The 2011 Cloud Networking Report

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

Produced by:

Webtorials

Part 4: The Management of Cloud Computing

Featured Sponsors of Part 4:

ca

Infoblox

NETSCOUT.

Packet Design

VISUAL NETWORK SYSTEMS

Additional Sponsors:

A10 Networks

Alcatel·Lucent Enterprise

AVAYA The Power of We™

BlueCoat

certeon Accelerate Your Business

cisco

CITRIX

DELL Force10

exinda

extreme networks Make Your Network Mobile

hp

IBM

ipanema Technologies

LSI

radware

riverbed

Silver Peak

TALARI NETWORKS

VYATTA Open Networking

Table of Contents

The Management of Cloud Computing

Executive Summary

The **2011 Cloud Networking Report** will be published both in its entirety and in a serial fashion. This is the fourth of the serial publications. One goal of this publication is to describe the management challenges created by cloud computing and to identify the importance that IT organizations place on these challenges. Another goal of this publication is to describe how the adoption of cloud computing increases the difficulty of traditional management processes such as root cause analysis. The third goal of this publication is to describe some of the possible approaches that IT organizations can take to better manage in a cloud computing environment and to identify the value that IT organizations see in those approaches.

In order to quantify how IT organizations are responding to the challenges of managing a cloud computing environment, this publication includes the results of surveys that were given to the subscribers of Webtorials.com in 2010 and 2011. Throughout this publication, those two groups of respondents will be respectively referred to as The 2010 Webtorials Respondents and The 2011 Webtorials Respondents.

Importance of Managing Cloud Computing

One of the questions that was administered to The 2011 Webtorials 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, many of which were included in a similar question that was administered to The 2010 Webtorials Respondents. 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. Three of the twenty tasks that were included in The Question were managing private, managing hybrid and managing public cloud computing solutions. The responses of The 2011 Webtorials Respondents are summarized in Table 1.

Table 1: Importance of Managing Cloud Solutions

	Private Cloud	Hybrid Cloud	Public Cloud
Extremely	16.5%	9.2%	5.3%
Very	35.7%	31.1%	23.9%
Moderately	21.7%	25.2%	23.9%
Slightly	11.3%	15.1%	23.9%
Not at All	14.8%	19.3%	23.0%

One observation that can be drawn from the data in Table 1 is that

The majority of IT organizations believe that getting better at managing private cloud computing solutions is either very or extremely important.

Another observation that can be drawn from the data in Table 1 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 Evolving Management Environment

The Increased Focus on Services

Just as IT organizations are getting somewhat comfortable with managing the performance of applications, they are being tasked with managing the performance of services. IT professionals use the term *service* in a variety of ways. For example, the ITIL definition of service[1] states that a service:

- Is based on the use of Information Technology.
- Supports one or more of the customer's business processes.
- Is comprised of a combination of people, processes and technology.
- Should be defined in a Service Level Agreement (SLA).

In part because the ongoing adoption of virtualization and cloud computing has created the concept of everything as a service (XaaS), the term service as used in this section of the report will sometimes refer to services that IT organizations acquired from a public cloud computing provider; e.g., compute, storage, applications.

In order to quantify the interest that IT organizations have in managing this type of service, three of the twenty tasks that were included in The Question were:

- Effectively monitoring and managing compute services acquired from a third party such as Rackspace.

- Effectively monitoring and managing storage services acquired from a third party such as Rackspace.

- Effectively monitoring and managing applications acquired from a software-as-a-service provider such as Salesforce.com.

The responses of The 2011 Webtorials Respondents are summarized in Table 2.

Table 2: Importance of Effectively Monitoring and Managing Cloud Solutions

	Compute Services	Storage Services	SaaS Based Applications
Extremely	8.1%	2.9%	9.4%
Very	20.7%	20.0%	30.8%
Moderately	25.2%	28.6%	23.9%
Slightly	22.5%	20.0%	18.8%
Not at All	23.4%	28.6%	17.1%

As shown in Table 2, 28.6% of The Webtorials Respondents responded with “not at all important” when asked about the importance of getting better at monitoring and

[1] ITIL definition of service

managing storage services that they acquire from a public cloud computing vendor; a.k.a., an Infrastructure as a Service (IaaS) vendor.

The 28.6% was the largest percentage to respond with "not at all important" for any of the twenty management tasks that were presented to The Webtorials Respondents. Given that, it is possible to conclude that monitoring and managing the services obtained from an IaaS vendor is not an important task. However, that conclusion is contradicted by the fact that almost a quarter of The Webtorials Respondents indicated that getting better at monitoring and managing storage services acquired from an IaaS vendor was either very or extremely important. A more reasonable conclusion is based on the observation that many companies don't make any use of storage and compute services from an IaaS vendor and the ones that do often make only minor use of such services. Based on that observation, the data in Table 2 suggests that if a company makes significant use of the services provided by an IaaS vendor, then monitoring and managing those services is indeed an important task.

The term service as used in this section of the report will sometimes refer to business services that involve multiple inter-related applications. One of the management tasks that was included in The Question was "Manage a business service, such as CRM, that is supported by multiple, inter-related applications." The answers of The 2011 Webtorials Respondents are summarized in Figure 1.

Figure 1: Importance of Getting Better at Managing Services

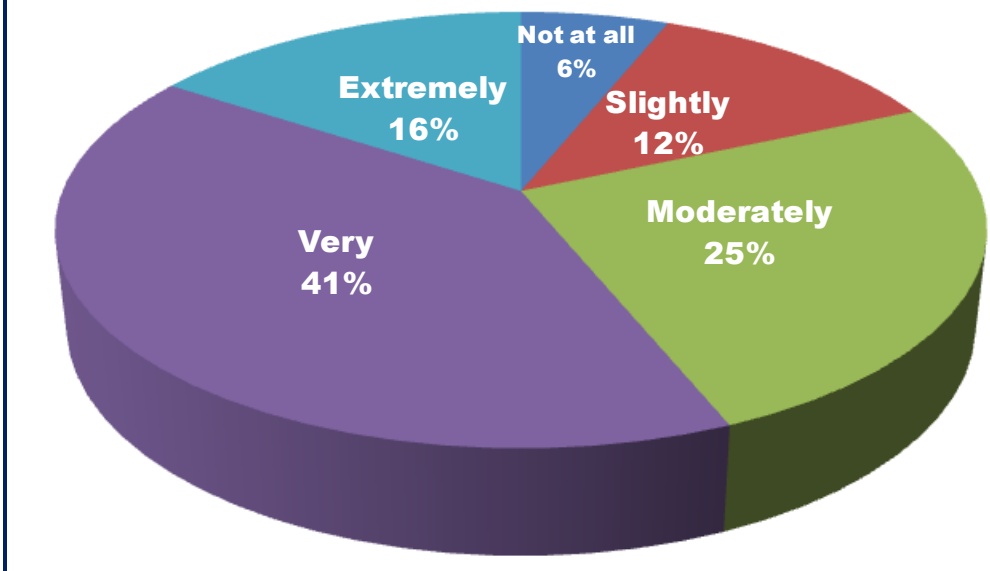

One observation that can be drawn from Figure 1 is that:

The majority of IT organizations believe that getting better at managing inter-related applications that comprise a business service is either very or extremely important.

Unfortunately, the adoption of cloud computing will further complicate the task of managing the inter-related applications that comprise a service. That follows because in a cloud computing environment, the applications that comprise the service will increasingly be supported by an infrastructure that is virtual. The challenges that are associated with managing server virtualization are discussed below. In addition, as is also discussed below, managing application performance in a cloud computing environment is extremely complex.

The 2010 Webtorials Respondents were asked to indicate the approach their organization takes to management. They were given the following choices and allowed to choose all that applied to their environment.

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

Their responses are summarized in Figure 2.

Figure 2: Focus of Management

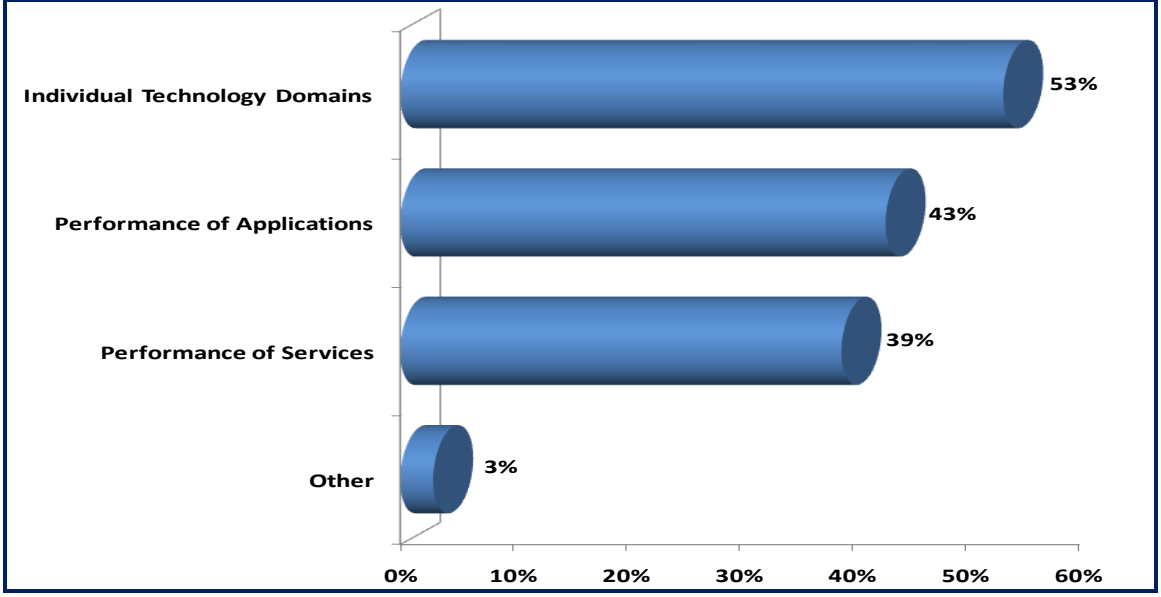

The data in Figure 2 indicates that the most frequent approach that IT organizations take to management is to focus on individual technology domains. However:

A significant percentage of IT organizations focus their management activities on the performance of applications and/or services.

