructure that are used by SmartCompany's employees and supported by the IT organization. In the case of the second option (hybrid cloud computing) similar management data also needs to be gathered on both the on-site infrastructure as well as the web and application servers that are supported by the IT organization. In either case, effective tools are also necessary in order to process all of this data so that IT organizations can identify when the performance of the application is degrading before end users are impacted and can also identify the root cause of that degradation.

A fundamental issue relative to managing either a public or hybrid cloud computing service is that the service has at least three separate management domains: the enterprise, the WAN service provider(s) and the various cloud computing service providers.

The section of The Report entitled *The Emergence of Cloud Computing and Cloud Networking* discussed the advantages of a particular form of hybrid cloud computing: cloud balancing. Until recently IT management was based on the assumption that users of an application accessed that application in one of the enterprise's data centers and that the location of that data center changed very infrequently over time. The adoption of Infrastructure-as-a-Service (IaaS) solutions in general, and the adoption of cloud balancing in particular demonstrates the fact that

IT organizations need to adopt an approach to IT management that is based on gathering management data across myriad data centers, including ones that are owned and operated by a third party.

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

Importance of Managing Cloud Computing

Three of the twenty tasks that were included in The Question were managing private, hybrid and public cloud computing solutions in an end-to-end manner. The responses of the **Survey Respondents** are summarized in **Table 2**.

Table 2: Importance of Managing Cloud Solutions			
	Private Cloud	Hybrid Cloud	Public Cloud
Extremely	16%	11%	9%
Very	25%	25%	19%
Moderately	25%	28%	23%
Slightly	25%	24%	29%
Not at All	10%	13%	19%

One observation that can be drawn from the data in **Table 2** is that

A majority of IT organizations believe that getting better at managing all forms of cloud computing solutions is at least moderately important.

Another observation that can be drawn from the data in **Table 2** is that managing a private cloud is more important than managing a hybrid cloud which is itself more important than managing a public cloud. One of the reasons for this phenomenon is that enterprise IT organizations are making more use of private cloud solutions than they are of either public or hybrid cloud solutions. Another reason for this phenomenon is that as complicated as it is to manage a private cloud, it is notably more doable than is managing either a hybrid or public cloud and IT organizations are placing more emphasis on activities that have a higher chance of success.

The Traditional Management Environment

Network Performance Management Systems

Most Network Performance Management Systems (NPMS) had their origins in monitoring the performance of telecommunication carriers to verify that organizations were getting the services they paid for. These systems are based on a combination of the Simple Network Management Protocol (SNMP) and the Internet Control Message Protocol (ICMP, also known as “ping”). Traditional NPMS measured how long it took a packet to travel from the data center to the branch office network and back - thus determining the Round Trip Time (RTT). If the return packet did not arrive within a few seconds, the original packet was deemed lost and this is how packet loss was measured.

These early NPMS solution worked acceptably well for traditional client/server applications and other centrally hosted applications. However, as technology and applications evolved, the limitations of these systems became apparent. Those limitations include the fact that early NPMS systems:

- Only measured from the central data center to the edge of the branch office network. Problems inside the branch office network went unreported until end users complained
- Had difficulty measuring network paths outside of the data center, such as those used by VoIP, IP video and other peer-to-peer communication traffic
- Measured performance across the entire path, but did not isolate which network segments had performance issues

Application Performance Management

As application architectures evolved from client/server to n-tier web-based applications, application functionality on the server was usually divided up into two or three segments. These segments are the web front-end (presentation tier or tier 1), business logic processes (logic tier or tier 2) and database operations (data tier or tier 3).

In an n-tier web-based application, the user interacts with the presentation tier and the presentation tier in turn communicates to the logic tier, which in turn communicates to the data tier. Each tier uses servers that are optimized to the characteristics of their tier. A presentation tier server, for example, is optimized for network I/O and web traffic, e.g. multiple network cards, large network buffers, etc. A logic tier server is optimized for logic computations, e.g. high-speed CPU, large memory size, etc. A data tier server is optimized for database operations, e.g. multiple disk I/O controllers, large disk cache, large memory size, etc.

Traditional application performance management was typically performed separately from network performance management. For example, when application degradation occurs, the triage process typically assigns the incident to either the network or server areas for resolution. Each area then examines their basic internal measurements of network and server performance and a pronouncement is made that the source of the issue is either the network or the application server or both or neither. Since these tasks are typically done by different parts of

the IT organization using different tool sets and management frameworks, it is quite possible that conflicting answers are given for the source of application performance issues.

Similar to traditional NPMS, traditional application performance management solutions have limitations. Those limitations include the fact that that traditional application performance management solutions:

- Only describe the performance within a single server, not the combined performance across all tiers of an application.
- Cannot attribute CPU, disk I/O, network I/O nor memory utilization to specific classes of transactions. Only aggregate server performance information is available.
- Do not integrate network performance data between tiers to monitor and analyze application performance problems.

Synthetic Transactions

Synthetic transactions provide a somewhat more realistic measurement of application performance than traditional NPMS and application performance management solutions. While synthetic transactions have the advantage of being a better representation of the end user's experience, they also have several disadvantages, including:

- The application being monitored has to be constructed to allow transactions that have no business impact. For example, a banking application would have to have a special account so that when money was added or subtracted from this special account, it would not count towards the banks total assets.
- Synthetic transactions frequently originate from the same data center in which the application servers reside and are not subject to the typical network latencies and availabilities that are present in branch office networks.
- Frequently exercising a synthetic transaction can cause the transaction to perform notably differently than a real production transaction would. For example, a frequently exercised transaction may have its related data in cache all the time and not loaded from disk. As a result, the synthetic transaction would occur notably quicker than a production transaction would.

Internal SLAs

As recently as two or three years ago, few IT organizations offered an SLA to the company's business and functional managers; a.k.a., an internal SLA. However, that situation has changed and now it is common for IT organizations to offer internal SLAs. To understand the prevalence and effectiveness of internal SLAs, The **Survey Respondents** were asked to indicate their agreement or disagreement with three statements. The three statements and the percentage of the **Survey Respondents** that agreed with the statement are shown in **Table 3**.

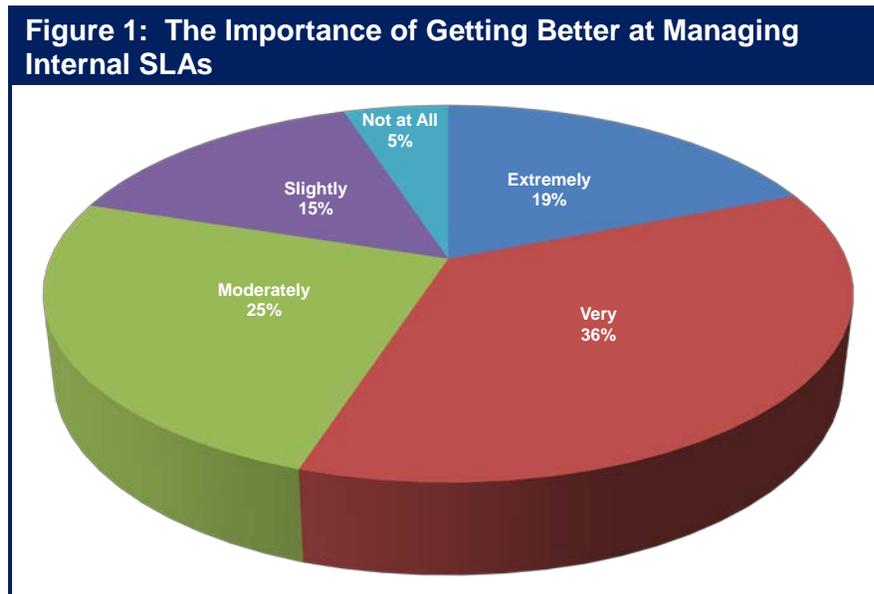
The data in **Table 3** highlights the fact that:

The vast majority of IT organizations provide an internal SLA for at least some applications, but that only half of all IT organizations are successful managing those SLAs.

Table 3: Status of Internal SLAs	
Statement	Percentage
We provide an SLA internally for every application that we support	30.0%
We provide an SLA internally for at least some applications	69.9%
We do a good job of managing our internal SLAs	55.8%

One of the answers to The Question was “managing internal SLAs for one or more business-critical applications”. The responses of the **Survey Respondents** are summarized in **Figure 1**.

The data in **Figure 1** leads to the conclusion that:



Two thirds of IT organizations believe that it is either very or extremely important to get better at effectively managing internal SLAs.

The conclusion stated above is a direct result of the importance of internal SLAs combined with the difficulty that IT organizations currently have with successfully managing those SLAs.

Unfortunately, the movement to utilize public cloud computing services greatly increases the difficulty associated with managing an internal SLA. That follows in part because as discussed previously in this section of The Report, the adoption of cloud computing in general and of virtualization in particular, creates significant management challenges. It also follows in part because it is common for Cloud Computing Service Providers (CCSPs) to deliver their services over the Internet and no vendor will provide an end-to-end performance guarantee for services and applications that are delivered over the Internet.

The lack of meaningful SLAs for public cloud services is a deterrent to the Global 2000 adopting these services for delay-sensitive, business-critical applications.

Delay Sensitive Traffic

Over the last few years the majority of IT organizations have adopted VoIP and video, which are examples of applications that have high visibility and which are very sensitive to transmission impairments. To identify the emphasis that IT organizations place on managing this type of traffic, the **Survey Respondents** were asked to indicate how important it was over the next year for their IT organization to get better at ensuring acceptable VoIP quality. Their answers are shown in **Table 4**.

	Percentage
Extremely Important	14%
Very Important	32%
Moderately Important	32%
Slightly Important	15%
Not at all Important	8%

The data in **Table 4** shows that almost 50% of the **Survey Respondents** indicated that getting better at managing VoIP quality is either very or extremely important to their IT organization.

In the traditional approach to IT management, one set of tools is used to manage enterprise data applications and a different set of tools is used to manage voice and video traffic. That approach is expensive and leads to a further hardening of the technology domains that often exist within an IT organization, which then leads to a lengthening of the time it takes to resolve problems. The reality for most IT organizations is that voice and video traffic is becoming an increasing percentage of the overall traffic on their networks. This reality is one of the reasons why

IT organizations need to adopt an approach to management in which one set of tools is used to manage enterprise data applications as well as voice, video and complex interrelated applications.

As part of the traditional approach to IT management, it is common to use network performance measurements such as delay, jitter and packet loss as a surrogate for the performance of applications and services. A more effective approach is to focus on aspects of the communications that are more closely aligned with ensuring acceptable application and service delivery. This includes looking at the application payload and measuring the quality of the voice and video communications. In the case of unified communications (UC), it also means monitoring the signaling between the components of the UC solutions.

