The 2012
Cloud Networking Report

Executive Summary

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

Platinum Sponsors:

Gold Sponsors:

Produced by:

Executive Summary

Background

On a going forward basis, IT organizations are expected to spend significantly more money on cloud computing initiatives than they are on other types of IT initiatives. Throughout this report, the phrase ***cloud networking*** refers to the LAN, WAN and management functionality that must be in place to enable the ongoing adoption of cloud computing. As is discussed in this report, in order to support cloud computing, a cloud network must be dramatically more agile and cost effective than a traditional network is. To help IT organizations deploy a network that can enable cloud computing, the primary goal of this report is to describe the challenges and solutions that are associated with cloud networking.

The first section of this report will identify what cloud computing is today and will also describe how cloud computing is likely to evolve in the near term. Subsequent sections focus on the key components of a cloud network: Data Center LANs, WANs, and Network Management. There is also a separate section on Software Defined Networking (SDN). This year's edition of the cloud networking report leverages last year's edition of the report[1]. However, every section of The 2011 Cloud Networking Report has been significantly updated to reflect the changes that have occurred in the last year.

As noted, the primary goal of this report is to describe the challenges and solutions that are associated with cloud networking. A secondary goal of this report is to identify how IT organizations are currently approaching cloud networking and where possible, indicate how that approach is changing. To accomplish that goal, this report includes the results of surveys that were recently given to the subscribers of Webtorials.com.

The Emergence of Cloud Computing and Cloud Networking

The goal of cloud computing is to enable IT organizations to achieve a dramatic improvement in the cost effective, elastic provisioning of IT services that are good enough. In order to demonstrate the concept behind the phrase ***good enough***, consider just the availability of an IT service. In those cases in which the IT service is business critical, *good enough* could mean five or six 9's of availability. However, in many other cases *good enough* has the same meaning as *best effort* and in these cases *good enough* could mean two or three 9's of availability if that approach results in a notably less expensive solution.

In most instances the SLAs that are associated with public cloud computing services such as Salesforce.com are weak and as such, it is reasonable to say that these services are delivered on a best effort basis. For example, most of the SLAs that are associated with public cloud computing services don't contain a goal for the end-to-end performance of the service in part because these services are typically delivered over the Internet and no provider will give an end-to-end performance guarantee for the Internet. While this situation will not change in the near term, as discussed in the WAN section of this report, there are technologies and services that can improve the performance of the Internet.

[1] http://www.webtorials.com/content/2011/11/2011-cloud-networking-report.html

Over the next year, the interest that IT organizations have in acquiring applications from Software-as-a-Service (SaaS) providers will grow significantly and will include the increased use of applications such as project and portfolio management, office productivity and collaboration. IT organizations are also beginning to make use of cloud computing service providers (CCSPs) for a number of applications that have historically been provided by IT organizations. This includes unified communications, VoIP, network management and network optimization.

The Infrastructure-as-a-Service (IaaS) market is going through some significant transformations. The initial set of IaaS solutions that were brought to market was basic compute and storage services. However, many IaaS providers are deploying myriad new services including:

- Disaster Recovery
- Virtual Private Data Centers
- High Performance Computing

In part because of the changes that are occurring in the IaaS market, the survey data indicates that roughly half of all IT organizations are currently in the process of developing a strategy for how they will use public and private IaaS solutions. As IT organizations develop those strategies, their concern about the security and confidentiality of data is the primary impediment to the broader adoption of both public and private IaaS solutions.

The survey data indicates that IT organizations expect that the IaaS solutions that they acquire will come with a wide variety of supporting network services, including load balancers, firewalls and IDS/IPS functionality. The most importance criterion that IT organizations use to evaluate those network services is the agility of the network service itself and the ability of the service to enable the agility of the IaaS solution.

A form of hybrid cloud computing that is growing in importance is cloud balancing. The phrase ***cloud balancing*** refers to routing service requests across multiple data centers based on myriad criteria. The advantages of cloud balancing are that it enables IT organizations to maximize performance, minimize cost and manage risk. Some of the challenges that are associated with cloud balancing are discussed in the WAN section of this report.

As much as cloud computing is about technologies, it is also about changing the culture of the IT organization. One of the cultural shifts that is associated with the adoption of cloud computing is that IT organizations become less of a provider of IT services and more of a broker of IT services. In their role as a broker of IT services, IT organizations can facilitate contract negotiations with CCSPs. IT organizations can also ensure that the acquired application or service doesn't create any compliance or security issues, can be integrated with other applications as needed, can scale, is cost effective and can be managed.