The Growing Importance of Application Performance Management

In order to quantify how successful IT organizations are with their growing focus on managing the performance of applications and services, The 2011 Webtorials Respondents were given a set of statements and were asked to indicate which of the statements described their organization's approach to application performance management (APM). They were allowed to indicate all that applied.

Only about fifteen percent of The 2011 Survey Respondents indicated that their organization currently does a good job of APM. In addition, The 2011 Survey Respondents indicated by a significant margin that the approach that their organization takes to APM is that each technical discipline does its own thing vs. their using an approach that is top down and pretty tightly coordinated.

There is growing discussion in the industry about the best technical approach to implement APM. One approach is to be able to infer application performance based on management data, such as NetFlow, that is routinely collected by the network elements. An alternative approach is to use specialized agents to gather more sophisticated management data. Approximately twelve percent of The 2011 Survey Respondents indicated that their approach to APM makes heavy use of specialized agents to monitor the status of the various components of the application delivery chain.

One conclusion that can be drawn from the data discussed in the preceding two paragraphs is that

> ***APM is a work in progress. By that is meant that in spite of its importance, the vast majority of IT organizations don't do a good job of it.***

The statement that APM is both important and a work in progress is supported by the fact that a third of The 2011 Survey Respondents indicated that it was important to their organization to get better at APM over the next year.

Communications Based Applications

Communications based applications are an important class of application in part because these applications tend to be highly visible and their performance can degrade quickly if they experience impairments such as undo delay, jitter or packet loss. These applications are also important because as explained in the section of this report entitled ***The Wide Area Network***, over the next year almost 80% of IT organizations will increase their use of video, and in many cases the increased use of video will be substantial.

Another reason why communications based applications are an important class of applications in general and important relative to cloud computing in particular is that as discussed in the section of this report entitled ***The Emergence of Cloud Computing and Cloud Networking***, services such as VoIP and unified communications are now available from a cloud computing service provider (CCSP). As that section of the report also discussed, there is significant interest on the part of IT organizations to acquire both VoIP and unified communications from a CCSP.

To quantify the challenges associated with supporting a range of communications traffic, The 2011 Webtorials Respondents were asked to indicate how important it was over the next year for their IT organization to get better at managing the use of VoIP, traditional video traffic and telepresence. Their answers are summarized in Table 3.

Table 3: Importance of Managing the Use of Communications Based Traffic			
	VoIP	**Traditional Video Traffic**	**Telepresence**
Extremely Important	13.4%	6.8%	4.8%
Very Important	33.9%	20.3%	25.6%
Moderately Important	29.9%	29.7%	25.6%
Slightly Important	14.2%	28.0%	24.8%
Not at all Important	8.7%	15.3%	19.2%

The data in Table 3 shows that almost 50% of The Survey respondents indicated that getting better at managing the use of VoIP traffic is either very or extremely important to their IT organization. This is a significant percentage, particularly given that VoIP is not a new application. The challenge of managing VoIP will increase in those situations in which VoIP is acquired from a CCSP. In those instances, the IT organization will have to be able to gather and correlate management data from the CCSP, the network or networks that carry the VoIP traffic and the users' devices.

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 2010 Webtorials Respondents were asked to indicate their agreement or disagreement with three statements. The three statements and the percentage of The Webtorials Respondents that agreed with the statement are shown in Table 4.

Table 4: 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%

The data in Table 4 highlights the growing interest that IT organizations have in providing internal SLAs for at least some applications.

The vast majority of IT organizations provide an internal SLA for at least some applications.

One of the answers to The Question was managing internal SLAs for one or more business-critical applications. The responses of The 2011 Webtorials Respondents are summarized in Figure 3.

Figure 3: The Importance of Getting Better at Managing Internal SLAs

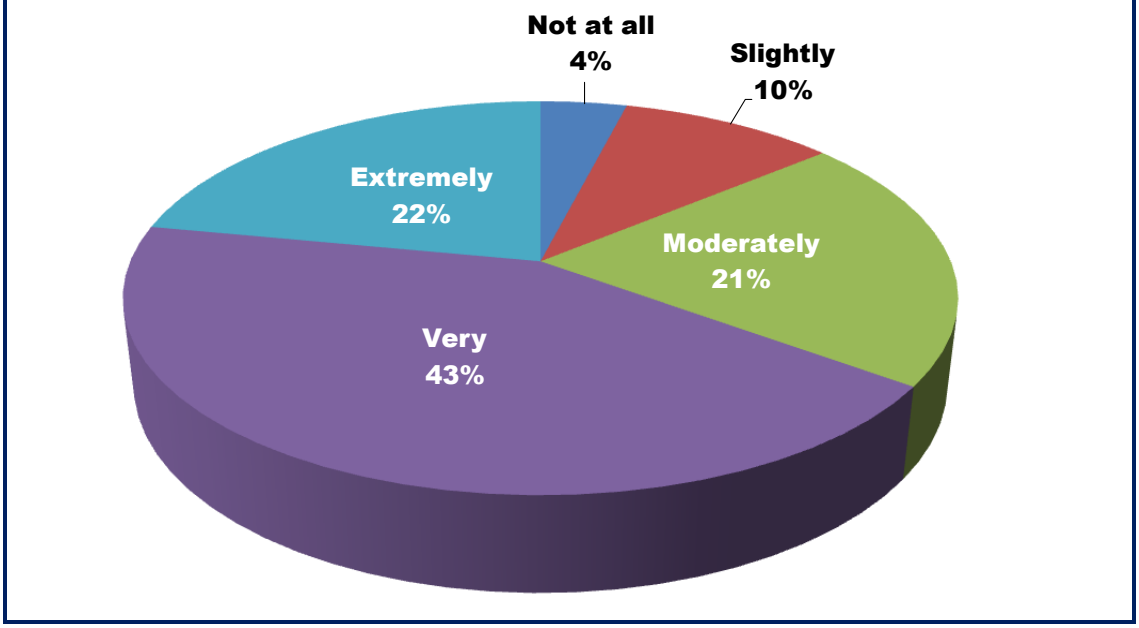

The data in Figure 3 leads to two related conclusions. One conclusion is that

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

A somewhat more subtle conclusion is that managing internal SLAs is difficult or else the majority of IT organizations would already be doing a good job of managing these SLAs and hence would not be striving to get better at the task. Unfortunately, the movement to utilize public cloud computing services greatly increases the difficulty associated with managing an internal SLA. That follows in part because of the difficulty of gathering all of the management data on an end-to-end basis that is necessary to effectively monitor an SLA. It also follows because as pointed out in the section of this report entitled ***The Emergence of Cloud Computing and Cloud Networking***, it is common for 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.

Root Cause Analysis

The 2011 Webtorials Respondents were asked how important it was over the next year for their organization to get better at rapidly identifying the causes of application degradation. Their responses are shown in Figure 4.

Figure 4: The Importance of Getting Better at Root Cause Analysis

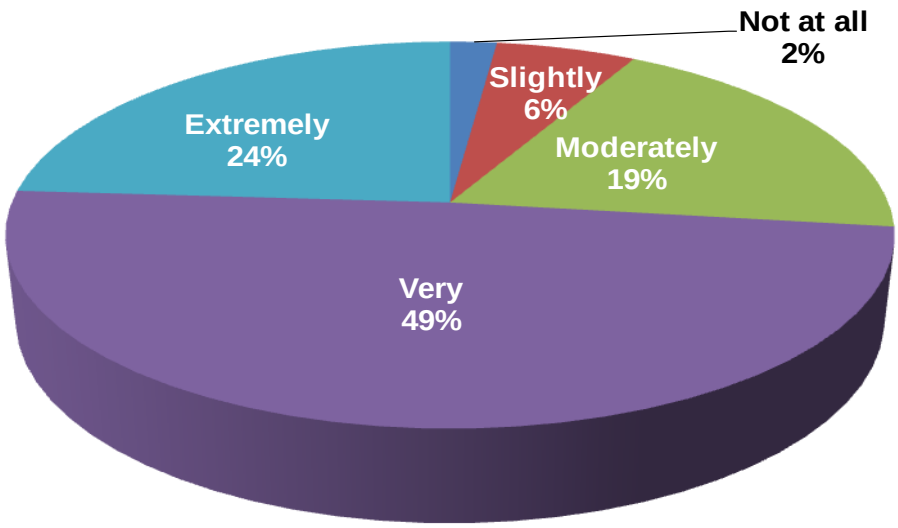

Comparing the answers that The 2011 Webtorials Respondents gave to the this management task to the other nineteen management tasks shows that:

Getting better at doing root cause analysis is the most important management task facing the vast majority of IT organizations.

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.

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. 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. In addition, 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.

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.

The 2011 Webtorials Respondents were asked how important it was over the next year for their organization to get better at identifying the components of the IT infrastructure that support the company's critical business applications. Their responses are shown in Figure 5.