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

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

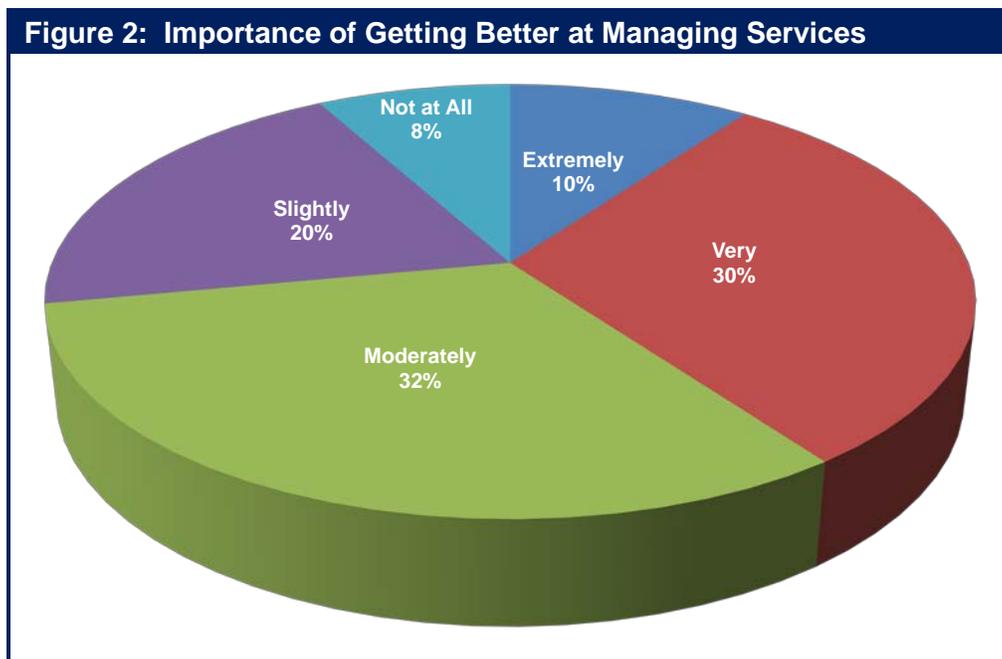
The Emerging Management Environment

The Evolving Focus on Services

Over the last five to ten years, IT organizations have placed a growing emphasis on managing applications in addition to the components of the IT infrastructure that support those applications. While this is still a critical task, IT organizations are coming under increasing pressure to manage not just an individual application such as email, but also a set of interrelated applications (e.g., product lifecycle management, sales order processing, supply chain management, financials and decision support systems) that comprise a business process such as Enterprise Resource Planning (ERP). In order to successfully respond to this pressure, IT organizations need to adopt an approach to service management that enables them to holistically manage the four primary components of a service:

- A multi-tier application and / or multiple applications
- Supporting protocols
- Enabling network services, e.g., DNS, DHCP
- The end-to-end network

To quantify this shift in thinking on the part of IT organizations, the **Survey Respondents** were asked to indicate how important it was over the next year for their organization to get better at managing a business service, such as ERP, that is supported by multiple, interrelated applications. Their responses are shown in **Figure 2**.



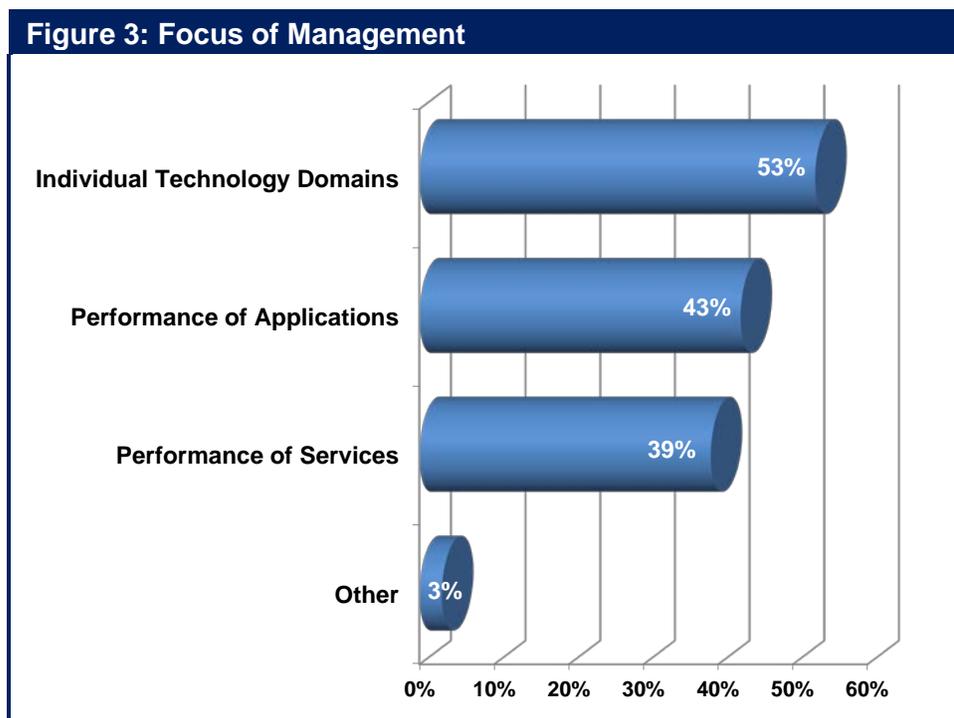
The fact that a significant majority of the **Survey Respondents** indicated that it is at least moderately important for their IT organization to get better at managing a service underscores the fundamental transformation that is underway whereby IT organizations place increasing emphasis on managing services. However, similar to the situation with managing internal SLAs,

the adoption of cloud computing will further complicate the task of managing the inter-related applications that comprise a service. As was the case with SLAs, that follows because the adoption of cloud computing in general and of virtualization in particular, creates significant management challenges.

Another way to measure this transformation is to identify how IT organizations currently focus their management efforts. To that end, the **Survey Respondents** were asked to indicate the approach their organization takes to service or performance 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, in which service refers to multiple, interrelated applications
- Other

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



The data in **Figure 3** indicates that the most frequent approach that IT organizations take to management 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.

Service Delivery Management

In order to respond to the previously described management challenges and to also overcome the limitations of traditional approaches to management, IT organizations must build on the growing emphasis of the last five to ten years to focus on managing application delivery and must establish a more top-down view of the applications that are being delivered. However, they must also broaden this view to include not just managing the delivery of individual applications, but managing the delivery of services as previously defined. In addition, in order to overcome the drawbacks that are associated with the traditional approaches to application performance management

IT organizations should adopt an approach to service delivery management that is unified across the various IT domains so that IT organizations have visibility across all of the applications, services, locations, end users and devices.

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

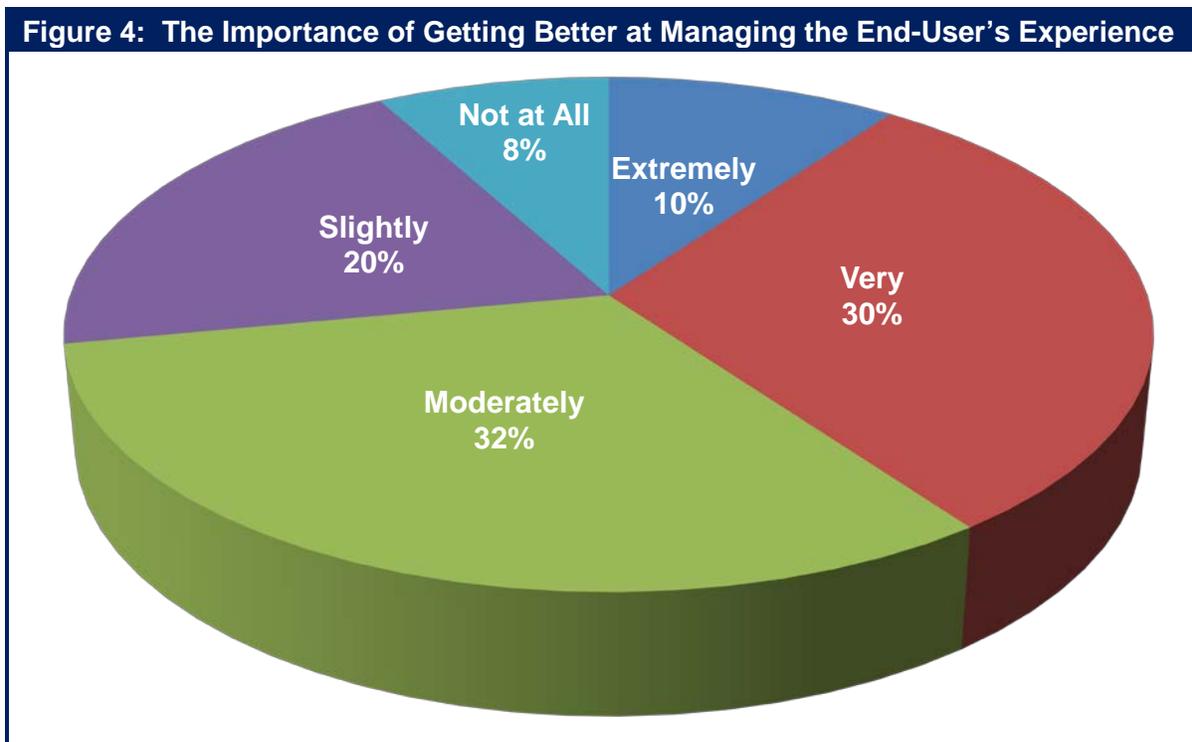
IT organizations must also be able to provide a common and consistent view of both the network and the applications that ride on the network to get to a service-oriented perspective. The level of granularity provided needs to vary based on the requirements of the person viewing the performance of the service or the network. For example, a business unit manager typically wants a view of a service that is different than the view wanted by the director of operations, and that view is often different than the view wanted by a network engineer.

One of the reasons why it is important to get better at managing the end-user's experience was highlighted in The 2012 Application and Service Delivery Handbook². That handbook presented market research that highlighted the fact that in spite of all of the effort that has gone into implementing IT management to date, that it is the end user, and not the IT organization who typically is the first to notice when the performance of an application begins to degrade.

The data in **Figure 4** demonstrates the growing importance that IT organizations place on managing end-user experience. That figure shows the results of a question in which the

² <http://www.webtutorials.com/content/2012/08/2012-application-service-delivery-handbook-2.html>

Survey Respondents were asked how important it was over the next year for their organization to get better at monitoring the end-user's experience and behavior. As shown in **Figure 4**, getting better at managing end-user's experience is either very or extremely important to roughly half of all IT organizations.



Dynamic Infrastructure Management

A traditional environment can benefit from implementing dynamic infrastructure management. However, due to the challenges that are associated with cloud computing:

A dynamic virtualized environment can benefit greatly from a highly scalable and integrated DNS/DHCP/IPAM solution, which is also well integrated with the virtual server management system.

Where DNS/DHCP/IPAM share a common database, the integration obviates the need to coordinate records in different locations and allows these core services to accommodate any different addressing and naming requirements of physical and virtual servers. Potential advantages of this approach include the automated generation of IP addresses for newly created VMs, the automated allocation of subnets for new VLANs, and the population of an IP address database with detailed information about the current location and security profiles of VMs. The integration of infrastructure utilities with the virtual server management system can also facilitate automated changes to the DHCP and DNS databases.

Virtualized Performance and Fault Management

In a traditional IT environment it is common to implement adaptive performance thresholding solutions that can identify systemic deviations from normal patterns of behaviour as well as time over threshold violations and can also automatically update thresholds based on changes to historic levels of utilization. As previously discussed, that same capability is needed in a virtualized environment so that IT organizations can monitor the performance of individual VMs.

Virtual switches currently being introduced into the market can export traffic flow data to external collectors in order to provide some visibility into the network flows between and among the VMs in the same physical machine. Performance management products are currently beginning to leverage this capability by collecting and analysing intra-VM traffic data. Another approach to monitoring and troubleshooting intra-VM traffic is to deploy a virtual performance management appliance or probe within the virtualized server. This approach has the advantage of potentially extending the fault and performance management solution from the physical network into the virtual network by capturing VM traffic at the packet level, as well as the flow level.

While changes in the virtual topology can be gleaned from flow analysis, a third approach to managing a virtualised server is to access the data in the server's management system. Gathering data from this source can also provide IT organizations with access to additional performance information for specific VMs, such as CPU utilization and memory utilization.

Converged Infrastructure Management

An increasingly popular approach to building cloud data centers is based on pre-integrated and certified infrastructure packages from a broadly-based IT equipment vendor, a group of partners or a joint venture formed by a group of complementary vendors. These packages typically are offered as turn-key solutions and include compute, server virtualization, storage, network, and management capabilities. Other data center functions such as WOCs, ADCs, application performance management and security functionality may also be included.