The Emerging Data Center LAN

One of the key characteristics of a traditional data center LAN is that it was usually designed around a three-tier switched architecture comprised of access, distribution and core switches. These LANs were also characterized by the use of the spanning tree protocol to eliminate loops, the use of Ethernet on a best effort basis and the separation of the data network from the storage network. Today, a number of factors are causing IT organizations to rethink their approach to data center LAN design. One of the primary factors driving change in the data center LAN is the ongoing virtualization of servers. Server virtualization creates a number of

challenges including the requirement to manually configure parameters such as QoS settings and ACLs in order to support the dynamic movement of VMs.

The deployment of server virtualization is just one of the on-going IT initiatives that are aimed at improving the cost-efficiency of the enterprise data center. In many cases these initiatives place a premium on IT organizations being able to provide highly reliable, low latency, high bandwidth communications among both physical and virtual servers. Whereas the hub and spoke topology of the traditional data center LAN was optimized for client-to-server communications, it is decidedly sub-optimal for server-to-server communications. As discussed in this report, one approach for improving server-to-server communications is to flatten the network from three tiers to two tiers consisting of access layer and aggregation/core layer switches. The survey data contained in this report indicates that there is significant desire on the part of IT organizations to flatten their data center LANs, but that there is also significant uncertainty relative to how flat those LANs will become in the next two years.

One of the key design considerations relative to the next generation data center LAN is what technologies, if any, will IT organizations use to replace the spanning tree protocol (STP), as this protocol only allows for a single active path between any two network nodes. One way to avoid the limitations of STP is to use switch virtualization and multi-chassis Link Aggregation Group (MC LAG) technologies. With switch virtualization, two or more physical switches are made to appear to other network elements as a single logical switch or virtual switch. MC LAG is not the only alternative to STP. One of the other alternatives is TRILL (Transparent Interconnection of Lots of Links), which is based on an Internet Engineering Task Force project to develop a Layer 2 shortest-path first routing protocol for Ethernet. Another alternative is Shortest Path Bridging (SPB). This protocol was defined by the IEEE 802.1aq working group which was chartered with defining a standard for the shortest path bridging of unicast and multicast frames and which supports multiple active topologies.

As mentioned, one of the characteristics of the current generation of data center LANs is the separation of the data and storage networks. However, a possible characteristic of the next generation of data center LANs will be the convergence of block-level storage and data traffic over a common high-speed Ethernet data center switching fabric. Traditional Ethernet, however, only provides a best effort service that relies on upper level protocols such as TCP to manage congestion and to recover lost packets through re-transmissions. In order to emulate the lossless behavior of a Fibre Channel (FC) SAN, Ethernet needs enhanced flow control mechanisms that eliminate buffer overflows for high priority traffic flows, such as storage access flows. Lossless Ethernet is based on a set of emerging standards, which are commonly referred to as IEEE Data Center bridging (DCB).

One of the challenges facing IT organizations as they attempt to deploy a flatter data center LAN is the scalability of the data center LAN architecture. The scalability of a data center LAN architecture is determined by the number of server ports that can be supported with a given level of redundancy and over-subscription at different points within the LAN topology. Many of the data center LANs that are being deployed today are based on a two-tier design that provides high levels of redundancy and low over-subscription levels for server-to-server traffic. This report develops a model for two tier switched LANs that takes into account both connections for redundancy and connections to the LAN core. IT organizations can use this model to estimate the TCO of alternative data center LAN designs.

As was also mentioned, one of the primary factors that is driving IT organizations to redesign their data center LANs is the requirement to support server virtualization. In particular, when

virtual machines (VMs) are migrated between servers, the network has to accommodate the constraints imposed by the VM migration utility; e.g., VMotion. Typically the VM needs to be on the same VLAN when migrated from source to destination server. An emerging approach that addresses some of the major limitations of live migration of VMs across a data center network is some form of network virtualization. Currently, the most common approach to automating the manual processes involved in VM provisioning and migration is based on communication between the Hypervisor Management system and the switch element management system (EMS) via APIs supported by both vendors. This type of solution is commonly referred to as Edge Virtualization.

One approach to edge virtualization is the Distributed Virtual Switch (DVS). With DVS, the control and data planes of the embedded hypervisor vSwitch are decoupled. This allows the data planes of multiple vSwitches to be controlled by an external centralized management system that implements the control plane functionality. Another approach is the IEEE 802.1Qbg standard that addresses both edge virtualization and some of the potential issues with vSwitches. This standard includes Edge Virtual Bridging (EVB) in which all the traffic from VMs is sent to the physical network access switch. If the traffic is destined for a VM on the same physical server, the access switch returns the packets to the server over the same port on which it was received.