Figure 5: Importance of Identifying Components of Critical Applications

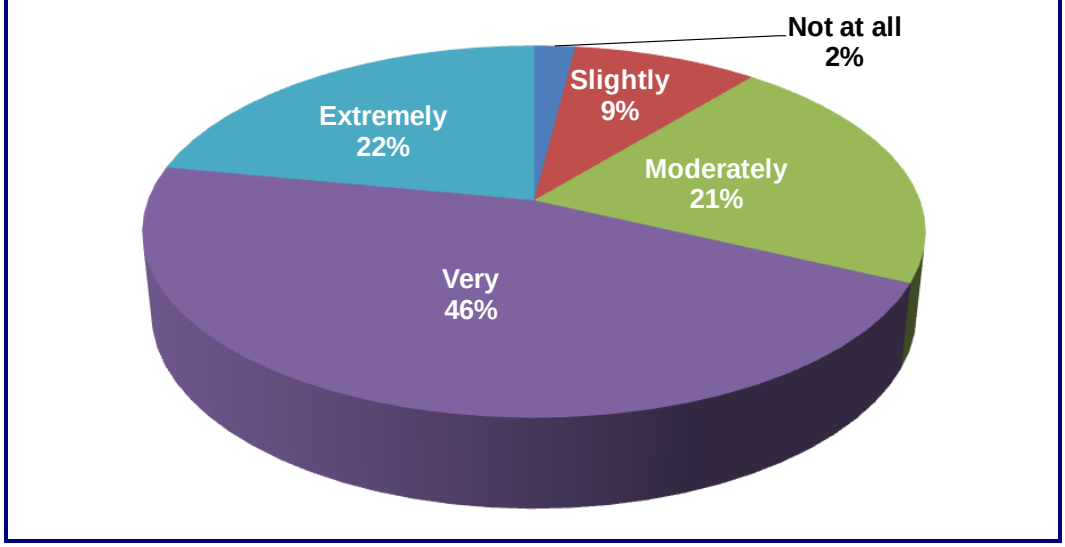

A clear observation that can be drawn from Figure 5 is that

Getting better at identifying the components of the IT infrastructure that support the company's critical business applications and services is one of the most important management tasks facing IT organizations.

Server Virtualization

As discussed in the section of this report entitled ***The Emergence of Cloud Computing and Cloud Networking***, there isn't a universally accepted definition of what is meant by cloud computing. That section of the report included a number of characteristics of a cloud computing solution, but also pointed out that there is not a litmus test to determine if a particular service is indeed a cloud computing service based on how many of the characteristics it supports. That said, the vast majority of private, public and hybrid cloud computing solutions are based at least in part on server virtualization. Hence, the management challenges that are associated with server virtualization can reasonably be regarded as management challenges for cloud computing.

As pointed out in Virtualization: Benefits, Challenges and Solutions, server virtualization creates a number of management challenges. For example, the need to manually reconfigure the network to support VM migration that was discussed in the section of the report entitled ***The Emerging Data Center LAN*** can be regarded as either a LAN challenge or a management challenge. 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 date 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 2011 Webtorials Respondents are summarized in Table 5.

Table 5: Importance of Managing Server Virtualization

	Traffic Between VMs	Move VMs Between Servers	Manage on a per VM Basis
Extremely	7.3%	15.4%	12.9%
Very	29.0%	32.5%	37.9%
Moderately	29.8%	20.5%	29.8%
Slightly	17.7%	18.8%	16.1%
Not at All	16.1%	12.8%	3.2%

One conclusion that can be drawn from the data in Table 5 is that managing the traffic that goes between VMs on a single physical server is not a very important task for the majority of IT organizations. Another conclusion is that

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 environment[2] 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.

One of the challenges associated with managing in any environment is that it is difficult to know the end-to-end path that packets take across a network. This management complexity comes in part from the distributed nature of IP. In particular, routers exchange reachability information with each other via a routing protocol such as OSPF (Open Shortest Path First). Based on this information, each router makes its own decision about how to forward a packet. There is, however, no single repository of routing information in the network. This lack of knowledge complicates tasks such as troubleshooting. The difficulty of knowing the path from origin to destination is greatly increased in a cloud computing environment because applications and services can be dynamically moved between servers both within and between data centers.

One of the fundamental issues relative to managing in a cloud computing environment is that the network topology becomes even more complex and hence understanding the end-to-end path becomes even more difficult.

In order to illustrate some of the other 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 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

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

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 looses 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 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 also 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.

Another 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.

Cloud Management Solutions

The Growing Use of Cloud Networking Services

As pointed out in the section of this 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 management, security, network and application optimization, VoIP, Unified Communications (UC) and virtualized desktops. This new class of solutions is referred to as Cloud Networking Services (CNS). That section of this report also presented the results of a survey in which The 2011 Webtorials Respondents were asked to indicate how likely it was over the next year that their company would acquire specific CNSs. The survey respondents were given nine types of services. Table 6 below highlights the interest that The 2011 Webtorials Respondents have in acquiring three specific CNSs.

Table 6: Interest in Cloud Networking Services

	Will Not Happen	Might Happen	50/50 Chance	Will Likely Happen	Will Happen
Security	39.0%	16.9%	16.9%	14.0%	13.2%
Network Management	38.8%	26.6%	7.2%	17.3%	10.1%
Application Performance Management	35.8%	28.4%	15.7%	12.7%	7.5%

One observation that can be drawn from the data in Table 6 is that:

Over the next year, more than a quarter of IT organizations will either likely acquire or will acquire security and/or management functionality from a CCSP.

Security as a Cloud Networking Service

Security is a very broad topic. That said, one of the largest, if not the largest sources of security vulnerabilities is Web based applications. 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 security based CNS 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.

Another one of the security challenges associated with the use of Web based applications that is rapidly increasing in importance is the growth of malware. To protect against malware, a security based CNS should be able to identify sites that are either suspecious or are known to distribute malware. In order to be effective, a CNS that provides Web content filtering or malware protection needs a source of intellectual capital that identifies known and suspected vulnerabilities. To be effective, this source needs to be both dynamic and as extensive as possible.

One component of the value proposition of a CNS that provides web filtering and/or malware protection is the standard value proposition of any cloud based service. That value proposition is that a cloud based service has the potential to lower the cost of providing the service, reduce the time it takes to implement the service and give the company that is using the service access to functionality that they couldn't otherwise acquire. Another component of the value proposition of a CNS that provides web filtering and/or malware protection is that

Unlike a traditional security solution that relies on the implementation of a hardware based proxy, a security based CNS can also protect mobile workers.

The security based CNS does this by leveraging functionality that it provides at its cloud data centers as well as functionality in a software agent that is deployed on each mobile device.

In many cases, the best use of a CNS is as part of a hybrid solution. For example, in some cases, the IT organization already has functionality such web filtering or malware protection deployed in CPE at some of their sites. In this case, the IT organization may choose to implement a CNS 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 CNS as part of a defense in depth strategy.

Other situations in which a security centric CNS 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.

Management as a Cloud Networking Service

As is the case with security, management is a very broad topic and hence it is possible to find a CNS that provides almost any possible form of management capability. For example, the preceding subsection discussed how a security based CNS could support mobile employees. In a similar fashion, a management based CNS can add value by helping IT organization to manage the burgeoning deployment of mobile devices.

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.

Route Analytics

As was previously mentioned, due to the distributed nature of IP it is sometimes difficult to know the end-to-end path that packets take across a network. While that is a challenge in any IT environment, it is a particularly difficult challenge in a cloud computing environment due to the dynamic nature of creating and moving virtual machines.

As shown in Figure 6, route analytics provides IT organizations and service providers with insight into the routing layer.

Figure 6: The Positioning of Route Analytics

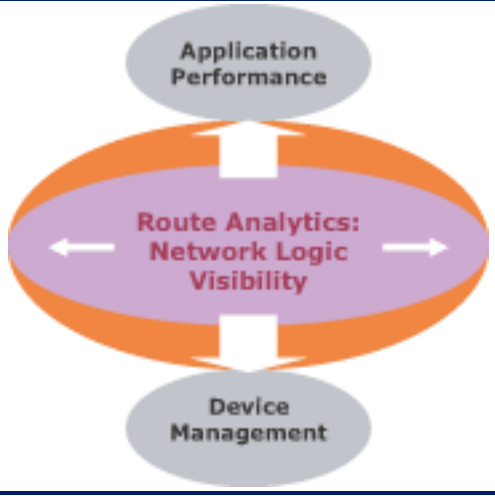

The value proposition of route analytics is that

Route analytics provides visibility, analysis, and diagnosis of the issues that occur at the routing layer in complex, meshed networks.

A route analytics appliance draws its primary data directly from the network in real time by participating in the IP routing protocol exchanges. This allows the route analytics device to compute a real-time Layer 3 topology of the end-to-end network, detect routing events in real time and correlate routing events or topology changes with other information, including application performance metrics. As a result, route analytics can help both IT organizations and service providers determine the impact on performance of both planned and actual changes in the Layer 3 network.

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. 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 at the flow level.

While changes in the virtual topology can be gleaned from flow analysis, a third approach to managing a virtualized 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.

Management Solutions Packaged with Converged Infrastructure

An increasingly popular approach to building cloud data centers is based on pre-integrated and certified infrastructure packages from either 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, APM 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 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, in order 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 a 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 Forum[3]. 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”.

IF-MAP is another emerging standard proposed by the Trusted Computing Group[4] and implemented by a number of companies in the security and network industries. It is a

[3] http://www.gridforum.org/
[4] http://www.trustedcomputinggroup.org/

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.

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.

A10 Networks

The Fastest Growing Application Networking Company

64-bit AX Series

Application Delivery

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

IPv6 Migration

- LSN, CGN, NAT444
- DS-Lite, 6rd
- NAT64 & DNS64
- SLB-PT, SLB-64

Cloud Computing & Virtualization

- SoftAX & AX-V
- AX Virtual Chassis
- AX Virtualization (Multi-tenancy)

Advanced Core Operating System (ACOS)

AX Series Advantage

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

Application Solutions

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

Blackboard · Microsoft · Microsoft Exchange · Microsoft Lync · IMPERVA · Infoblox · Juniper Networks · ORACLE · vmware READY

Cloud Networking – the not-so-quiet revolution

Avaya's vision for the Enterprise calls for a new level of synergy between people, the collaborative real-time applications they use, and the underlying network. A key building block for this vision is the foundational networking technology. As real-time communications continue the evolution to IP, the data network becomes completely integrated into the delivery of communications-enabled business services and mission critical business applications.

Avaya Networking provides advanced enterprise-class reliability, performance, and security that organizations throughout the world depend on to run their businesses. Because our solutions are streamlined to better utilize and manage networking resources, an Avaya data network can uniquely deliver both mission critical dependability and superior return on investment.