One of the primary reasons why IT organizations implement a converged IT infrastructure is to reduce the overall complexity of a pervasively virtualized infrastructure. The reduction in complexity makes it feasible for IT organizations to fully capitalize on the virtualized infrastructure's inherent potential to serve as an agile, demand-driven platform that can deliver dynamic IT services with unprecedented levels of control, security and compliance, reliability, and efficiency. In order to realize the full potential of the converged IT infrastructure, the management system must provide a unified, cross-domain approach to automated element management, provisioning, change management and operations management. Some of the most critical aspects of managing a cloud data center include:

- **Integrated and Automated Infrastructure and Service Management:** Integrated management reduces the number of management interfaces that are involved in implementing administrative workflows. Automation allows services to be dynamically provisioned, modified or scaled without requiring time-consuming manual configuration across the various technology domains of the data center; e.g., compute, network, storage and security. The management suite should also include the application and service level management capabilities that will support end-to-end SLAs. From an operational management perspective, the management system should provide

additional capabilities, such as cross-domain root cause analysis and service impact analysis, to support the highest levels of service reliability.

- **Secure Multi-tenancy:** A robust multi-layer security architecture is required to ensure confidentiality and integrity of the services and the subscriber's data, particularly in a multi-tenant environment.
- **Support for Enterprise Co-Management:** The service management system should provide a web portal supporting the self-service provisioning of new services or the scaling of existing services. The portal should also include dashboards that provide real-time visibility of application and service performance as well as the consumption of on-demand services. The service management system should also facilitate turning off resources such as VMs that are acquired from a CCSP when they are not needed so that the company using the resources does not incur unnecessary expenses.
- **Compatibility with Enterprise Cloud Implementations:** The efficiency of hybrid clouds is optimized where there is a high degree of consistency across the private and public portions of the solution in terms of the cloud management systems, the hypervisors and the hypervisors' management systems. This consistency facilitates the movement of VMs between enterprise data centers and service provider data centers, and this movement also enables the dynamic reallocation of cloud resources.

Management systems for converged infrastructure typically support APIs for integration with other management systems that may be currently deployed in order to manage the end-to-end data center. These APIs can provide integration with enterprise management systems, automated service provisioning systems, fault and performance management systems and orchestration engines.

While IT departments or CCSPs can themselves achieve some degree of cross-domain management integration by leveraging available element manager plug-ins and APIs, ad hoc automation and integration across the end-to-end infrastructure is quite time-consuming and involves considerable specialized programming expertise. Therefore, the completeness and effectiveness of pre-integrated management functionality are likely to be two of the key differentiators among converged infrastructure solutions.

Cross-domain integrated management of the converged infrastructure will bring added benefits in those situations in which a single administrator has the authority to initiate and complete cross-domain tasks, such as provisioning and modifying infrastructure services. The use of a single administrator can eliminate the considerable delays that are typical in a traditional management environment in which the originating administrator must request other administrators in the other domains to synchronize the configuration of elements within their domains of responsibility. However, a well-known cliché describes the difficulty of realizing these benefits.

Culture eats strategy for breakfast.

That cliché refers to the fact that in many cases the culture of an IT organization resists any changes that involve changing the roles of the members of the organization. Exacerbating the challenge of the IT organization's resistance to change is the fact that, as was pointed out in the

section of this report entitled *The Emergence of Cloud Computing and Cloud Networking*, the culture of an IT organization typically changes very slowly.

Orchestration and Provisioning

Service orchestration is an operational technique that helps IT organizations automate many of the manual tasks that are involved in provisioning and controlling the capacity of dynamic virtualized services. Orchestration engines are available as standalone management products or as part of complete suites of management tools that are focused on the data center. In addition, the management systems that are integrated with converged infrastructure solutions typically include some orchestration capabilities.

By automatically coordinating provisioning and resource reuse across servers, storage, and networks, service orchestration can help IT organizations streamline operational workloads and overcome technology and organizational silos and boundaries. The value proposition of an orchestration engine is that

Orchestration engines use business policies to define a virtual service and to translate that service into the required physical and virtual resources that are needed for deployment.

The orchestration engine then disseminates the needed configuration commands to the appropriate devices across the network in order to initiate the requested service. The orchestration engine can automatically initiate the creation of the required virtual machines while simultaneously deploying the network access and security models across all of the required infrastructure components. This includes routers, switches, security devices and core infrastructure services. The entire process can allow for the setup and deployment of network routes, VPNs, VLANs, ACLs, security certificates, firewall rules and DNS entries without any time consuming manual entries via device-specific management systems or CLIs.

Orchestration engines are available that are pre-configured to interface with certain families of infrastructure devices. Therefore, it is possible to think of the orchestration engine as providing some degree of management integration for non-converged infrastructure. As such, orchestration engines might be a highly desirable approach in those instances in which an existing heterogeneous (i.e., non-converged) data center infrastructure is being transitioned to perform as a cloud data center.

Orchestration solutions would benefit greatly from the emergence of an open standard for the exchange of information among the full range of devices that may be used to construct a dynamic virtual data center. In the Cloud Computing arena there are a number of standards under development, including the Open Cloud Computing Interface (OCCI) from the Open Grid Forum³. These standards activities may also provide value within the enterprise virtual data center, since the stated scope of the specification is to encompass “all high level functionality required for the life-cycle management of virtual machines (or workloads) running on virtualization technologies (or containers) supporting service elasticity”.

³ <http://www.gridforum.org/>

IF-MAP is another emerging standard proposed by the Trusted Computing Group⁴ and implemented by a number of companies in the security and network industries. It is a publish/subscribe protocol that allows hosts to lookup meta-data and to subscribe to service or host-specific event notifications. IF-MAP can enable auto-discovery and self-assembly (or re-assembly) of the network architecture. As such, IF-MAP has the potential to support the automation and dynamic orchestration of not only security systems, but also other elements of the virtual data center. For example, IF-MAP could facilitate the automation of the processes associated with virtual machine provisioning and deployment by publishing all of the necessary policy and state information to an IF-MAP database that is accessible by all other elements of the extended data center.

⁴ <http://www.trustedcomputinggroup.org/>

Application Performance Management

Impediments

Application performance management has been deployed for several years and yet only a small percentage of the **Survey Respondents** indicated that their organization did a good job of managing application performance. To understand why IT organizations are not more successful with application performance management, the **Survey Respondents** were asked to indicate the two primary impediments to their organization being more successful with application performance management. The impediments and the percentage of the **Survey Respondents** who indicated that the impediment was one of the two primary impediments to successful application performance management are shown in **Table 5**.

Impediment	Percentage of the Survey Respondents
Our organization tends to be more reactive than proactive	33%
We focus too much on managing technology domains and not enough on managing business transactions	32%
The tools we use don't give us an end-to-end view of the user's experience	29%
The various sub-groups within the IT organization don't work effectively to identify and resolve problems	26%
We don't have the ability to manage the performance of applications and services acquired from cloud service providers	17%
The tools we use don't allow us to perform rapid root cause analysis	14%
The tools we use don't give us the ability to link the performance of a transaction as seen by the user with all of the various applications that comprise that application	13%
The tools we use don't give us the ability to link the performance of a transaction as seen by the user with the components of the infrastructure that support those transactions	13%
We don't have the ability to gather management data across both the physical and the virtual components of the infrastructure	12%
Other	10%

One observation that can be drawn from the data in **Table 5** is that there isn't a single impediment that is the primary reason why IT organizations aren't successful with application performance management. Rather, there is a wide range of impediments that limit the ability of

IT organizations to be successful with application performance management. Another observation is that

Organizational impediments are more likely to limit an IT organization's success with application performance management than are technical impediments.

A Top Down Approach

The subsection of The Report entitled "Service Delivery Management" discussed the importance of having an approach to managing that is unified across all of the various IT domains. In spite of the importance of having a holistic approach to management in general and to application performance management in particular, only about 15% of the **Survey Respondents** indicated that their organization's approach to application performance management was both top down and tightly coordinated.

Only a small minority of IT organizations has a top down, tightly coordinated approach to application performance management.

As part of an effective approach to application performance management, the automated generation of performance dashboards and historical reports allows both IT and business managers to gain insight into SLA compliance and performance trends. The insight that can be gleaned from these dashboards and reports can be used to enhance the way that IT supports key business processes, help the IT organization to perform better capacity and budget planning, and identify where the adoption of new technologies can further improve the optimization, control and management of application and service performance. Ideally, the dashboard is a single pane of glass that can be customized to suit different management roles; e.g., the individual contributors in the Network Operations Center, senior IT management as well as senior business management.

Root Cause Analysis

As previously mentioned, one of the questions (The Question) that was administered to the **Survey Respondents** was "Please indicate how important it is to your organization to get better at each of the following tasks over the next year." The question included twenty wide-ranging management tasks. **Table 6** lists the three management tasks that were the most important to the **Survey Respondents** and the percentage of the **Survey Respondents** that indicated that getting better at those tasks was either very or extremely important.

Table 6: Primary Management Challenges

Management Task	Percentage
Rapidly identify the root cause of degraded application performance	68%
Identify the components of the IT infrastructure that support the company's critical business applications	63%
Obtain performance indicator metrics and granular data that can be used to detect and eliminate impending problems	52%

It is not surprising that rapidly identifying the root cause of degraded application performance is so important to IT organizations in part because on an ever increasing basis a company's key business processes rely on a handful of applications. That means that if those applications are not running well, neither are those key business processes.

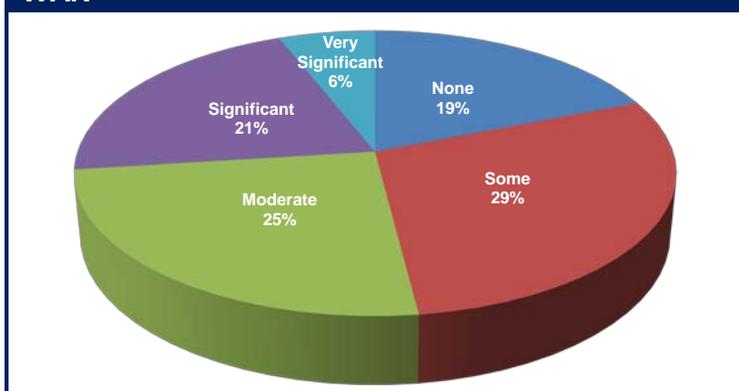
As alluded to in the preceding section of The Report, a prerequisite to being able to perform effective root cause analysis is the automatic discovery of all the elements in the IT infrastructure that support each service or application. That explains why the **Survey Respondents** indicated that this is the second most important management task. For example, if IT organizations can effectively identify which components of the infrastructure support a particular application or service, monitoring can much more easily identify when services are about to degrade due to problems in the infrastructure. As part of this approach, predictive techniques such as heuristic-based trending of software issues and infrastructure key performance indicators can be employed to identify and alert management of problems before they impact end users – a task that the **Survey Respondents** indicated that this is the third most important management task.

In addition, if the IT organization can identify which elements of the IT infrastructure support each service and application, outages and other incidents that generate alerts can be prioritized based on their potential business impact. Prioritization can be based on a number of factors including the affected business process and its value to the enterprise, the identity and number of users affected and the severity of the issue. Another benefit of this approach is that once the components of the infrastructure that support a given application or service has been identified, triage and root cause analysis can be applied at both the application and the infrastructure levels. When applied directly to applications, triage and root cause analysis can identify application issues such as the depletion of threads and pooled resources, memory leaks or internal failures within a Java server or .NET server. At the infrastructure level, root cause analysis can determine the subsystem within the component that is causing the problem.