However, most protocols for network virtualization are based on creating virtual network overlays using tunneling and encapsulation techniques. This includes the Virtual eXtensible LAN (VXLAN), the Network Virtualization using Generic Router Encapsulation (NVGRE) and the Stateless Transport Tunneling (STT) protocols.

VXLAN[2] virtualizes the network by creating a Layer 2 overlay on a Layer 3 network via MAC-in-UDP encapsulation. The VXLAN segment is a Layer 3 construct that replaces the VLAN as the mechanism that segments the data center LAN for VMs. The VXLAN segment has a 24 bit VXLAN Network identifier and VXLAN is transparent to the VM, which still communicates using MAC addresses. NVGRE[3] uses the GRE tunneling protocol defined by RFC 2784 and RFC 2890. NVGRE is similar in most respects to VXLAN with two major exceptions. While GRE encapsulation is not new, most network devices do not parse GRE headers in hardware, which may lead to performance issues and issues with 5-tuple hashes for traffic distribution in multi-path data center LANs. The other exception is that the current IETF NVGRE draft does not address the control plane question, leaving that for a future draft or possibly as something to be addressed by (Software Defined Networking) SDN controllers. STT[4] is a third overlay technology for creating Layer 2 virtual networks over a Layer 2/3 physical network within the data center. Conceptually, there are a number of similarities between VXLAN and STT. However, STT encapsulation differs from NVGRE and VXLAN in two ways. First, it uses a stateless TCP-like header inside the IP header, which allows tunnel endpoints within end systems to take advantage of TCP segmentation offload (TSO) capabilities of existing TOE server NICs. STT also allocates more header space to the per-packet metadata, which provides added flexibility for the virtual network control plane.

[2] http://searchservervirtualization.techtarget.com/news/2240074318/VMware-Cisco-propose-VXLAN-for-VM-mobility
[3] http://tools.ietf.org/html/draft-sridharan-virtualization-nvgre-00
[4] http://tools.ietf.org/html/draft-davie-stt-01

Security is generally considered by enterprise IT departments to be the primary concern in today's highly virtualized data centers and in the implementation of private or public cloud computing environments. In the traditional data center, internal security has generally been implemented by deploying dedicated physical security appliances at the aggregation layer of a 3-tier or a 2-tier network. This approach has been successful in relatively static non-virtualized environments that require infrequent changes to the location and configuration of both servers and physical security appliances. With the advent of server virtualization and the dynamic migration of workloads within and between data centers, there is a growing need to make the workload's complete security environment as easily provisioned and migrated as the VMs themselves. In addition to being dynamic and virtualization-aware, the security solution needs to be both scalable and automated to the degree possible.

One way to achieve this goal is to deploy a virtualized physical security appliance that can support a large number of instances of virtual security devices, such as firewalls, IDS/IPS, WAF, etc. Potentially these instances could be implemented as VMs running on the security device's hypervisor. This type of integrated security device can also include its own physical Layer 2 and Layer 3 switching functionality, which allows the device to be installed in line between the access and aggregation layers of the physical data center LAN. The VLANs used by the virtualized servers are trunked to the virtualized security appliance via the hypervisor vSwitches and the physical access switches.

Software Defined Networks (SDN)

As is typical of emerging technologies and new approaches to networking, there is currently somewhat of a broad definition relative to how the industry, particularly vendors, define SDN. The most common way that SDN is described is based on a layered architecture as shown in **Figure 1**. In **Figure 1**, the control plane function is centralized in SDN Controller software that is installed on a server or on a redundant cluster of servers for higher availability and performance. The SDN controller is used to control the actions of the subtending networked elements.