Virtualization within the Data Center is now taken for granted, with some declaring that 'Cloud Computing' will be the choice of most enterprises and that applications and information will become commodities. Experience has proved one thing; the Data Center of the future cannot be built on the technology of the past. General-purpose products, outmoded techniques, and legacy designs cannot be re-packaged as 'Data Center-ready'. The industry will take the best and leave the rest. Ethernet is readily available, cost-effective, extensible, and – as the 40/100 Gigabit developments prove – seamlessly without limitation of scale, however many of the underlying deployment methodologies are no longer an option.

Today's Enterprise network must be flatter, less tree centric, and able to support sustained east-west flows between multiple servers, in addition to traditional client/server transactions. Factors driving the transformation of enterprise networks include the transition to composite application architectures, an adoption of business operations intelligence applications (based on communications-enabled business processes and complex-event processing), and an increase in live virtual machine migrations. With each factor creating a unique challenge for the Data Center network, ranging between sensitivity to latency and loss, increased traffic levels (background noise), and risk of extended saturation of the common I/O connection, what's required is an agile, high-performance, latency-optimized networking solution that delivers exceptionally high performance.

To support the transition to a multi-dimensional environment the underlying network also needs to change. Provisioning needs to be simpler, and availability and performance need to scale seamlessly. Empowering a truly commoditized approach to service delivery requires a solution that is characterized by simplification, and a standards-based approach will help ensure an open architecture that avoids costly or inflexible lock-in.

Avaya is able to clearly demonstrate a set of differentiating benefits:

- Reduction in the configuration burden by up to 25X over the techniques traditionally implemented in large Data Centers
- Simplification of application implementation and number of devices affected, thereby reducing chances for configuration errors; it's these human-errors that account for up to 40% of all network downtime
- Data Center resiliency that delivers millisecond convergence times during failover and recovery

Enabling Enterprises to build a Private Cloud infrastructure that is extensible from Data Center to Campus and ultimately to the Branch Office; end-to-end network virtualization is an important element of the Avaya Virtual Enterprise Network Architecture (VENA). Designed for next-generation networking, Avaya VENA is a flexible solution that can be tailored to fit current business needs while providing a smooth migration path that accommodates business evolution. Addressing crucial Data Center requirements, Avaya VENA creates self-aware network infrastructures that simplify the logical provisioning of network services and provide the components required to create an Ethernet fabric featuring active/active connectivity for all attached servers, and service-orientated networking from Top-of-Rack to Core. Chief among Avaya VENA components are our innovative Switch Clustering and the IEEE's 802.1aq Shortest Path Bridging virtualization technologies – enhanced with enterprise-friendly, Layer 3 functionality, authenticated network access, and a network management toolset that simplifies deployment, monitoring, and troubleshooting.

Avaya, uniquely positioned based on decades of networking experience, helps ensure that the transition to the next-generation of fabric-based infrastructure is low-risk, seamless, and evolutionary. Avaya's pedigree of proven, ground-breaking innovation delivers a truly fit-for-purpose Cloud-ready solution that encompasses both the Data Center and the Campus; ensuring simplified yet optimized end-to-end connectivity between users and their content.

Giving Cloud Applications a Lift

Software-as-a-service (SaaS) enables businesses to quickly innovate and compete in worldwide markets while lowering IT costs. From ubiquitous Microsoft SharePoint for robust content management, to customer relationship management by Salesforce.com, to hosted Windows Media Services pushing out multimedia communications and training; SaaS powers business productivity anywhere in the world.

As SaaS platforms, however, these applications are beyond the bounds of IT control. As a result, latency, chatty protocols, and packet loss easily impact SaaS performance, inhibiting business productivity and competiveness. WAN optimization is essential to assure reliable user experiences, yet the conventional symmetric approach of deploying appliances at both ends of the transaction does not apply for cloud-based applications.

Asymmetric Advantage

Cloud-based SaaS utilizes common protocols (e.g. HTTP, SSL, TCP, FTP) to deliver applications and data over WAN/Internet connections. Protocols, such as HTTP/SSL that secures data over the internet/WAN, can be chatty. When combined with the latency, limited bandwidth and packet loss associated with the WAN, these protocols can reduce SaaS performance. Add to this the bandwidth demands caused by the rapid growth in the use of video and other multi-megabyte files within the SaaS infrastructure; it's easy to see how performance degradation compounds particularly for remote sites with limited bandwidth.

Blue Coat offers a different WAN optimization approach. With its asymmetric CloudCaching Engine, Blue Coat is able to overcome latency, chatty protocols, limited bandwidth and packet loss with 3x -110x performance improvements on the same WAN/Internet connectivity. Blue Coat's asymmetric WAN Optimization technology is able to reduce transfer times and the bandwidth consumed by accessing SaaS infrastructure.

Accelerate and Optimize Cloud SaaS Applications

Blue Coat MACH5 WAN optimization is a combination of five separate application management and tuning technologies that provide unrivaled improvements in application performance and bandwidth utilization. These technologies include:

1. **Bandwidth Management** assigns priority and network resources based not only on port or device, but on users, applications and content to more accurately reflect corporate policies on the network. Whether alone or integrated with network QoS, bandwidth management provides application intelligence to the packet switching network.
2. **Protocol Optimization** improves the performance of protocols that are inefficient over the WAN through specific enhancements that improve tolerance to the higher latencies occurring there. Blue Coat offers multiple improvements for TCP, CIFS, HTTP, HTTPS, MAPI and most streaming video and IM protocols.
3. **Byte Caching** stores repetitive traffic found in the byte stream and serves it locally to reduce the amount of traffic that traverses the WAN. Requests can be served from a local WAN optimization appliance, accelerating updates to existing files with significantly less bandwidth; leading to dramatic bandwidth savings.
4. **Object Caching** further reduces demands on bandwidth by storing and serving files, videos and web content locally, without the overhead and risk of traditional wide area file services. For content delivery, no technology does more to improve the end user experience by reducing latency and bandwidth.

5. Finally, inline **Compression** can reduce predictable patterns even on the first pass, making it an ideal complement to byte caching technology.

Blue Coat's asymmetric WAN optimization technology can also be combined with industry leading QoS and visibility, to ensure that SaaS application data and files are transferred with the right priority and sufficient bandwidth to avoid competition with recreational or non-critical traffic.

Setting the SaaS Example

To demonstrate these performance gains, tests were conducted using a simulated 1.544 Mbps (T1) WAN link with 100ms latency to simulate a Cloud SaaS (data center) to branch office environment. Results show access times between the branch and Cloud for Salesforce.com are reduced by up to 33x. SharePoint BPOS access via the data center or direct-to-net is reduced by 7x – 110x. WAN bandwidth demand required to stream video from Microsoft Media Services Servers was reduced by 65% - 95%, and enabled up to 100 users to view streaming video simultaneously. Even files saved over WAFS show dramatically reduced transfer times (see Figure 1):

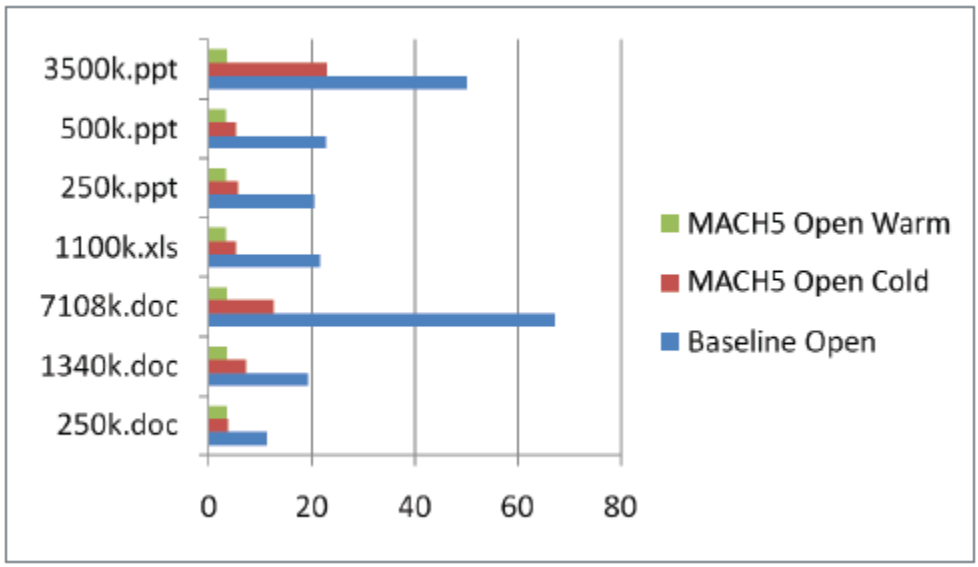

Figure 1: PowerPoint, Excel and Word file open times with and without MACH5 WAN optimization

The right choice for Cloud SaaS

Blue Coat Systems is a leading provider of web security and WAN Optimization solutions that secure and optimize the flow of information to any user, on any network, anywhere. Video and SaaS application delivery are today's IT's challenges, and with the right acceleration strategy you can gain superior business value from your internal and external infrastructure. Find out how Blue Coat can help you at bluecoat.com.

Cisco Unified Fabric

Converged. Scalable. Intelligent.

Cisco Unified Fabric is a flexible, innovative, and proven platform for physical, virtual or cloud deployments. It provides the foundational connectivity within and across data centers so resources are highly available wherever and whenever they are needed.

A key building block for cloud-based environments and virtualized data centers, the Cisco Unified Fabric brings unmatched architectural flexibility and scale to meet the diverse requirements of massively scalable data centers, bare-metal infrastructures, high performance and big data applications.