Designing for Application Performance

One of the traditional challenges to effective application performance is that most IT organizations don't place much emphasis on application performance during application development. To quantify that phenomenon, the **Survey Respondents** were asked "When your IT organization is in the process of either developing or acquiring an application, how much attention does it pay to how well that application will perform over the WAN?" Their answers are shown in **Figure 5**.

Figure 5: The Emphasis on Performance over the WAN



The data in **Figure 5** shows that almost three quarters of all IT organizations place at most moderate emphasis on performance while either developing or acquiring an application.

The lack of emphasis on an application's performance over the WAN during application development often results in the development and implementation of applications that run poorly once they are placed into production. One of the reasons for that phenomenon is that due to factors such as chatty protocols (**Figure 6**), an application can run well over a high-speed, low latency LAN in a development environment but run poorly over a relatively low-speed, high latency WAN in a production environment.

Figure 6: Chatty Protocol



To exemplify the impact of a chatty protocol or application, let's assume that a given transaction requires 200 application turns. Further assume that the latency on the LAN on which the application was developed was 5 milliseconds, but that the round trip delay of the WAN on which the application will be deployed is 100 milliseconds. For simplicity, the delay associated with the data transfer will be ignored and only the delay associated with the application turns will be calculated. In this case, the delay over the LAN is 1 second, which is generally not noticeable. However, the delay over the WAN is 20 seconds. The best case is that a delay of this magnitude results in very unhappy users. In the worst case, it results in the application not being usable. In either instance, the IT organizations will have to devote significant additional time and resources to improving the performance of the application.

Application Performance Engineering

Ideally the issue of application performance would be addressed at all stages of an application's lifecycle, including multiple iterations through the design/implement/test/operate phases as the application versions are evolved to meet changing requirements. However, the vast majority of IT organizations don't have any insight into the performance of an application until after the application is fully developed and deployed. In addition, the vast majority of IT organizations have little to no insight into how a change in the infrastructure, such as implementing server virtualization, will impact application performance prior to implementing the change.

Application Performance Engineering (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, application performance management is used to monitor, diagnose, and report on application performance. Application performance management and APE are therefore highly complementary disciplines. For example, once an application performance management 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. Possible solutions include modifying the application code or improving application performance by making changes in the supporting infrastructure, such as implementing more highly performing

servers or deploying WAN Optimization Controllers (WOCs). Throughout this section of The Report, implementing products such as WOCs will be referred to as a Network and Application Optimization (NAO) solution. Independent of which remedial option the IT organization takes, the goal of APE can be realized – performance bottlenecks are identified, root causes are determined, alternative remedies are analyzed and bottlenecks are eliminated.

An IT organization could decide to ignore APE and just implement NAO in a reactive fashion in an attempt to eliminate the sources of the degraded application performance. Since this approach is based on the faulty assumption that NAO will resolve all performance problems, this approach is risky. This approach also tends to alienate the company's business unit managers whose business processes are negatively impacted by the degraded application performance that isn't resolved until either WOCs are successfully deployed or some other solution is found. A more effective approach was described in the preceding paragraph. This approach calls for NAO to be a key component of APE – giving IT organizations another option to proactively eliminate performance problems before they impact key business processes.

The key components of APE are described below. The components are not typically performed in a sequential fashion, but in an iterative fashion. For example, as a result of performing testing and analysis, an IT organization may negotiate with the company's business unit managers to relax the previously established performance objectives.

- **Setting Performance Objectives**
This involves establishing metrics for objectives such as user response time, transaction completion time and throughput. A complex application or service, such as unified communications, is comprised of several modules and typically different objectives need to be established for each module.
- **Discovery**
Performance modeling and testing should be based on discovering and gaining a full understanding of the topology and other characteristics of the production network.
- **Performance Modeling**
APE modeling focuses on creating the specific usage scenarios to be tested as well as on identifying the performance objectives for each scenario. A secondary focus is to identify the maximum utilization of IT resources (e.g., CPU, memory, disk I/O) and the metrics that need to be collected when running the tests.
- **Performance Testing and Analysis**
Test tools can be configured to mimic the production network and supporting infrastructure, as well as to simulate user demand. Using this test environment, the current design of the application can be tested in each of the usage scenarios against the various performance objectives. The ultimate test, however, is measured performance in the actual production network or in a test environment that very closely mimics the actual production environment.
- **Optimization**
Optimization is achieved by identifying design alternatives that could improve the performance of the application and by redoing the performance testing and analysis to quantify the impact of the design alternatives. In conjunction with the testing, an ROI

analysis can be performed to facilitate cross-discipline discussion of the tradeoffs between business objectives, performance objectives, and cost. This component of APE is one of the key ways that APE enables an IT organization to build better relationships with the company's business unit managers.

Application Performance Management Tools

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

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

The ultimate goal of application performance management is have a single screen that integrates the information from all of the tools in both categories. The idea being that a dashboard on the screen would indicate when user response time or transaction completion time begins to degrade. Then, within a few clicks, the administrator could determine which component of the infrastructure was causing the degradation and could also determine why that component of the infrastructure was causing degradation; e.g., high CPU utilization on a router.

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

Management as a Cloud Provided Service

As pointed out in the section of The Report entitled *The Emergence of Cloud Computing and Cloud Networking*, a new class of solutions has begun to be offered by CCSPs. These are solutions that have historically been provided by the IT infrastructure group itself and include VoIP, network management, security, network and application optimization, application performance management, Unified Communications (UC) and virtualized desktops. This new class of solutions is referred to as [Cloud Networking Services](#) (CNS). That section of The Report also presented the results of a survey in which The **Survey Respondents** were asked to indicate the CNSs that their organization currently acquires from a CCSP and the CNSs that they would like acquire from a CCSP in the next year. Their responses are shown in **Table 7**.

Table 7: Current and Planned Adoption of CNSs		N = 142
	Currently Acquire	Will Likely Acquire
VoIP	20.4%	17.6%
Network Management	19.7%	8.5%
Security	18.3%	9.9%
Unified Communications	15.5%	23.2%
Application Performance Management	10.6%	10.6%
Network and Application Optimization	8.5%	9.2%
Virtual Desktops	7.0%	19.0%

The data in **Table 7** shows that

IT organizations have a significant interest in acquiring network management functionality for a cloud service provider.

In the current environment it is possible to find a CNS that provides almost any possible form of management capability. For example, one class of management based CNS is focused on managing specific types of devices, such as branch office routers, WiFi access points, mobile devices or security devices. In some cases, the CNS supports customer-owned CPE from a wide range of vendors. In other cases, the CNS could be bundled with CCSP-owned devices located at the customer's premise. A variation on the latter approach involves a CNS vendor that provides devices, such as branch office routers, that have been specifically designed to be centrally managed from the cloud via a web portal. In this case, the vendor can move the device's control plane into the cloud in a manner analogous to the separation of control plane and data plane provided by OpenFlow, as discussed in the section of this report entitled *The Emerging Data Center LAN*.

A second class of management based CNS is focused on managing other CNS services provided by a CCSP. These services typically are aimed at addressing the weaknesses in management capability generally associated with early CCSP provided services. For example, the initial wave of CCSP services came with little if any commitment on the part of the service provider relative to an SLA. One example of this class of management based service is a CNS that provides an enhanced level of management for a VoIP service that an IT organization acquires from a CCSP.

Security

The Current Environment for Security Breaches

The security landscape has changed dramatically in the last few years. In the very recent past, the typical security hacker worked alone, relied on un-sophisticated techniques such as dumpster diving, and was typically motivated by the desire to read about their hack in the trade press. In the current environment, sophisticated cyber criminals have access to malware networks and R&D labs and can use these resources to launch attacks whose goal is usually to make money for the attacker. National governments and politically active hackers (hacktivists) are engaging in cyber warfare for a variety of politically motivated reasons.

Over the last few years, the sophistication of hackers has increased by an order of magnitude.

The shift in the security landscape has been documented in a number of reports. For example, IBM's X-Force 2011 Trend and Risk Report⁵ made a number of observations relative to the current environment for security breaches. Some of the key observations made in that report are:

- **Mobile Devices**

The IBM report stated that in 2011 there was a 19 percent increase over 2010 in the number of exploits publicly released that can be used to target mobile devices such as those that are associated with the BYOD movement. The report added that there are many mobile devices in consumers' hands that have unpatched vulnerabilities to publicly released exploits, creating an opportunity for attackers.

- **Social Media**

With the widespread adoption of social media platforms and social technologies, this area has become a target of attacker activity. The IBM report commented on a surge in phishing emails impersonating social media sites and added that the amount of information people are offering in social networks about their personal and professional lives has begun to play a role in pre-attack intelligence gathering for the infiltration of public and private sector computing networks.

- **Cloud Computing**

The IBM report stated that there were many high profile cloud breaches affecting well-known organizations and large populations of their customers. IBM recommended that IT security staff should carefully consider which workloads are sent to third-party cloud providers and what should be kept in-house due to the sensitivity of data. The IBM X-Force report also noted that the most effective means for managing security in the cloud may be through Service Level Agreements (SLAs) and that IT organizations should pay careful consideration should be given to ownership, access management, governance and termination when crafting SLAs.

⁵ [X-Force 2011 Trend and Risk Report](#)

Blue Coat Systems' 2012 Web Security Report⁶ also made a number of observations relative to the current environment for security breaches. According to the Blue Coat report, "In 2011, malnets emerged as the next evolution in the threat landscape. These infrastructures last beyond any one attack, allowing cybercriminals to quickly adapt to new vulnerabilities and repeatedly launch malware attacks. By exploiting popular places on the Internet, such as search engines, social networking and email, malnets have become very adept at infecting many users with little added investment." That report also noted the increasing importance of social networking and stated that, "Since 2009, social networking has increasingly eclipsed web-based email as a method of communications" and that, "Now, social networking is moving into a new phase in which an individual site is a self-contained web environment for many users – effectively an Internet within an Internet."

The Current Environment for Implementing Security

IT security systems and policies have evolved and developed around the traditional application delivery architecture in which branch offices are connected to application servers in a central corporate data centers. In this architecture, the central corporate data center is a natural location to implement IT security systems and policies that provide layered defenses as well a single, cost efficient location for a variety of IT security functions. With the adoption of public cloud computing, applications and services are moving out of the central corporate data center and there is no longer a convenient single location for security policies and systems.

IT security systems and policies have traditionally distinguished between people who were using IT services for work versus those who were using it for personal use. The use of an employer provided laptop was subject to the employer's IT security policies and systems. In this environment, the use that employees made of personal laptops was generally outside of the corporate IT security policy. With the arrival of smartphones and tablet computers, the ownership, operating systems and security capabilities of the end user devices have changed radically. IT security policies and standards that were developed for PCs are no longer effective nor optimal with these devices. Most corporations have embraced the BYOD movement and end users are less willing to accept strict corporate security policies on devices they own. Additionally, strict separation of work and personal usage for security on an employee owned device is impractical.

The demands of governments, industry and customers have historically shaped IT security systems and policies. The wide diversity of organizations that create regulations and standards can lead to conflicts. For example, law enforcement requires access to network communications (Communications Assistance for Law Enforcement Act – CALEA) which may in turn force the creation of locations in the network that do not comply with the encryption requirements of other standards (e.g. Health Insurance Portability Accountability Act – HIPPA).

In order to determine how IT organizations are responding to the traditional and emerging security challenges, the **Survey Respondents** were asked a series of questions. For example, to get a high level view of how IT organizations are providing security, the **Survey Respondents** were asked to indicate which of a number of network security systems their organization supports. The **Survey Respondents** were asked to check all of the alternatives that applied in their environment. Their responses are shown in **Table 8**.