Figure 1: **Software Defined Network Controller**

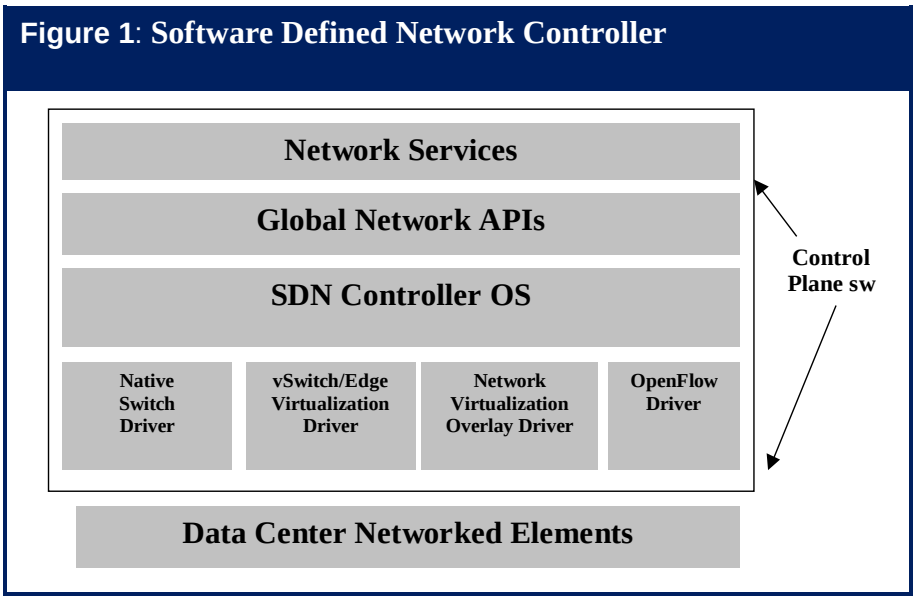

Most of the discussion of SDN includes the use of OpenFlow. OpenFlow is an open protocol between a central SDN/OpenFlow controller and an OpenFlow switch that can be used to program the forwarding behavior of the switch. Using pure OpenFlow switches, a single central controller can program all the physical and virtual switches in the network. All of the control functions of a traditional switch (e.g. routing protocols that are used to build forwarding

information bases (FIBs)) are run in the central controller. As a result, the switching functionality of the OpenFlow switch is restricted entirely to the data plane,

The organization that is most associated with SDN is the Open Networking Foundation (ONF). The ONF was launched in 2011 and has as its vision to make OpenFlow-based SDN the new norm for networks. To help achieve that vision, the ONF has taken on the responsibility to drive the standardization of the OpenFlow protocol. Unlike most IT standards groups or industry consortiums, the ONF was not by founded by suppliers of the underlying technologies, but by Deutsche Telekom, Facebook, Google, Microsoft, Verizon, and Yahoo! As such, the ONF is one of the very few IT standards groups or industry consortiums that were launched by potential users of the technologies on which the consortium focused.

One of the primary benefits of OpenFlow is the centralized nature of the FIB. This centralization allows optimum routes to be calculated deterministically for each flow leveraging a complete model of the end-to-end topology of the network. Based on an understanding of the service levels required for each type of flow, the centralized OpenFlow controller can apply traffic engineering principles to ensure each flow is properly serviced. The result can be much better utilization of the network without sacrificing service quality. Centralized route processing also allows the pre-computation of a set of fail-over routes for each possible link or node failure.

Most of the discussion in the industry about SDN focuses on its use in data centers. However, Google has implemented SDN in their G-Scale WAN, which is the WAN that links Google's various global data centers. The G-Scale WAN is a prime example of a production OpenFlow Layer 3 WAN that is realizing the benefits of FIB centralization.

The Wide Area Network

After a lengthy period in which the WAN underwent repeated fundamental change, there are currently no fundamental changes in store for the WAN. So on a going forward basis, IT organizations need to plan for WAN evolution based on the assumption that at least for the next few years, their WAN will be comprised primarily of intelligence added on top of two WAN services: MPLS and the Internet.

Driven by the adoption of initiatives such as cloud computing, virtual machine migrations, virtual desktops and collaboration, the amount of traffic that transits the typical WAN grows significantly each year. The WAN, however, doesn't follow Moore's Law and as a result, the price / performance of WAN services such as MPLS tends to improve by only a relatively small amount each year. The result of these two factors is that for most companies the cost of the WAN increases on an annual basis.

As previously discussed, the goal of cloud computing is to enable IT organizations to achieve a dramatic improvement in the cost effective, elastic provisioning of IT services that are good enough. In a growing number of instances, Internet-based VPNs that use DSL for access are *good enough* to be a cloud network. A somewhat related shift in terms of how IT organizations design their WAN to support cloud-based services is that in a growing number of instances IT organizations will avoid backhauling Internet traffic and will instead implement distributed access to the Internet from their branch offices.