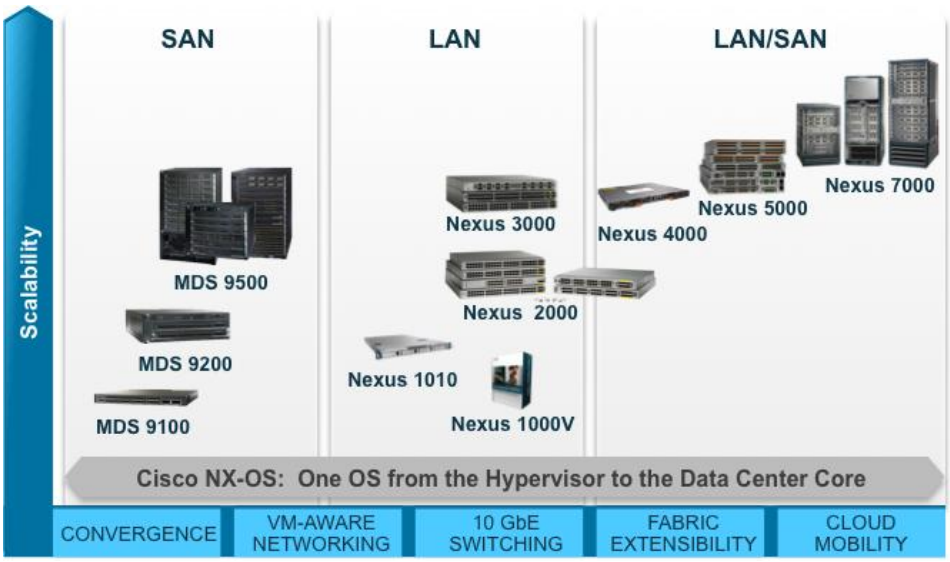

- Revolutionary fabric scale with over twelve thousand 10 GbE server connectivity with Cisco Nexus
- Highest 10Gb Ethernet density in the industry with Cisco Nexus 7000
- High performance and ultra-low latency networking at scale with Cisco Nexus
- Network services delivered in virtual and physical form factors with Cisco ASA, ASA 1000v, WAAS, vWAAS, VSG and more
- Virtual networking from the hypervisor layer on up with Cisco Nexus 1000v, VSS, VDC, and more
- High availability within and across devices with ISSU, VSS, vPC, and more.
- Flattened and scalable networking at Layer 2 and Layer 3 with Cisco FabricPath, TRILL, L3 ECMP, and more
- Overcome the challenges of expanding networks across locations and the limitations of network segmentation at scale with Cisco OTV, LISP, VXLAN, and more
- Unified operational, control, and management paradigms across the entire fabric with Cisco NX-OS, DCNM and open APIs
- Converged networking to carry every kind of traffic on a single fabric with DCB and FCoE with Cisco Nexus and MDS

Cisco Unified Fabric is a flexible, innovative, and proven platform for physical, virtual or cloud deployments with a non-disruptive, evolutionary approach to create future-proofed, service- and cloud-ready data centers and prevent 'rip and replace' for existing data centers. For more info: http://www.cisco.com/go/unifiedfabric

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.**

Consolidation and Cloud Computing Without Compromise

Citrix virtualization and cloud networking solutions accelerate, optimize, and secure application and service delivery from both the enterprise datacenter and the Cloud.

Starting Point

Server, storage, and other virtualization technologies are enabling organizations to consolidate infrastructure and transform to a dynamic, cloud computing model of IT service delivery. The result is a substantial reduction in capital and operating costs, *plus* a highly scalable and agile approach to meeting the computing needs of the business.

Next Step

To maximize gains, organizations should also extend virtualization and cloud computing principles to crucial networking components, including application delivery controllers (ADCs). Taking advantage of the flexibility and cost effectiveness of virtual appliance ADCs to more thoroughly ensure the performance, availability, and security of business-critical applications and services is a significant next step. Ideally, though, it should also be possible to consolidate numerous standalone ADCs to help reduce datacenter complexity and further control costs.

No Compromises

A new service delivery platform from Citrix, NetScaler SDX addresses this need by enabling multiple, independent instances of the NetScaler ADC to run on a single physical appliance. With NetScaler SDX, organizations gain the opportunity to reduce ADC footprint and total cost of ownership by maximizing consolidation of standalone ADC devices, across both different applications (i.e., horizontally) and different network zones (i.e., vertically).

> NetScaler SDX is a true multi-tenant platform that enables consolidation of core data center services. It delivers full functionality and meets the most demanding availability, security and performance SLAs.

Unique NetScaler Strengths

- **High consolidation density** – Up to 40 ADC instances can run independently on a single NetScaler SDX platform —more than double what competitors offer.
- **Complete isolation of ADC resources** – All critical system resources, including memory, CPU and SSL processing capacity, are assigned to individual NetScaler instances. Performance SLAs can thus be maintained on a per tenant basis.
- **Full ADC functionality** – Support for 100 percent of the NetScaler application delivery capabilities enables consolidation of all existing ADC deployments without any policy constraints or compromises.
- **Pay-As-You-Grow scalability** – An innovative, software-based Pay-As-You-Grow option provides essential elasticity, enabling organizations to scale performance and capacity on-demand without the need for expensive hardware upgrades.

For more information and a free NetScaler VPX download, please visit www.citrix.com/netscaler.

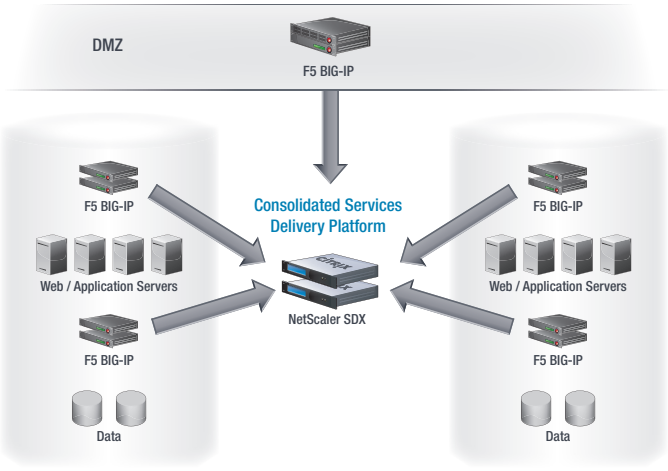

Network Services Consolidation

The power to do more

Unlock the full potential of your data center.

Servers
Intelligent compute architecture

Storage
Fluid data architecture

Networking
Open Cloud Networking

Application layer
Virtual Integrated System

Open Cloud Networking from Dell Force10 gives you the ability to achieve new levels of flexibility, performance and automation without having to rip and replace systems.

When we say "open" we mean it.

Open Architectures - build upon your existing infrastructure
Open Automation - simplify data tasks and vm management
Open Ecosystems - maximum choice and true architectural freedom

Is your network "open" for business? Find out more at **dell.com/OCN**

Scan this tag to unlock the full potential of your data center Dell Force10 Open Cloud Networking.

Get a free QR tag reader at your favorite app store.

Dell Force10
Data center switching, Top-of-rack, and next generation distributed core networks

HP FLEXNETWORK ARCHITECTURE

Meet the stringent performance, security, and agility demands of cloud computing

Enterprises are turning to the cloud to accelerate business innovation, improve business agility, and contain costs. Cloud computing reshapes the way applications are deployed and consumed and influences data center network designs. HP helps organizations build unified, virtualization-optimized networks that meet the rigorous performance, scalability, availability, and agility demands of the cloud.

HP FlexNetwork—an architectural blueprint for cloud-optimized networking

HP FlexNetwork architecture—HP's blueprint for cloud-optimized networking—lets enterprises securely deploy and centrally orchestrate cloud-optimized architectures that scale from the data center to the network edge.

HP FlexNetwork: Industry's only network architecture converging data center, campus, and branch office

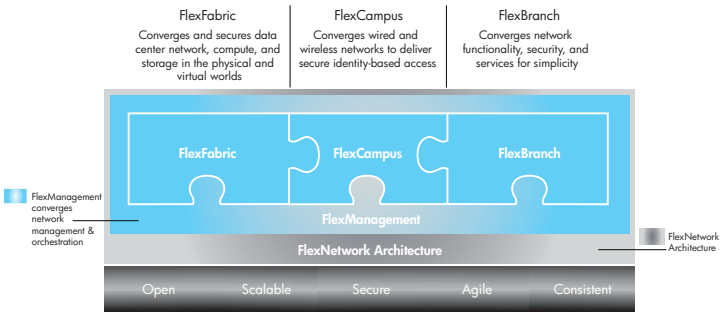

HP FlexFabric and **HP FlexCampus** enable the construction of flat, low-latency data center and campus networks with fewer layers, less equipment and cabling, and greater port densities.

HP FlexBranch includes comprehensive WAN optimization and routing solutions for delivering dynamic cloud-based services to geographically distributed enterprises.

HP FlexManagement provides a unified view into the virtual and physical network infrastructure, which accelerates application and service delivery, simplifies operations and management, and boosts network availability.

HP CloudSystem—a single platform for private, public, and hybrid clouds

HP CloudSystem is the industry's most complete, integrated, and open system for building and managing cloud services. Based on proven, market-leading HP Cloud Service Automation and Converged Infrastructure, HP CloudSystem combines servers, storage, networking, and security together with automated system and hybrid service delivery management. It enables organizations to build, manage, and consume cloud services across private clouds, public clouds, and traditional IT environments—without having to know, or care, whether those services come from HP CloudSystem's own "on-premises" resources or from the public domain.

HP CloudSystem and HP FlexNetwork networking solutions deliver:

- **Flatter and more efficient data center networks** with fewer layers, less equipment and cabling, and greater port densities
- **High-performance, low-latency intra-data-center connectivity** for virtual machine migration and bandwidth-intensive server-to-server communications
- **Virtualization-aware security** to partition multi-tenant environments and isolate virtual resources and intra-server communications flows
- **Optimal WAN performance** for the highest-quality end-user and application experiences and most efficient use of WAN resources
- **Unified administration and service orchestration** to accelerate the delivery of cloud-based applications and services
- **Multi-site, multi-vendor management** to connect and control thousands of physical and virtual resources from a single pane of glass

For more information

HP Networking Solutions: **www.hp.com/networking**

HP Cloud Solutions: **www.hp.com/go/cloud**

HP CloudSystem: **www.hp.com/go/cloudsystem**

To learn more about how HP can help you build a cloud-optimized data center network, please contact your HP account manager or reseller.