⁶ http://www.bluecoat.com/sites/default/files/documents/files/BC_2012_Security_Report-v1i-optimized.pdf

Table 8: The Network Security Systems in Use	
Network Security Systems	Percentage
Remote Access VPN	86.30%
Network Access Control	73.50%
Intrusion Detection/Protection Systems (IDS/IPS)	65.70%
Next Generation Firewalls (Firewall+IPS+Application Control)	56.90%
Secure Web Gateways	46.10%
Web Application and/or XML Firewalls	36.30%
Mobile Device Security/Protection	36.30%
Security Information Event Management	31.40%
Data Loss Prevention	24.50%
Password Vault Systems (either local or portal based)	12.70%
SAML or WS-Federation Federated Access Control	8.80%

One obvious conclusion that can be drawn from **Table 8** is that IT organizations use a wide variety of network security systems. A slightly less obvious conclusion is that

On average, IT organizations use 4.8 network security systems.

The **Survey Respondents** were asked to indicate the approach that best describes how their company uses data classification to create a comprehensive IT security environment. Their responses are shown in **Table 9**.

Table 9: Approach to Comprehensive IT Security	
Approach	Percentage
We have a data classification policy and it is used to determine application access/authentication, network and end user device security requirements.	42.90%
We do not have a data classification policy.	33.00%
We have a data classification policy and it is used to determine application security requirements.	13.20%
We have a data classification policy, but it is not used nor enforced.	11.00%

The data in **Table 9** represents a classic good news/bad news situation. The good news is that the majority of IT organizations have a data classification policy that they use to determine requirements. The bad news is that

Almost half of all IT organizations either don't have a data classification policy or they have one that isn't used or enforced.

In order to understand how IT organizations are responding to the BYOD movement, the **Survey Respondents** were asked, “If your organization does allow employee owned devices to connect to your network, please indicate which of the following alternatives are used to register employee owned devices and load authentication (e.g. certificate/private key) data onto those devices before they are allowed to connect to your company’s network.” The **Survey Respondents** were asked to check all of the alternatives that applied in their environment. Their responses are shown in **Table 10**.

Table 10: Alternatives to Support Employee Owned Devices	
Alternative	Percentage
Employees must install a VPN client on their devices for network access	53.90%
IT Administrator and/or Service Desk must register employee owned device for network access	47.40%
Employees can self-register their devices for network access	28.90%
Employees must generate and/or load X.509 certificates & private keys network access	13.20%
Employees must install a token authentication app on their devices for network access	10.50%

The data in **Table 10** indicates that while using a VPN is the most common technique that a wide range of techniques are used. VPN’s popularity comes in part from the fact that remote access VPN solutions implemented on new generation mobile devices have various capabilities to enforce security policies when connecting to the corporate network. Popular security checks include ensuring that a screen password is present, that anti-virus software is present and is up to date, that there is not rogue software on the device and that the operating system has not been modified.

Two different approaches have emerged to protect against lost devices. For the traditional PC, full disk encryption is typically used to protect data if the PC is lost or stolen. However, on new generation mobile devices, remote erase solutions are typically used to protect data. New generation mobile devices with smaller displays are often used more for content reading rather than content creation. As screen sizes and resolution improves, this situation may change. In order to understand how IT organizations have implemented full disk encryption, the **Survey Respondents** were asked to indicate which alternatives their organization implements relative to using full disk encryption on laptops and desktop PCs. Their responses are shown in **Table 11**.

Table 11: Techniques for Implementing Full Disk Encryption	
Alternative	Percentage
We do not use full disk encryption on PCs.	52.5%
We use software based disk encryption on PCs.	49.5%
We use hardware based self-encrypting rotating drives on PCs.	6.1%
We use hardware based self-encrypting Solid State Drives on PCs.	6.1%

The data in **Table 11** indicates that

Just over half of all IT organizations don't use full disk encryption on PCs.

The data also indicates that those IT organizations that do use full disk encryption do so by using a software solution and that a small percentage of IT organizations use multiple techniques.

The **Survey Respondents** were asked to indicate the approach that best describes their company's approach to Identity and Access Management (IAM). Their responses are shown in **Table 12**.

Table 12: How IAM is Implemented	
Approach	Percentage
We do not have a formal IAM program.	36.6%
We have an IAM program, but it only partially manages identities, entitlements and policies/rules for internal users.	25.8%
We have an IAM program and it manages identities, entitlements and policies/rules for all internal users.	20.4%
We have an IAM program and it manages identities, entitlements and policies/rules for end users for internal, supplier, business partner and customers.	17.2%

The data in **Table 12** indicates that only a minority of IT organizations has a IAM program that has broad applicability.

The **Survey Respondents** were asked to indicate how their company approaches the governance of network and application security. Their responses are shown in **Table 13**.

Table 13: Governance Models in Use	
Approach	Percentage
Network Security and Application Security are funded, architected, designed and operated together.	46.9%
Network Security and Application Security are funded, architected, designed and operated separately.	30.2%
Network Security and Application Security are funded jointly, but architected, designed and operated separately.	22.9%

The data in **Table 13** indicates that

In the majority of instances, network security and application security are architected, designed and operated separately.

Security as a Cloud Provided Service

As previously mentioned, IT organizations have shown a great interest in acquiring from CCSPs a wide range of functionality that historically has been provided by the IT infrastructure group; a.k.a., cloud networking services (CNS). This includes security. In particular, as was also previously discussed (**Table 7**), over a quarter of the **Survey Respondents** indicated that their company either currently acquires security functionality from a CCSP or they expect that their company will within the next year.

Security is clearly a very broad topic. That said, one of the largest, if not the largest sources of security vulnerabilities is Web based applications. As previously mentioned, a large part of the growing security challenge associated with Web based applications is the continually increasing business use of social media sites such as Facebook and of major Webmail services such as Yahoo. A company could implement a simple acceptable use policy that either allows or denies access to these sites. However, such a policy ignores the fact that these sites typically provide a variety of functions, some of which fall into the acceptable use policies of a growing number of organizations. To deal with the evolving use of multi-faceted social media sites

A cloud-based security service needs to be able to allow access to a social media site such as Facebook, but block specific activities within the site, such as gaming or posting.

Analogously, the CNS needs to have the granular controls to be able to allow users to send and receive mail using Yahoo, but block email attachments.

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

One part of the value proposition of a CBSS is the value proposition of any cloud based service. For example, a CBSS reduces the capital investment in security that an organization would have to make. In addition, a CBSS reduces the amount of time it takes to deploy new functionality. The speed at which changes can be made to a CBSS adds value in a variety of situations, including providing better protection against zero-day attacks⁷. Another part of the value proposition of a CBSS is that unlike a traditional security solution that relies on the implementation of a hardware based proxy, a CBSS can also protect mobile workers. The CBSS does this by leveraging functionality that it provides at its POPs as well as functionality in a software agent that is deployed on each mobile device. The use of a Cloud-based solution to provide mobile device management and security was discussed previously in this section.

In many instances, the best security solution is a hybrid solution that combines traditional on-premise functionality with one or more Cloud-based solutions. For example, in many cases IT organizations already have functionality such as web filtering or malware protection deployed in

⁷ http://en.wikipedia.org/wiki/Zero-day_attack

CPE at some of their sites. In this case, the IT organization may choose to implement a CBSS just to protect the sites that don't have security functionality already implemented and/or to protect the organization's mobile workers. Alternatively, an organization may choose to implement security functionality in CPE at all of their sites and to also utilize a CBSS as part of a defense in depth strategy.

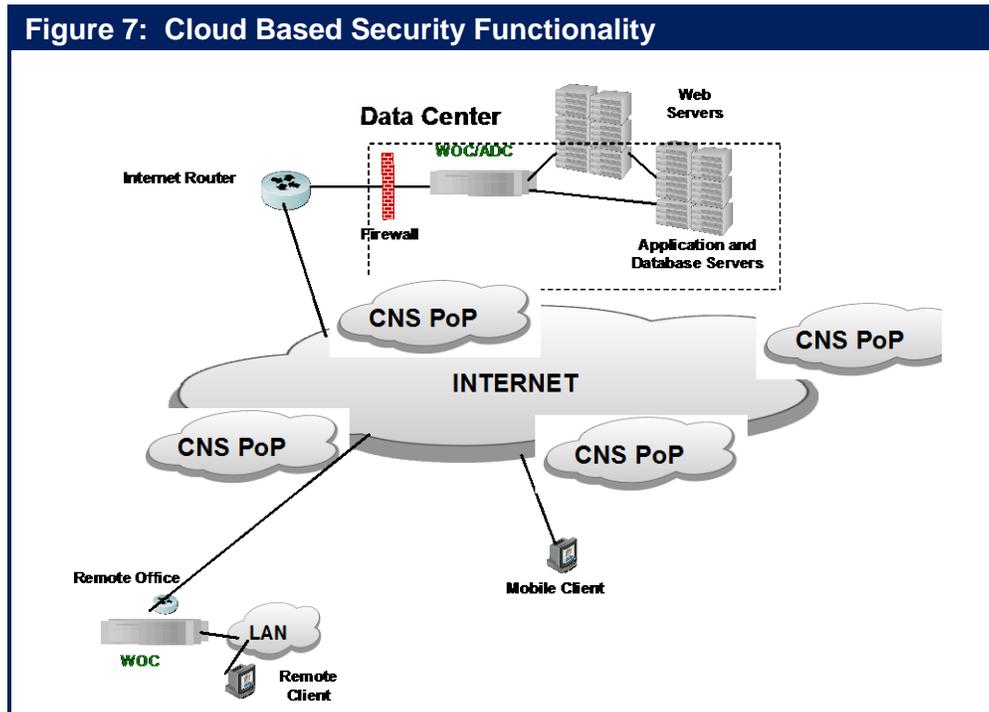
Other situations in which a CBSS can serve to either be the only source of security functionality, or to compliment CPE based implementations include cloud-based firewall and cloud-based IPS services. Such a service should support equipment from the leading vendors. Given the previously mentioned importance of hybrid solutions, the service should allow for flexibility in terms of whether the security functionality is provided in the cloud or from CPE as well as for flexibility in terms of who manages the functionality – a CCSP or the enterprise IT organization.

In addition to the specific security functionality provided by the CBSS, the CBSS should also:

- Provide predictive analytics whereby the CBSS can diagnose the vast majority of potential enterprise network and security issues before they can impact network health.
- Incorporate expertise, tools, and processes to ensure that the service that is provided can meet auditing standards such as SAS-70 as well as industry standards such as ITIL.
- Integrate audit and compliance tools that provide the necessary event-correlation capabilities and reporting to ensure that the service meets compliance requirements such as Sarbanes-Oxley, HIPAA, GLB and PCI.
- Provide the real-time notification of security events.

Web Application Firewall Services

The section of this report entitled *Wide Area Networking*, discussed how a Cloud-based service, such as the one shown in **Figure 7**, can be used to optimize the performance of the Internet. As will be discussed in this sub-section of the handbook, that same type of service can also provide some CCSBs.



The Role of a Traditional Firewall

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

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

Whereas network firewalls are focused on parameters such as IP address and port numbers, a more recent class of firewall, referred to as a Web application firewall, analyzes messages at layer 7 of the OSI model. Web application firewalls are typically deployed as a hardware appliance and they sit behind the network firewall and in front of the Web servers. They look for

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

The Role of a Web Application Firewall Service

There are fundamental flaws with an approach to security that focuses only on the perimeter of the organization. To overcome these flaws, most IT organizations have moved to an approach to security that is typically referred to as *defense in depth*. The concept of defense in depth is not new. What is new in the current environment is the use of a CBSS to provide Web application firewall functionality that is distributed throughout the Internet. This means that Web application functionality is close to the source of security attacks and hence can prevent many security attacks from reaching the organization.

In the current environment, high-end DDoS attacks can generate 100 Gbps of traffic or more⁸. Attacks of this magnitude cannot be prevented by onsite solutions. They can, however, be prevented by utilizing a CBSS that includes security functionality analogous to what is provided by a Web application firewall and that can identify and mitigate the DDoS-related traffic close to attack traffic origin.

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

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

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

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

⁸ [DDoS-attacks-growing-in-size](#)

- **Protecting DNS**

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

- **Implementing robust, multi-tiered failover**

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

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

- Block or redirect requests based on characteristics such as the originating geographic location and whether or not the originating IP addresses are on either a whitelist or a blacklist.
- Direct traffic away from specific servers or regions under attack.
- Issue slow responses to the machines conducting the attack. The goal of this technique, known as tarpits⁹, is to shut down the attacking machines while minimizing the impact on legitimate users.
- Direct the attack traffic back to the requesting machine at the DNS or HTTP level.

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

⁹ [Wikipedia Tarpit\(networking\)](#)

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.

For more information and for additional Webtorials® Editorial/Analyst Division products, please contact Jim Metzler at jim@webtorials.com or Steven Taylor at taylor@webtorials.com.

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

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

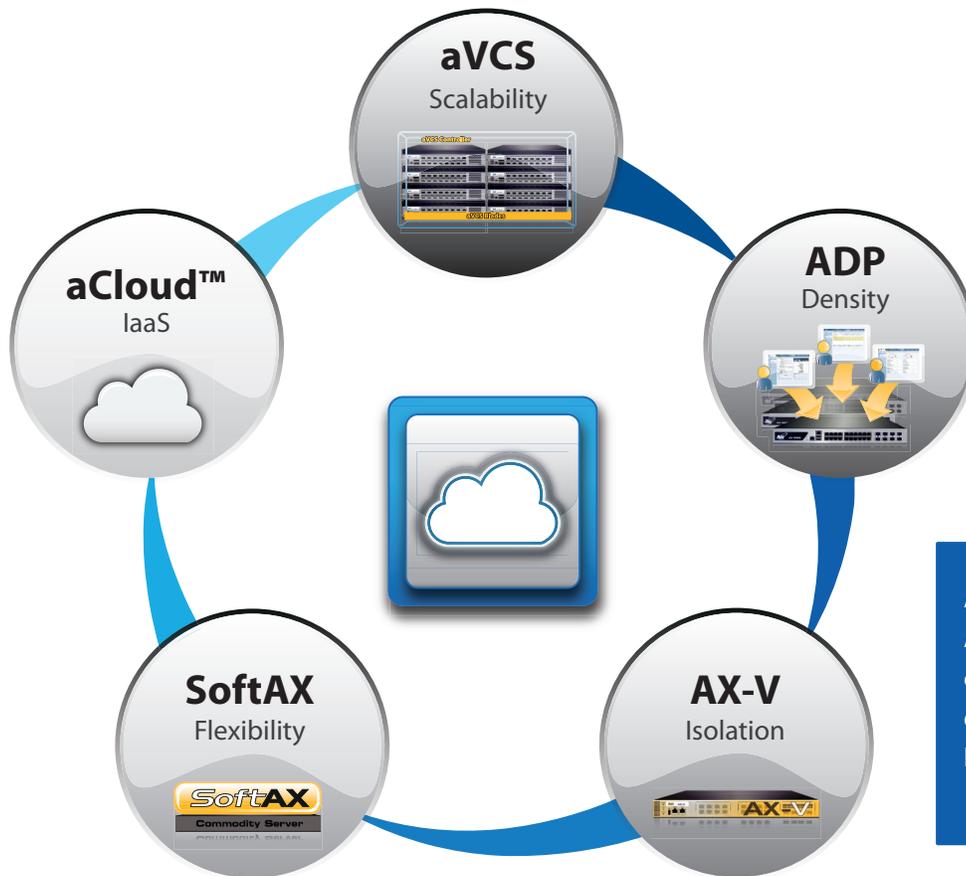
Professional Opinions Disclaimer

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

Copyright © 2012, 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.

Cost Effective Cloud Networking | Virtualization as the Enabler



A10 also offers a powerful choice of AX Series ADC form factors with comprehensive management options, delivering flexibility and efficiency for large scale deployments.

AX Series Virtualization Products & Solutions

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

SoftAX™

- SoftADC: AX virtual machine (VM) on commodity hardware
- Rapidly scale with commodity hardware
- Reduce hardware costs and upload to compatible cloud providers

AX-V Appliance

- SoftADC: AX virtual machine (VM) on AX Series hardware
- SoftAX flexibility with AX hardware performance and reliability
- Guaranteed performance, certifications, support and optimized hardware

AX Virtual Chassis System (aVCS™)

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

Application Delivery Partitions (ADPs)

- Divide the AX platform resources for individual applications
- Enables quality multi-tenancy with granular resource allocation
- Reduce the number of appliances to host multiple applications

The Application Fluent Data Center Fabric

Introduction

The rise of virtualization and cloud computing requires the selection of a best-of-breed data center switching solution as part of an enterprise's overall data center strategy. And at the heart of this strategy is the need to deliver a high quality user experience with new virtualized applications, including video, on new devices such as smart phones and tablets. However, the traditional 3-layer networks designed for a client/server communication model cannot meet the requirements of these new applications and devices, nor can it address the new requirements of virtualized servers and desktops.

Application Fluency for the Data Center

Resilient Architecture

- Simplified 10 & 40 GigE network with low latency and ready for 100 GigE
- Multi-path data center network extends between data center sites and to public cloud
- Supports definition of virtual data centers
- Ready for storage convergence with lossless Ethernet

Automatic Controls

- Application profiles ensure that the network is aware of application provisioning, security and QoS requirements
- The network will automatically sense virtual machine location and movement
- The network will automatically adjust to VM motion within and between data center sites

Streamlined Operations

- Applications are automatically provisioned
- Core switches automatically configure top of rack switches
- Converged management for data center network and virtual machine mobility
- Low power consumption

The Alcatel-Lucent Mesh

Alcatel-Lucent provides a unique Application Fluent approach to maximize the benefit from virtualization technologies for servers, the desktop, as well as the network. Alcatel-Lucent's application fluent data center fabric can scale from several hundred to over 14,000 server facing ports while keeping aggregate latency at 5ms, and can automatically adapt to virtual machine movement no matter which server virtualization platform is used.

The Alcatel-Lucent Virtual Network Profile (vNP), embedded in the Alcatel-Lucent Mesh, includes the critical information the fabric needs to understand each application, including provisioning requirements, security profiles, and expected quality of service levels. With this knowledge, the network can manage applications as services, including automatically discovering the location of each virtual machine, modifying the network configuration to follow virtual machine moves and providing an integrated view on visibility on VM movement and current location from a network perspective.

Application fluency in the corporate data center includes its transformation into a multi-site private cloud by extending layer 2 connectivity between data center sites and allowing for seamless delivery of public cloud-based services on the corporate network.

The Alcatel-Lucent Mesh enables enterprises to provide a high quality user experience with mission critical, real-time applications, and to improve agility in deploying new applications while significantly reducing data center costs.

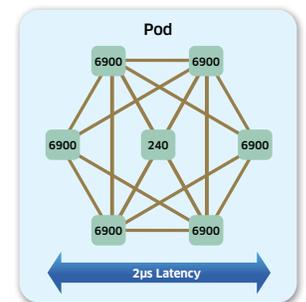
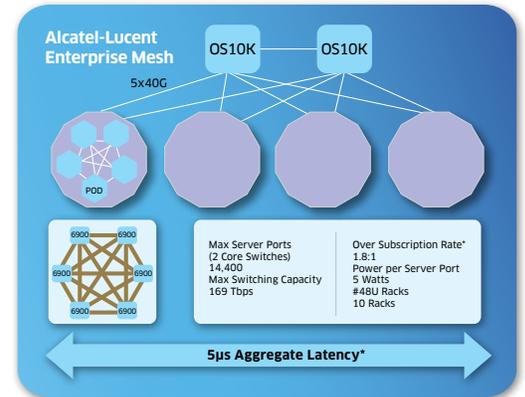
Open Ecosystems and Market Success

Alcatel-Lucent Enterprise is committed to open standards, allowing enterprises to select best-of-breed suppliers for their complete data center solution: servers, storage, data center fabric, and data center interconnect.

- Winner: Best of Interop 2011 for Data Center Switching and Storage
- Data center ecosystem partners include Emulex, NetApp, VMware, Citrix, and QLogic
- Participant in IEEE sponsored Shortest Path Bridging interoperability test with Avaya, Huawei, Solana and Spirent
- Over 20 million Ethernet ports shipped

For More Information

[Alcatel-Lucent Data Center Switching Solution](#)
[Alcatel-Lucent Application Fluent Networks](#)
[Alcatel-Lucent Enterprise](#)



*Assuming Server to Server Traffic 70% within a Pod, 20% between Pods and 10% Via Core

Visibility. Control. Optimize SaaS, BYOD, and Social Media

How to Lower Networking Costs and Safely Improve Performance

So many of the dominant trends in applications and networking are driven from outside the organization, including software-as-a-service (SaaS), bring-your-own-device (BYOD), Internet streaming video, and social networking. These technologies of an Internet connected world are fundamentally changing how we live and work every day. Yet, Network Administrators struggle to see and control these traffic streams from the Internet.

As businesses have opened their networks to SaaS applications, users are quickly starting using business bandwidth to access recreational websites and download BYOD updates, applications, and upload photos, videos and backups. This has created overburdened networks and slows the response of both cloud-based and internally delivered applications.

But with Visibility and Control from Blue Coat, Network Administrators can see all traffic on their networks and apply policies that can separate and control application traffic, and ensure internal and SaaS application performance.

First: Visibility of all traffic on all ports – Understand what is on your network

Blue Coat PacketShaper leverages Blue Coat WebPulse™, an

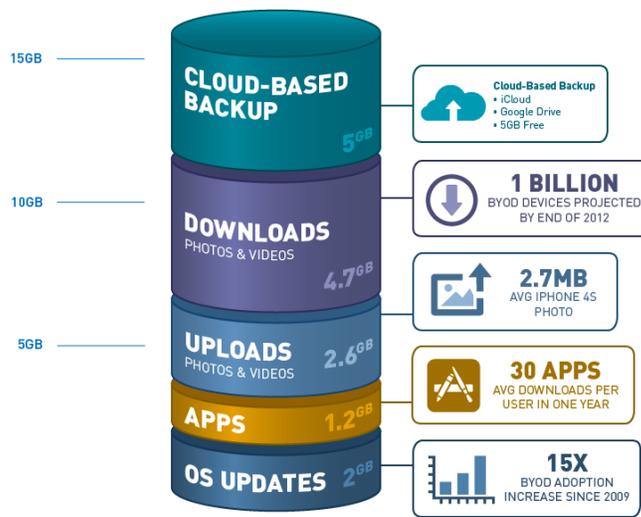
Internet Intelligence Service powered by a global community of 75 million users, the Cloud Service is able to deliver real-time categorization of Internet applications and web traffic.

WebPulse is based on sound analysis-system design principles:

- Massive input: WebPulse analyzes up to 1 billion web requests per day.
- In-depth analysis: 16 layers of analysis support over 80 categories in 55 languages.
- Granular policy: Up to 4 categories can be applied to each web request for multi-dimensional ratings.

- Speed: Automated systems process inputs – in most cases, in real time.
- Results: This collective intelligence allows WebPulse to categorize new Internet applications and websites quickly to PacketShaper without software updates/upgrades.

BYOD BANDWIDTH CONSUMPTION - JUNE 2011 TO JUNE 2012



The graphic details the impact of BYOD and Recreational video traffic can have on a network if left unchecked.