Created November 2011

EXECUTIVE VIEWPOINT

Maximizing Your IT Resources

Network Service Automation Rightsizes IT Staff and Delivers "Time to Value"

Steve Nye

EXECUTIVE VICE PRESIDENT, PRODUCT STRATEGY AND CORPORATE DEVELOPMENT, INFOBLOX, INC.

Steve Nye is the Executive Vice President of Product Strategy and Corporate Development for Infoblox, Inc. He is responsible for formulating the Company's longer-term strategy for portfolio and market expansion. Within his organization he directs all product management, marketing and business development activities. He oversees corporate development, which includes strategic alliances, both technical and marketing, as well as M&A activity.

Infoblox®

www.infoblox.com
1-866-463-6256
info@infoblox.com

WHAT IS THE BIGGEST CHALLENGE YOU SEE DRIVING IT DEPARTMENTS THESE DAYS?

Our customers and business partners say complexity is on the rise, which is putting more demands on IT to respond faster to business changes. However, because their budgets and staff are constrained, most companies cannot move quickly. They need help with scaling in an environment in which technology is moving faster than IT talent. We think new solutions that help manage the growing chaos surrounding IP initiatives will help increase network availability by reducing errors or delays in rolling out new services.

WHAT IS THE IMPACT OF VIRTUALIZATION ON NETWORK STAFF?

Virtualization breaks the traditional "one server, one application" architecture, and that creates new management challenges. For example, troubleshooting and seeing which virtual machine is connected to which port have become more difficult. Businesses need new discovery and visualization tools that automatically collect configuration information and automate repetitive and high-responserate chores such as assigning IP addresses and server names in a virtual environment. The task of issuing IP addresses and names for virtual machines should happen just as fast as a virtual machine can be provisioned. The network team in a virtualized environment must be as dynamic as the server team's ability to provision new systems. This type of automation is a critical part of any private cloud strategy.

HOW DOES THE INFLUX OF NEW MOBILE CONSUMER DEVICES CORRELATE WITH THE NEED FOR MORE NETWORK AUTOMATION?

IT managers are often not informed when new mobile devices come into the company. Employees bring them to work, or business units buy new systems because they do not want to wait for funds to be allocated to fulfill a critical business need. The IT department needs to know what is being attached to the enterprise network, because the impact of these devices can be significant. This shift to a more mobile and dynamic computing environment puts a strain on mission-critical network services such as Domain Name Service (DNS). As a result, IT needs simple-to-use, intuitive tools that monitor network activity while proactively managing and securing connections from a single central console.

HOW DOES THE MOVEMENT TO IPV6 AFFECT NETWORK STAFF?

The migration has already begun. T-Mobile is delivering IPv6 support in its phones, and these new IPv6 devices still need to connect to IPv4 networks. In the past, address management was done on spreadsheets, but 128-bit-IPv6 addressing brings an entire new set of challenges. When you add virtualization and cloud to this challenge, managing IP addresses with just a spreadsheet becomes impossible. IT teams will need automated network services.

WHERE SHOULD A COMPANY START AND HOW CAN YOU GAUGE SUCCESS?

Automation is a new "big idea." To some, it means ripping and replacing—or making significant investments in professional services and/or integration work. At Infoblox, we strive to make automation compelling by demonstrating that we can make adoption simple. By using automation, companies can reduce a 40-step process to a few clicks of a mouse. As a result, companies can make huge productivity gains and save money—many of our customers see an immediate increase in network availability and savings of millions of dollars annually by embracing automation.

Once companies see such results, they can expand their use of these tools and dramatically increase IT staff productivity. Infoblox's heritage is in automating network services such as DNS and IP address management. We anticipate that both automation and next-generation network services will be key elements powering the next 10 years of IT.

nano|engine

Full application control at 10% of the cost

www.ipanematech.com

A unique technology that breaks the price/performance barrier to guarantee business application performance in branch offices

- For the first time it is possible to guarantee application performance with a device compatible with branch office constraints;
- The nano|engines fully integrate with the other components of Ipanema's ANS solution;
- Plug-and-Play devices, nano|engines are managed under SALSA;
- Real-time changes in network performance and each user's behavior are taken into account in real-time.

Algorithms embedded in the nano|engine automatically adapt to real-time changes as they happen on the network:

- Traffic from private data centers mixed with traffic from external public clouds;
- Hybrid networks combining MPLS and Internet;
- Unified Communications branch-to-branch flows;
- Virtual desktops and rich media delivery…

The nano|engine's ability to guarantee application performance at the branch maximizes productivity, prevents brownouts and protects the business.

Ultra compact **nano|engine** appliances are tailored for providing full application control with unmatched performance/price ratio in broadband branch offices.

The **nano|engine** devices target broadband branch offices and provide:

- Application aware, **per connection Control and dynamic QoS** for public and private application flows to guarantee an excellent and stable Quality of Experience to each user;
- **End-to-end visibility** of application performance of each flow with comprehensive KPIs and application quality scores;
- **Dynamic WAN path selection** among up to 3 networks for optimized control of multi-attached branches, local Internet breakouts and hybrid networks.

Self-managed, nano|engines are installed at the edge locations of the WAN, typically between the CPE router and branch office LAN. Fully "Plug and Play," nano|engines require no on-site configuration. They operate under control of the central management software, SALSA. Customers simply need to plug the nano in, and configuration and provisioning are managed by SALSA.

The nano|engine family fits particularly well in B to C sectors like retail, finance and hospitality, where slow response times to access customer data or delays in processing an order lead to customer dissatisfaction and loss of productivity. Nano|engines' ability to guarantee application performance prevents any brownouts and protects the business.

The nano|2 addresses branch offices with up to 20 users and 4 Mbps while the nano|5 targets branch offices with up to 50 users and 20 Mbps.

Packet Design Solutions:

Packet Design's IP routing and traffic analysis solutions empower network management best practices in the world's largest and most critical enterprise, Service Provider and Government OSPF, IS-IS, BGP, EIGRP and RFC2547bis MPLS VPN networks, enabling network managers to maximize network assets, streamline network operations, and increase application and service up-time.

Route Explorer: Industry-Leading Route Analytics Solution

Optimize IP Networks with Route Explorer

- Gain visibility into the root cause of a signification percentage of application performance problems.
- Prevent costly misconfigurations
- Ensure network resiliency
- Increase IT's accuracy, confidence and responsiveness
- Speed troubleshooting of the hardest IP problems
- Empower routing operations best practices
- Complement change control processes with real-time validation of routing behavior
- Regain network visibility across outsourced MPLS VPN WANs

Deployed in the world's largest IP networks

400+ of the world's largest enterprises, service providers, government and military agencies and educational institutions use Packet Design's route analytics technology to optimize their IP networks.

Overview of Route Explorer

Route Explorer works by passively monitoring the routing protocol exchanges (e.g. OSPF, EIGRP, IS-IS, BGP, RFC2547bis MPLS VPNs) between routers on the network, then computing a real-time, network wide topology that can be visualized, analyzed and serve as the basis for actionable alerts and reports. This approach provides the most accurate, real-time view of how the network is directing traffic, even across MPLS VPNs. Unstable routes and other anomalies – undetectable by SNMP-based management tools because they are not device-specific problems – are immediately visible. As the network-wide topology is monitored and updated, Route Explorer records every routing event in a local data store. An animated historical playback feature lets the operator diagnose inconsistent and hard-to-detect problems by "rewinding" the network to a previous point in time. Histograms displaying past routing activity allow the network engineer to quickly go back to the time when a specific problem occurred, while letting them step through individual routing events to discover the root cause of the problem. Engineers can model failure scenarios and routing metric changes on the as-running network topology. Traps and alerts allow integration with existing network management solutions. Route Explorer appears to the network simply as another router, though it forwards no traffic and is neither a bottleneck or failure point. Since it works by monitoring the routing control plane, it does not poll any devices and adds no overhead to the network. A single appliance can support any size IP network, no matter how large or highly subdivided into separate areas.

Traffic Explorer: Network-Wide, Integrated Traffic and Route Analysis and Modeling Solution

Optimize IP Networks with Traffic Explorer

- Monitor critical traffic dynamics across all IP network links
- Operational planning and modeling based on real-time, network-wide routing and traffic intelligence
- IGP and BGP-aware peering and transit analysis
- MPLS VPN service network traffic analysis
- Network-wide and site to site traffic analysis for enterprise networks utilizing MPLS VPN WANs
- Visualize impact of routing failures/changes on traffic
- Departmental traffic usage and accounting
- Network-wide capacity planning
- Enhance change control processes with real-time validation of routing and traffic behavior

Traffic Explorer Architecture:

Traffic Explorer consists of three components:

- **Flow Recorders:** Collect Netflow information gathered from key traffic source points and summarize traffic flows based on routable network addresses received from Route Explorer
- **Flow Analyzer:** Aggregates summarized flow information from Flow Recorders, and calculates traffic distribution and link utilization across all routes and links on the network. Stores replayable traffic history
- **Modeling Engine:** Provides a full suite of monitoring, alerting, analysis, and modeling capabilities

Traffic Explorer Applications

Forensic Troubleshooting: Traffic Explorer improves application delivery by speeding troubleshooting with a complete routing and traffic forensic history.

Strengthened Change Management: Traffic Explorer greatly increases the accuracy of change management Processes by allowing engineers to model planned changes and see how the entire network's behavior will change, such as if there will be any congestion arising at any Class of Service.

Network-Wide Capacity Planning: Using its recorded, highly accurate history of actual routing and traffic changes over time, Traffic Explorer allows engineers to easily perform utilization trending on a variety of bases, such as per link, CoS, or VPN customer. Traffic Explorer ensures application performance and optimizes capital spending by increasing the accuracy of network planning.

Disaster Recovery Planning: Traffic Explorer can simulate link failure scenarios and analyze continuity of secondary routes and utilization of secondary and network-wide links.

Overview of Traffic Explorer

Traffic Explorer is the first solution to combine real-time, integrated routing and traffic monitoring and analysis, with "what-if" modeling capabilities. Unlike previous traffic analysis tools that only provide localized, link by link traffic visibility, Traffic Explorer's knowledge of IP routing enables visibility into network-wide routing and traffic behavior. Powerful "what-if" modeling capabilities empower network managers with new options for optimizing network service delivery. Traffic Explorer delivers the industry's only integrated analysis of network-wide routing and traffic dynamics. Standard reports and threshold-based alerts help engineers track significant routing and utilization changes in the network. An interactive topology map and deep, drill-down tabular views allow engineers to quickly perform root cause analysis of important network changes, including the routed path for any flow, network-wide traffic impact of any routing changes or failures, and the number of flows and hops affected. This information helps operators prioritize their response to those situations with the greatest impact on services. Traffic Explorer provides extensive "what-if" planning features to enhance ongoing network operations best practices. Traffic Explorer lets engineers model changes on the "as running" network, using the actual routed topology and traffic loads. Engineers can simulate a broad range of changes, such as adding or failing routers, interfaces and peerings; moving or changing prefixes; and adjusting IGP metrics, BGP policy configurations, link capacities or traffic loads. Simulating the affect of these changes on the actual network results in faster, more accurate network operations and optimal use of existing assets, leading to reduced capital and operational costs and enhance service delivery.

For more information, contact Packet Design at:

Web: http://www.packetdesign.com
Email: info@packetdesign.com
Phone: +1 408-490-1000

IT Organizations Find Key Enabling Technologies for Adopting Cloud Architectures

With engineers and administrators at companies like Google and Amazon redefining the standards for application, infrastructure, and data center efficiency, IT organizations have begun to reexamine their internal operations in order to apply the lessons of cloud computing. What they have discovered is that cloud computing is not a technology that can be applied, but an architecture that is built from many existing components and key enabling technologies.

Those key technologies support the centralization and consolidation of infrastructure, as well as the automation of IT processes, such as provisioning and scaling capacity. It is an architecture that favors economies of scale – such scale that for certain types of workloads, the most attractive and cost-effective deployment option is with third-party cloud providers. While many organizations initially hesitate to deploy their applications and store their data on the shared infrastructure of a public cloud provider, organizations that ultimately adopt third-party services recognize that shifting the burden of infrastructure administration to a provider operating at massive scale not only yields cost savings, but frees IT personnel to focus on more differentiated technology efforts.

As a result, the advent of cloud computing offers new choices in architecting IT infrastructure for the best possible blend of performance, availability, cost, and control. Finding that optimal balance will require both consolidation to fewer data centers and migration of selected applications and data to more cost efficient public cloud services.

> “Regardless of whether an organization chooses a private, public or hybrid cloud approach, they will likely experience performance problems as they encounter challenges caused by distance and the sheer growth of data.”
>
> *Eric Wolford, executive vice president of marketing and business development, Riverbed*

After identifying which applications are candidates to centralize into a consolidated private data center and which are candidates to move to a public cloud service, organizations must consider what their existing infrastructure supports and what new requirements will emerge. For example, centralizing resources and adopting public cloud services inherently requires users to depend on a network connection when accessing data and applications. However, migrating data and accessing applications across the WAN or public Internet is negatively impacted by distance, which introduces latency, as well as bandwidth congestion. For that reason, WAN optimization, with it’s ability to reduce data traffic and accelerate applications, is one of those key enabling technologies of cloud computing, by supporting the movement of infrastructure from inefficient, distributed models, to highly-automated, centralized cloud models. Not all WAN optimization vendors support the full spectrum of deployment scenarios that may make up an organization’s mix of private and public resources, but Riverbed Technology is one such vendor that has made it’s Steelhead product available in cloud deployments as well as traditional private WAN environments. “Regardless of whether an organization chooses a private, public or hybrid cloud approach, they will likely experience performance problems as they encounter challenges caused by distance and the sheer growth of data,” said Eric Wolford, executive vice president of marketing and business development at Riverbed.

Another challenge is in shifting legacy stove-piped application deployments to take full advantage of virtualized, scaled-out cloud architectures. Converting the application into a virtual machine is a critical step, but that alone does not ensure that an application can scale to more capacity and seamlessly migrate across available cloud resources. Application delivery controllers, which encompass traditional load-balancing functionality, can add a point of flexibility in an application’s architecture to allow organizations to seamlessly and automatically add additional server capacity to an application without disrupting its availability. Similarly, organizations can take advantage of hybrid cloud cost efficiencies by deploying an application across multiple public and private cloud data centers and using global load balancing technologies to manage application traffic across these multiple cloud deployments, reducing risk and improving the performance and capacity of applications.

However, a physical application delivery controller appliance tethers an application, even a virtualized one, to a limited set of resources in the data center. Thus, only a software-based virtual application delivery controller provides the flexibility necessary to enable cloud computing. As a virtual appliance, it can seamlessly migrate with a virtual application across available resources, within a single data center or between cloud data centers operated by different entities.

Transitioning to a cloud architecture, whether public or private, means applications run in massive, virtualized data centers. There are necessarily fewer of them and they will be farther apart and farther from end users. Thus, part of the transition to cloud computing is using enabling technologies to overcome the inherent challenges to running in virtual environments across wide distance.

WANop Anywhere

Highest Capacity, Any Platform, Lowest Cost

Accelerate cloud adoption with the leader in data center class WAN optimization. Silver Peak delivers the highest-capacity, lowest-cost, and most comprehensive portfolio of virtual and physical WAN optimization appliances.

DOWNLOAD FOR FREE TODAY!

www.vx-xpress.com

Silver Peak

www.silver-peak.com

WAN Virtualization Reduces Costs by 40% to 90%, Significantly Increases Bandwidth and Improves Reliability

As a CIO or IT manager responsible for network architecture, you may have connected more branch offices over recent years, or consolidated your data centers. As a result, you've witnessed first-hand the phenomenon that as servers move farther away from users, more WAN traffic is generated.

Also adding to your WAN traffic are the increased use of latency-sensitive applications, like VoIP, videoconferencing and desktop virtualization.

Because you don't want to hear unnecessary complaints when VoIP calls drop or applications perform poorly, you've likely purchased very expensive leased lines or MPLS services to ensure scalable, reliable and predictable WAN connectivity. Although alternative connectivity choices (e.g., Internet, DSL, etc.) are extremely attractive from a cost point of view, they simply don't provide the necessary four nines reliability to keep your business-critical applications up and running 24X7.

Into this carrier-pricing environment where a price/performance factor of 2x is enormous enters WAN Virtualization via Adaptive Private Networking (APN) technology from Talari Networks. WAN Virtualization brings Moore's Law and Internet economics to enterprise WAN buyers for the first time in 15-plus years. Further, Talari's Mercury appliances do this incrementally and seamlessly on top of existing networks – no forklift upgrades required.

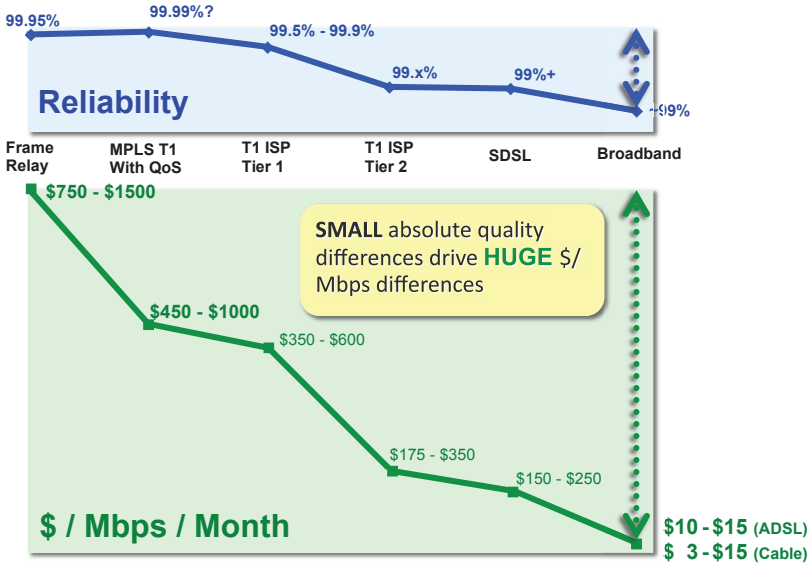

Figure 1: Private / Public WAN Pricing Disparity

Talari Networks Customer's 'AHA' Moment

Tim Hays at Lextron Inc. has used what is now called "cloud computing" in his network for over a decade. After he deployed Talari's solution, he said, "That was an 'aha' moment for me because I thought, 'Somebody finally gets it.' Talari's Adaptive Private Networking technology allows me to route each packet over the best, most reliable route, over multiple paths, including private lines, MPLS, DSL, and cable modem. By using WAN Virtualization, we've essentially created our own, big, private tunnel that aggregates different types of connectivity transparently across the Internet."