Second: Optimize Performance

SaaS, BYOD, Video and Social Media present challenges to network capacity and user patience. Blue Coat WAN Optimization helps overcome these challenges.

Chatty protocols and multi-megabyte files can hurt SaaS performance. Video requirements destroy capacity plans.

Blue Coat's asymmetric, on-demand video caching and live stream splitting boost video capacity up to 500x – whether it's corporate or recreational video. For SaaS, our CloudCaching Engine improves performance by 3-93x, dramatically raising productivity for SaaS users at branch locations.

And now Blue Coat ProxySG/MACH5 technology secures SaaS applications as it accelerates their performance. ProxySG/MACH5 connects directly to the Blue Coat Cloud Service, enforcing SaaS user policies and leveraging WebPulse to scan and filter cloud traffic. Branch users can access applications like SAP, Salesforce, and RightNow without the burden of bandwidth slowdowns or risk of malware threats.

agility
made possible™



On The Road To The Cloud?

With Converged Infrastructure Management and Network Automation, CA Technologies' allows you to transform your IT management functionality...reduce complexity and proactively optimize infrastructure while reducing costs...for a superior customer experience.

The Cloud Challenge... Increasingly CIO's and CEO's are looking to the IT organization to help deliver differentiation to the marketplace through innovation. As well, some organizations are looking to the Cloud to help them become more agile. Today "Cloud is synonymous with "Agility" but can you ensure your business services and guarantee application performance and availability in the cloud? How can you be proactive and optimize your infrastructure for lower costs while still delivering the highest quality user experience?

Cloud-Enable Your Network... CA Technologies Converged Infrastructure Management delivers ease of use and simple deployment while getting you up and running quickly with prescriptive OOTB capabilities- the benefit of IT organizations that say "It works as advertised." As well as functionality that can go deeper for dedicated IT teams showing them visibility into the infrastructure they specifically manage.



Access a single user interface for actionable performance, availability, flow capacity and application response information for all Layer 2 and Layer 3 technologies.

CA Technologies Converged Infrastructure Management delivers up to 25X Faster Problem Resolution While Reducing Total Cost by as Much as 50%. It helps you deliver a superior, differentiated customer experience – quickly and economically while -

Speeding proactive triage and remediation with less effort

- Analytics translate disparate data into intelligent views for up to 25x faster problem resolution

Meeting massive scalability demands cost-effectively

- Monitoring leading nationwide voice and video network with only two management servers

Shifting operations costs to innovation

- Converged infrastructure management reduces total costs by as much as 50%

Improving revenue streams

- Generate differentiated new sources of revenue and onboard new clients faster

The Cloud and Network Automation...CA Technologies

Network Automation enables cloud-readiness all across your network, making your operation more efficient, more cost-effective and safer. Automation allows your workers to be more productive, improves your compliance and security issues, diminishes the risk of failure and ensures safe and immediate disaster recovery.



Automated dashboard for data collection and analysis to improve remediation options like manual time and level of effort.

Just some of the ways Network Automation helps enable Cloud is:

- Tasking over manual, error-prone processes of provisioning network devices.
- Detecting network changes and addressing their impact with troubleshooting and notifying in real time when issues are detected.
- Knowing and showing who is on the network, where and when at any given time, as well as archiving historical configurations.
- Updating network configuration changes on a wide number of devices from a central location automatically.
- Obtaining a current inventory of all components on the network and detecting policy and compliance failures in real time.
- Backing up all network configuration on a near real time basis, allowing restoration to take place in a matter of minutes.

Whether you are looking for ease-of-use, enterprise scalability or automation on your journey to the cloud, CA Technologies will help you deliver the innovation and agility that today's business services demand.

Visit us at <http://www.ca.com/converge> or <http://www.ca.com/us/it-automation.aspx>

Simplify and Accelerate Private Cloud Deployments with Cisco's Virtual Networking Portfolio

Cisco and a Multi-Vendor Ecosystem Provide Cloud-ready Network Solutions

ROLE OF THE NETWORK PLATFORM IN CLOUD
<p>Access to Critical Data, Services, Resources and People</p> <ul style="list-style-type: none"> • Core fabric connects resources within the data center and data centers to each other • Pervasive connectivity links users and devices to resources and each other • Network provides identity- and context-based access to data, services, resources and people
<p>Granular Control of Risk, Performance and Cost</p> <ul style="list-style-type: none"> • Manages and enforces policies to help ensure security, control, reliability, and compliance • Manages and enforces SLAs and consistent QoS within and between clouds, enabling hybrid models and workload portability • Meters resources and utilization to provide transparency for cost and performance
<p>Robustness and Resilience</p> <ul style="list-style-type: none"> • Supports self-healing, automatic redirection of workload and transparent rollover • Provides scalability, enabling on-demand, elastic computing power through dynamic configuration
<p>Innovation in Cloud-specific Services</p> <ul style="list-style-type: none"> • Context-aware services understand identity, location, proximity, presence, and device • Resource-aware services discover, allocate, and pre-position services and resources • Comprehensive insight accesses and reports on all data that flows in the cloud

The Power of Cloud for the Enterprise

Business and IT executives are confronted daily by conflicting and exaggerated claims of how cloud will transform their industries, but the lure of transformative efficiency and agility is hard to ignore. Understanding the objectives and obstacles to cloud, as well as the solutions to overcome those obstacles is the key to achieving cloud-readiness.

Defining Cloud

In the simplest terms, cloud is IT delivered as a service over the network. Going a level deeper, cloud is a model in which IT resources and services are abstracted from the underlying infrastructure and provided on demand and at scale in a multi-tenant environment.

- *On demand* means that resources can be provisioned immediately when needed, released when no longer required, and billed only when used.
- *At scale* means the service provides the experience of infinite resource availability to meet whatever demands are made on it.
- *Multi-tenant environment* means that the resources are provided to many consumers - for example, business units -from a single physical infrastructure.

Note that the physical location of resources (on or off premises) is not a part of this statement. From the perspective here, that aspect has more to do with the way the cloud is sourced than with what the cloud does.

CISCO VIRTUAL NETWORK PORTFOLIO

Routing and Switching

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

Security and VPN

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

WAN Optimization

- Cisco Virtual Wide Area Application Services (vWAAS)

Network Analysis and Monitoring

- Cisco Prime Virtual Network Analysis Module (NAM)

Application Delivery Controllers

- Citrix NetScaler VPX virtual application delivery controller

Virtual Services Deployment Platform

- Cisco Nexus 1100 Series Virtual Services Appliance

Cloud Orchestration and Management

- Cisco Intelligent Automation for Cloud
- Cisco Virtual Network Management Center (VNMC)

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

Barriers to Adoption

While most enterprises have recognized the potential benefits of cloud, practical concerns and perceived challenges have hampered the widespread adoption of cloud technologies and services. Many of these barriers can be understood as questions of trust: Can the cloud be trusted to deliver the same capabilities at the same service levels in the same controlled way as traditional IT?

- **Security:** Can the same security available to applications be applied in the cloud?
- **Compliance:** Can applications in the cloud meet the same regulatory compliance requirements?
- **Reliability and quality of service (QoS):** Can the same service-level agreements (SLAs) for reliability and QoS be met in the cloud, especially given the multi-tenant use of the underlying IT infrastructure?
- **Control:** Can application owners still have the same amount of control over their applications and the infrastructure supporting them in the cloud?
- **Fear of vendor lock-in:** Will use of a particular vendor for cloud services or infrastructure prevent use of a different one in the future, or will the enterprise's data and applications be tightly locked into a particular model?

These concerns represent questions of technology and governance, but do not address any potential organizational friction that might arise from adopting cloud. For example, who will manage which part of the cloud or who will determine which applications to migrate to the cloud. Cisco believes that all these concerns can be met with the right technology, architecture, and approach.

Practical Solutions for Cloud-ready Virtual Networks and Infrastructure

The Cisco Virtualized Multi-Tenant Data Center (VMDC) architecture provides an end-to-end architecture and design for a complete private cloud providing IaaS capabilities. VMDC consists of several components of a cloud design, from the IT infrastructure building blocks to all the components that complete the solution, including orchestration for automation and configuration management. The building blocks are based on stacks of integrated infrastructure components that can be combined and scaled: Vblock™ Infrastructure Packages from the VCE coalition developed in partnership with EMC and VMware and the Secure Multi-Tenancy (SMT) stack developed in partnership with NetApp and VMware. Workload management and infrastructure automation is achieved using BMC Cloud Lifecycle Management (CLM). Clouds built on VMDC can also be interconnected or connected to service provider clouds with Cisco DCI technologies. This solution is built on a service delivery framework that can

be used to host other services besides IaaS on the same infrastructure: for example, a virtual desktop infrastructure (VDI).

These solutions for building private clouds are also being used by service providers to build cloud infrastructures on which to provide public, hybrid, and virtual private clouds to their enterprise customers. With service providers and enterprises, Cisco is developing an ecosystem of cloud providers, builders, and consumers. This ecosystem will be able to take advantage of common approaches to cloud technology, management, interconnection, and operation.

Where to Begin Your Cloud Journey

Cisco is working with its broad ecosystem of partners to assist some of the world's leading institutions in their initial cloud deployments. Cisco will have a central role in the unique journeys of enterprises, small and medium-sized businesses (SMBs), public-sector organizations, and service providers as they move to cloud.

When the topic of cloud comes up, the conversation often focuses on the newest technologies and the latest service provider offerings. However, Cisco believes that every conversation needs to begin with an understanding of the expected business outcomes. Is the goal lower total cost of ownership (TCO) or greater agility and innovation, or some blend of the two? The journey to cloud has many paths; starting the journey without a clear understanding of the destination can lead to disappointing results.

Enterprises should start the journey to cloud by answering some basic questions:

- What is the expected impact of cloud on my business?
- Which applications can and should I move to the cloud?
- What cloud deployment model is best suited for each of my applications?
- How do I maintain security and policy compliance in the cloud?
- How do I transition my organization to best take advantage of cloud?

The answers to these questions will fundamentally shape your cloud strategy. We are helping customers define and implement a pragmatic approach to cloud. We deliver solutions that address our customers' unique business architecture and needs, align with regulatory constraints, and are optimized according to the customer's individual preferences for performance, cost, and risk.

For More Information

As you begin your own journey to the cloud, we invite you to discuss the right approach for your organization with your Cisco account manager, channel partners, and other IT advisors. For additional information about cloud, please visit: <http://www.cisco.com/go/cloud>.



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.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: 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. (1110R)



Application Performance for Business Efficiency

The unique way to guarantee business application performance over the WAN, increase IT productivity and save on IT costs.

82%*

of organizations suffer application performance problems.

63%*

of organizations don't know the number of apps using the network.

72%*

of organizations use very occasionally their network to its full data transmission capacity.

Business and IT performance are tightly coupled...

Losing 5 minutes per day for poor application performance means 1% of productivity drop which can turn down profitability by 10%.

**Ipanema Killer Apps survey 2012*

IT departments are witnessing change at a pace never seen before

Transformation is occurring as CIOs seek to access the benefits offered by Unified Communications, cloud computing, internet-based applications and consolidation, amongst many other strategic projects.

These initiatives are aimed at increasing enterprise's business efficiency. While they simplify the way IT is delivered to users, they increase the complexity and the criticality of corporate networking as applications and users rely more than ever on the continuous, reliable and consistent flow of data traffic.

In order to protect the business and the significant investments made in transformative applications such as Unified Communications and SaaS the network must be more intelligent, more responsive and more transparent. Ipanema's revolutionary self-learning, self-managing and self-optimizing Autonomic Networking System™ (ANS) automatically manages all its tightly integrated features to guarantee the application performance your business requires over the global network:

- Global Application Visibility
- Per connection QoS and Control
- WAN Optimization
- Dynamic WAN Selection
- SLA-based Network Rightsizing

Business efficiency requires guaranteed application performance

- Know which applications make use of your network...
- Guarantee the application performance you deliver to users...
- Manage cloud applications, Unified Communications and Internet growth at the same time...
- Do more with a smaller budget in a changing business environment, and to prove it...