Real-Time, Per-Packet Traffic Engineering

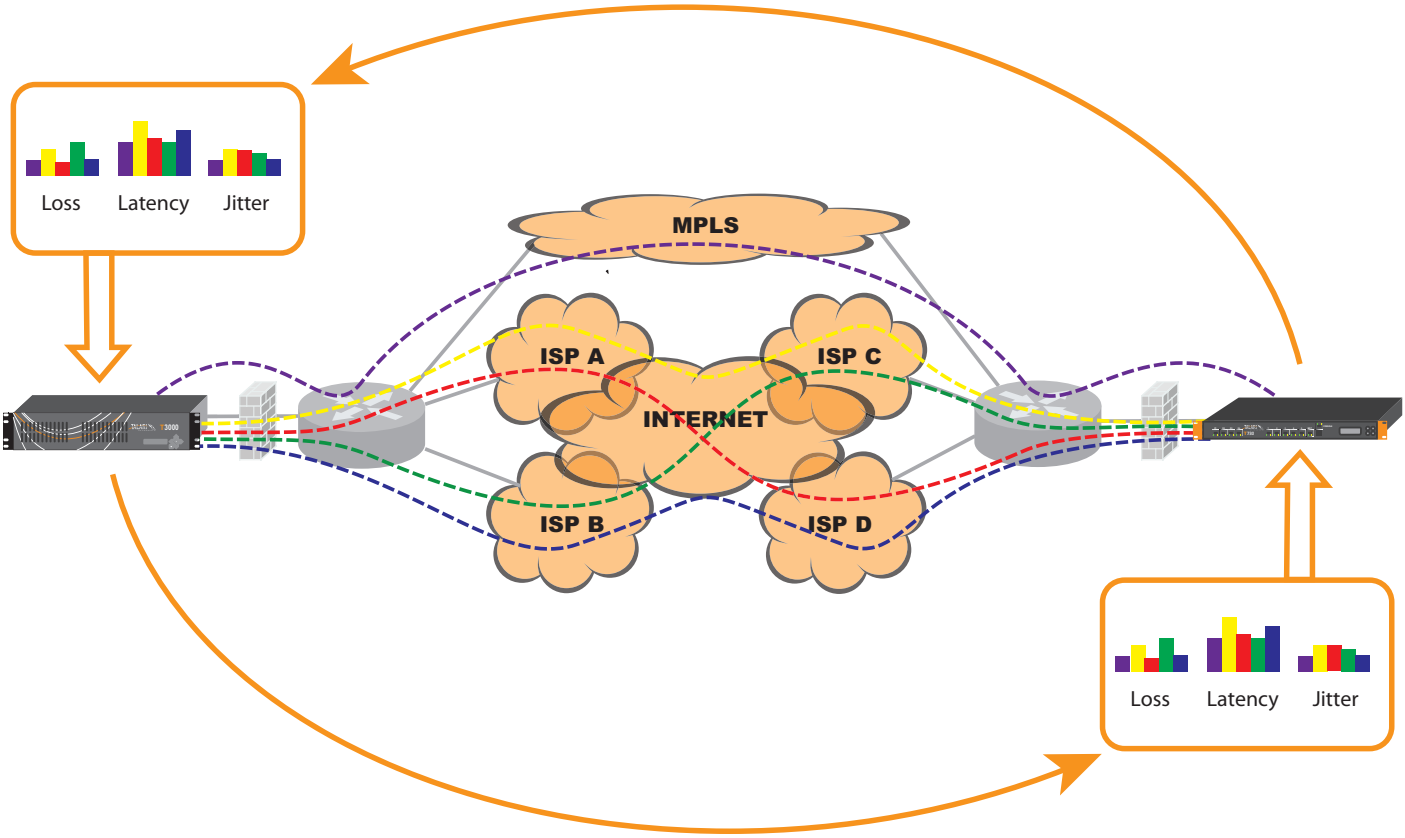

Figure 2: Continuous Measurment and Adaptation to Network Conditions

Requiring only two IP connections at each site which can include an existing private WAN connection, WAN Virtualization combines a variety of networks into a virtual WAN to deliver packets without being lost or excessively delayed 99.99% of the time. All network paths between locations are continually measured to determine current conditions. This allows each and every packet to be sent on the most appropriate path as determined by the type of traffic and available network resources. In addition, sub-second response to any congestion detected ensures predictable performance for all applications.

With this approach, Talari customers are building WANs where:

- **30 to 100 times more bandwidth can be purchased for every dollar spent**
- **Ongoing monthly WAN service charges can be reduced by 40% to 90%**
- **The resulting network is more reliable than any single MPLS private WAN**
- **Public cloud resources can be accessed with high reliability**

An APN Appliance for Every Situation

The Mercury family of APN appliances offer a wide range of perfomance points that span from large data centers to small remote offices and can be seamlessly added to your existing network in an overlay configuration to leave your current routed infrastructure intact. This allows you to introduce WAN Virtualization at your own pace to eventually migrate some or all of your locations off expensive private WAN connections.

Talari's customers see significant reductions in their ongoing monthly WAN expense that results in payback times for their WAN Virtualization deployments in the range of 6 to 12 months.

To learn more about how WAN Virtualization can transform the economics of your WAN please contact Talari Networks: **www.talari.com**.

VISUAL NETWORK SYSTEMS™

Application Performance Management in the Cloud

By 2016, 41% of all enterprise communications application users worldwide will have migrated to the cloud according to a study by ABI Research. That translates into trillions of dollars in business revenue depending on the delivery of these services. Application performance management will become even more critical to daily operations, but a surprisingly small number of cloud-based application users have adequate performance management software monitoring service delivery today.

In today's non-Cloud environment, with the current technology for application performance management, it is possible to instrument and collect run-time metrics, and provide access to management tools to analyze and report on the metrics. It is also possible to obtain a comprehensive view of an application including end user experience, specific transactions, and the supporting delivery infrastructure in order to manage the availability and performance of a business service.

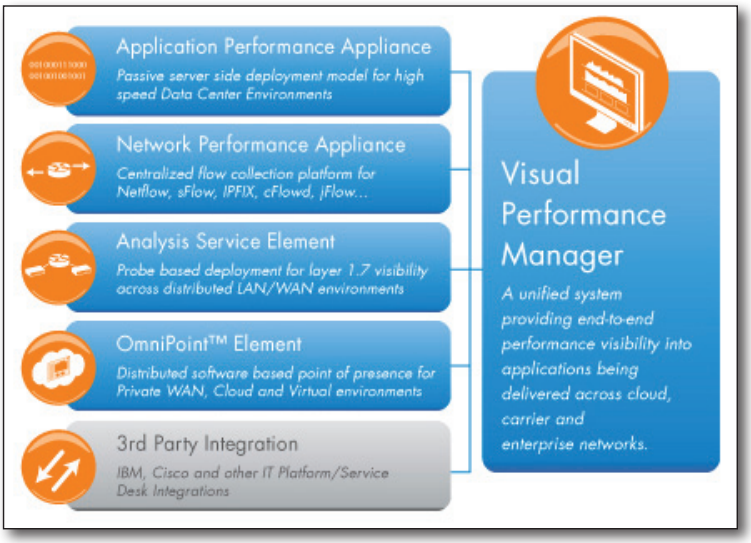

Visual Performance Manager is a unified system providing end-to-end performance visibility into applications being delivered across cloud, carrier and enterprise networks

When parts or all of an application moves to a Cloud, the view into the application is disrupted. One thing that doesn't change for both Cloud and non-Cloud environments is that the users, representing the business, expect the same level of availability, access to the applications and performance. Here in lies the challenge.

At a high level, the Cloud infrastructure includes the application delivery infrastructure which is made up of the applications running in virtualized environments and the network that supports the delivery of the applications to users. While the infrastructure itself is still made up of switches, routers, firewalls WAN optimization devices, VM Hosts and servers; the new part is that there could be more than one owner, such as cloud service providers and the private enterprise, for different parts of the delivery infrastructure.

For an Enterprise IT organization that is building a private cloud and virtualizing applications, or migrating to a hybrid cloud, the following are a few criteria that can help you find the right solution as you evaluate application performance management products.

- **Bridge application and network performance management between cloud and non-cloud environments –** More than likely, you are not moving all applications to the cloud. Whether you are in a transition to migrate applications or simply maintaining both cloud and non-cloud based applications, you are presented with the challenge of managing availability and performance for both sets of applications.

- **Flexible data collection instrumentation -** Within your private cloud, the challenge is visibility of applications in a virtualized environment. It is important that the instrumentation allows you to measure the performance of multi-tier applications as well as providing you with transaction level information for root cause analysis when performance degrades. This requires supporting deployment models to see the intra virtual machines traffic.

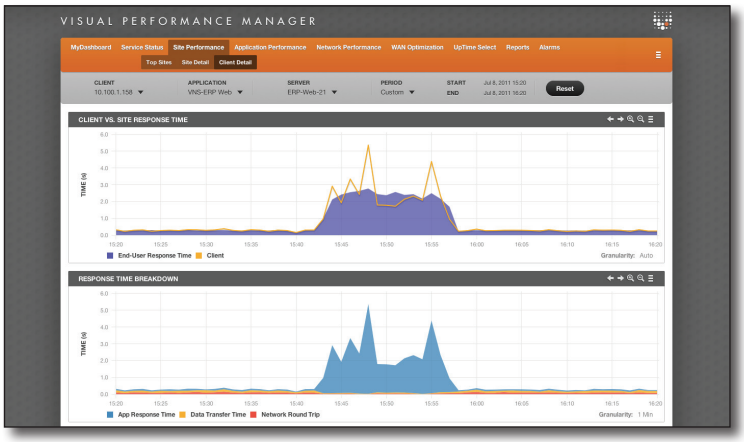

Managing end user experience with Visual Performance Manager. Quickly isolate user problem at a remote site.

- **Future proof and scalable architecture** - As with any new technology, you will need new information that you do not know about today. The chosen solution needs to be extensible to support new relevant performance metrics without having to do a mass rip and replace. A proven scalable architecture is important especially if you are managing many remote offices. For the IT team to be effective, the architecture needs to be able to support mediating a variety of data sources in your delivery infrastructure and correlating performance metrics to provide a comprehensive view of application performance.

- **Establish service level agreements with your Cloud provider** - For Enterprise IT using hybrid cloud environments, in addition to visibility of the application performance in your private cloud, you should be demanding that your Cloud service provider establish service level agreements and prove that application services are delivered according to availability and performance objectives.

SECURE & CONTROL ANY CLOUD ENVIRONMENT

The Vyatta Network OS is industry's most complete virtual networking & security solution. Optimized for VMware, XenServer, Red Hat KVM and in the Amazon Cloud, Vyatta enables you to segment, isolate and secure applications and data in any cloud environment.

- Stateful Firewall
- IPSec and SSL VPN
- Intrusion Prevention
- Dynamic Routing
- Web Filtering
- Management API

...and more

VYATTA®

http://www.vyatta.com