With Ipanema, control all your IT transformations!



For \$3/employee/month, you guarantee the performance of your business applications... and can save 10 times more!

Ipanema's global and integrated approach allows enterprises to align the application performance to their business requirements. With an average TCO of \$3/employee/month, Ipanema directly saves x10 times more and protects investments that cost x100 times more:

- **Application performance assurance:** Companies invest an average of \$300/employee/month to implement the applications that support their business. At a mere 1% of this cost, Ipanema can ensure they perform according to their application SLAs in every circumstance, maximizing the users' productivity and customers' satisfaction. While they can be seen as "soft money", business efficiency and investment protection are real value to the enterprise.
- **Optimized IT efficiency:** Ipanema proactively prevents most of the application delivery performance problems that load the service desk. It automates change management and shortens the analysis of the remaining performance issues. Global KPIs simplify the implementation of WAN Governance and allow better decision making. This provides a very conservative direct saving of \$15/employee/month.
- **Maximized network efficiency:** Ipanema's QoS & Control allows to at least doubling the actual capacity of networks, deferring upgrades for several years and saving an average of \$15/employee/month. Moreover, Ipanema enables hybrid networks to get access to large and inexpensive Internet resources without compromising the business, typically reducing the cost per Mbps by a factor of 3 to 5.

What our customers say about us:

Do more with less

"Whilst data volume across the Global WAN has increased by 53%, network bandwidth upgrades have only grown by 6.3%. With Ipanema in place we have saved \$987k this year alone."

Guarantee Unified Communications and increase network capacity

"Ipanema is protecting the performance our Unified Communication and Digital Signage applications, improving our efficiency as well as our customers' satisfaction. Moreover, we have been able to multiply our available capacity by 8 while preserving our budget at the same time."

Reduce costs in a cloud environment

"With Ipanema, we guaranteed the success of our cloud messaging and collaboration deployment in a hybrid network environment, while dividing per 3 the transfer cost of each gigabyte over our global network."





Enabling the cloud:

Award-winning NEC ProgrammableFlow® Open Software Defined Networking... ...delivering automated, efficient, and agile networks for the cloud

NEC's ProgrammableFlow network suite was the first commercially available SDN solution to leverage the OpenFlow protocol—enabling network-wide virtualization, allowing customers to easily deploy, control, monitor, and manage multi-tenant network infrastructure in a cloud environment. This architecture delivers better utilization of all IT assets, and helps provide ongoing investment protection as customers add functionality or upgrade their networks. NEC's approach simplifies network administration and provides a programmable interface for unifying the deployment and management of network services with the rest of IT infrastructure.

Specific functions customers prize include:

- **Drag and drop network design:** The GUI interface to the ProgrammableFlow Controller includes the familiar CLI found on most routers and switches today, so with minimal training a network admin can easily point and click to design an entire network from the single pane provided by the PF6800. This can radically reduce network programming and design time and errors caused previously by human intervention.
- **VM mobility:** With the ability to readily direct traffic throughout the data center—or throughout multiple data centers, it is possible to better manage all of the resources in a data center. For example, in NEC's own data centers in Japan, where they have recently implemented the ProgrammableFlow Fabric, it has enabled them to spread traffic between East and West Japan, offloading servers in East Japan that were nearing capacity, and postponing purchase of new servers, for a substantial saving. VM Mobility also enabled Nippon Express to complete a data center consolidation move that normally would have taken 2 months down to 10 days.
- **Bandwidth monitoring and traffic flow visualization:** This feature of the PF6800 provides performance monitoring of network flows and centralized management of network traffic, reducing bottlenecks and enabling smooth, streamlined network operations with substantially improved network admin productivity.
- **Secure, multi-tenant networks:** Secure, multi-tenant networks from the PF6800 enables customers like Genesis Hosting to expand their service offering with new sources of revenue potential. Genesis also reports software engineering investments were reduced by 100 hours each month with the advancements provided by ProgrammableFlow multi-tenancy.
- **Automation and administration of business policy to network management:** With network services aligned with business policy, automation such as prioritizing classes of applications or specific applications over other enterprise activity during peak loads is now possible with the ProgrammableFlow Network Suite, with multiple paths provided automatically. These capabilities offer significant value, particularly to enterprises engaged in heavy transaction loads.
- **Load balancing:** Traditional networking protocols often lead to performance-reducing bottlenecks. ProgrammableFlow uses path selection algorithms to analyze traffic flow across the network, check all available paths, and customize traffic flows to maintain performance and fully utilize network capacity. This increases the utilization of the network and improves application performance.

Backed by a 100-year history of technology innovation, NEC helps customers improve performance and solve their toughest IT challenges.

To learn more about how NEC can help you optimize your network for the cloud, visit necam.com/pflow or call your NEC Account Manager today.



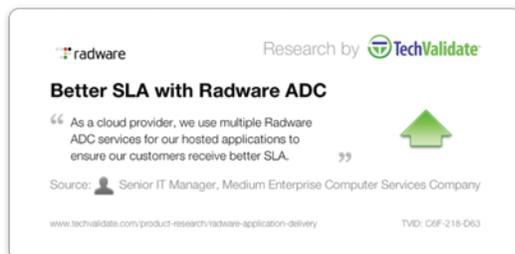
Expand Your Cloud Offering with Advanced Cloud ADC Solutions

Challenges in the Cloud Provider Business

The broad adoption of cloud based services by enterprise organizations and the multiple entrants into the cloud and hosting business challenges cloud providers to differentiate their service offerings and attract customers. Cloud providers face multiple challenges in establishing their business.



The first challenge is the infrastructure availability challenge. In an effort to provide uptime assurance at the base service level, or as a value added service offering, cloud providers must provide continuous availability of customer resources. One threat impacting the business availability is general connectivity: infrastructure outages and disruption events in which providers are dependent on external utilities and their running equipment. Failure to these can have significant adverse affect on the providers' business. Furthermore, part of the scalability value proposition of a cloud provider is the ability to scale-out application infrastructures – without load balancers, application scale-out is virtually impossible.



Above all, cloud providers are pressed to build solutions with minimal capital expenditure, maintain low operational costs and rapidly meet spikes in customer demand. Flexible procurement models by vendors and platforms that are easily scalable and centrally managed support the overall operational constraints faced by cloud providers.

Radware Solutions for Cloud Service Providers

Radware offers a set of fully integrated infrastructure availability and security solutions to meet the demands of cloud providers worldwide. Radware's solutions are comprised of the following components as illustrated in the figure below:

- **Radware ADC-VX™** – highly scalable ADC virtualization and consolidation solution offering high speed global and local load balancing, application acceleration and SSL offloading that supports dynamic availability requirements of cloud customers. ADC-VX can host multiple fully isolated, fully featured vADC instances.
- **Radware Alteon VA®** – flexible virtual ADC instance running atop most commercial, general purpose x86 server hypervisors.
- **Radware VADI®** – comprehensive virtual application delivery infrastructure solution including Alteon VA and ADC-VX-based virtual ADCs (vADC) and vDirect, an ADC service automation plug-in that simplifies ADC service deployment in cloud environments.

Radware's solutions enable cloud providers and hosts to offer more reliable and scalable infrastructure services to their customers. Resilience and scalability are key attributes of a cloud service as enterprises are contemplating the extent of cloud service adoption.

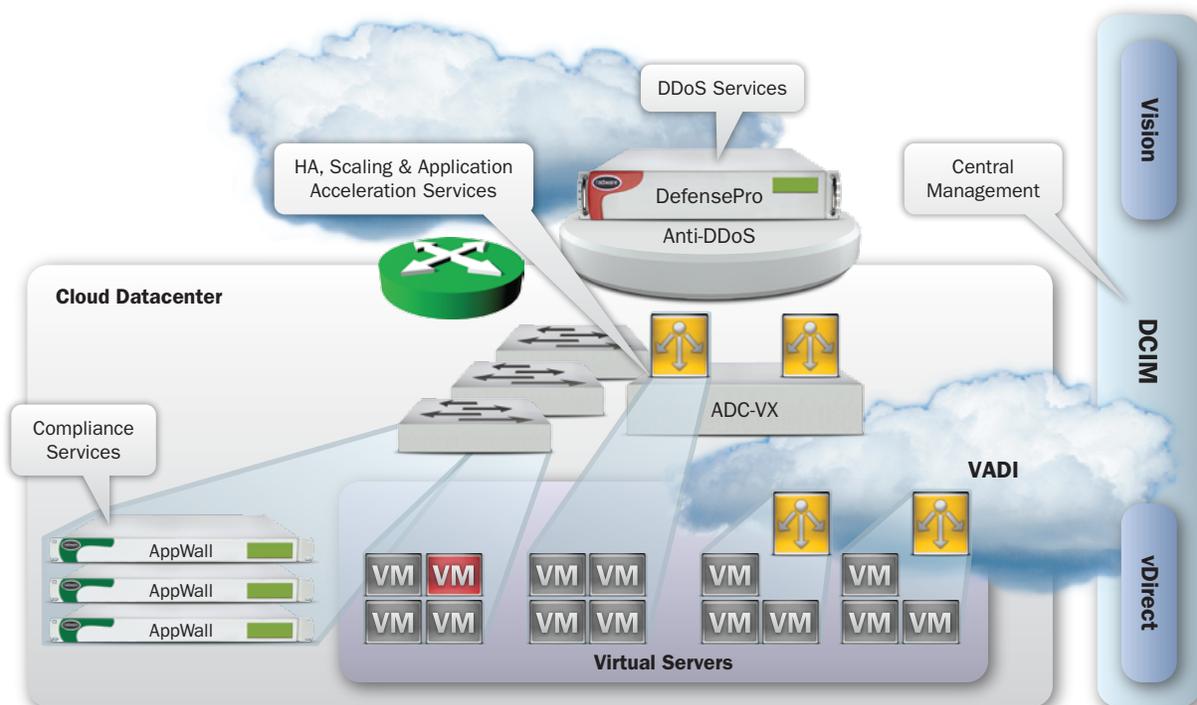


Figure 1 - Radware Service Architecture for Cloud

Benefits of Radware Solutions for Cloud Service Providers

1. Offer increased level of availability to cloud customers through highly available deployments of load balancing and application delivery services. High availability can be offered across any hardware form factor and location.
2. Seamlessly offer scale-out services to cloud customers inside cloud datacenters and across cloud datacenters by leveraging advanced health monitoring and KPI based global server load balancing.
3. Host a large scale of diverse services over a shared, purpose-built ADC infrastructure while fully isolating ADC instances associated with the different services.
4. Easily integrate application delivery and load balancing services into existing cloud service orchestration frameworks, home grown management tools and applications.
5. Simplify operations with a single management system controlling the entire set of Radware products in the cloud datacenter.
6. Cloud providers can offer additional value-add services such as application acceleration and application performance monitoring to their customers. All this while easily bundling the services into service packages and increasing customer confidence of rolling out applications in the cloud.

Summary

Radware application delivery and security solutions for cloud and hosting providers offer exceptional capabilities that greatly enhance the resilience, scalability and breadth of services offered by cloud and hosting providers. The value of the Radware is derived from 3 main benefits: (1) ability to enhance stability and scalability of cloud provider infrastructure (2) capability to help cloud providers build value added network services and offer these to their customers and (3) enabling these capabilities with minimal integration efforts and enhanced control.

Radware works with cloud providers globally addressing the key application delivery requirements presented in a cloud infrastructure through innovative cloud specific solutions.

For more information please visit <http://www.radware.com>