One of the key trends in network and application optimization is the deployment of virtual appliances; e.g., virtual WAN Optimization Controllers (vWOCs) and virtual Application Delivery

Controllers (vADCs). One of the compelling advantages of a virtualized appliance is that the acquisition cost of a software-based appliance can be notably less than the cost of a hardware-based appliance with same functionality. Another benefit of virtualized appliances is that in many instances the benefits of the dynamic movement of a VM from one server to another are maximized if the supporting infrastructure, including the WOCs and ADCs, is virtualized and can also be dynamically moved. In addition, there is significant interest in placing a WOC on premise at an IaaS provider's data centers.

One of the ways that an IT organization can get better performance out of the Internet is by using an Internet overlay. An Internet overlay leverages service provider resources that are distributed throughout the Internet in order to optimize the performance, security, reliability, and visibility of the enterprise's Internet traffic. Another approach that improves the performance and availability of the Internet is to combine multiple ISP connections and to share traffic over the connections using policy based routing (PBR).

Unfortunately PBR can be difficult to administer and manage and it also creates only a static allocation of WAN capacity. In order to overcome these limitations, another way that an IT organization can better leverage the Internet is by implementing an aggregated virtual WAN (avWAN). This technology enables IT organizations to implement WANs based on multiple WAN services (e.g., MPLS, Frame Relay and the Internet) and/or WANs based on just multiple Internet VPN connections. An aggregated virtual WAN transcends simple PBR by dynamically recognizing application traffic and allocating traffic across multiple paths through the WAN based on real-time traffic analytics.

As previously mentioned, cloud balancing provides a lot of benefits. There are, however, a number of challenges associated with cloud balancing. For example, the VLANs within which VMs are migrated must be extended over the WAN between and amongst the private and public data centers. This involves the creation of an overlay network that allows the Layer 2 VLAN traffic to be bridged or tunneled through the WAN. In addition, application performance must meet user expectations regardless of the location of the users or the IT resources that the users are accessing. This means that the public cloud data centers need to offer the same WAN optimization and application acceleration capabilities that are deployed within the enterprise.

The two biggest concerns that IT organizations have with the use of MPLS are its cost and the amount of time it takes to implement new circuits. An emerging WAN service, referred to as Network-as-a-Service (NaaS), is intended to avoid those concerns. NaaS is built using a core network that interconnects a distributed set of Points of Presence (POPs). The phrase ***NaaS*** implies that unlike MPLS, the service can be deployed rapidly – typically within a day by leveraging Internet links for the first and last mile connections while providing a reliable private core network and additional network intelligence. The service also allows IT organizations to add capacity on demand, rather than provisioning and paying for bandwidth to support future requirements. Another key feature of a NaaS is that it should allow a customer to quickly upgrade to add the optimization capabilities discussed in the following paragraph.

As previously mentioned, it is now possible for IT organizations to acquire network optimization from a CCSP. In this situation, instead of a physical or virtual WOC at each site, the WOC functionality is provided at the CCSP's cloud data centers or POPs, which ideally are in close proximity to the enterprise users, the data centers and the providers of other cloud services. The PoPs are interconnected by the CCSP's core network with customer access to each PoP provided via the Internet or via an enterprise WAN service.

Somewhat of a new class of WAN product is cloud optimized WOCs. These are purpose-built virtual WOC appliances for deployment in public cloud environments. Cloud optimized features include compatibility with cloud virtualization environments, SSL encryption and acceleration, and automated migration or reconfiguration of virtual WOCs in conjunction with VM provisioning or migration.

Another emerging class of product is hypervisor–based multi-tenant ADC Appliances. Partitioned ADC hardware appliances have for some time allowed service providers to support a multi-tenant server infrastructure by dedicating a single partition to each tenant. Enhanced tenant isolation in cloud environments can be achieved by adding hypervisor functionality to the ADC appliance and by dedicating an ADC instance to each tenant. Each ADC instance is then afforded the same type of isolation as a virtualized server instance, with protected system resources and address space.

Management & Security

Until recently, IT management was based on the assumption that IT organizations performed tasks such as monitoring, baselining and troubleshooting on a server-by-server basis. Now, given the widespread adoption of server virtualization, the traditional approach to IT management must change to enable management tasks to be performed on a VM-by-VM basis. Another assumption that underpinned the traditional approach to IT management was that an application resided on a given server, or set of servers, for very long periods of time. However, part of the value proposition that is associated with server virtualization is that it is possible to migrate VMs between physical servers. This ability to migrate VMs between physical servers is just one example of the fact that IT organizations need to adopt an approach to management that is based on the assumption that the components of a service, and the location of those components, can and will change frequently.

Part of the overall management challenge associated with any form of cloud computing are the challenges discussed in the preceding paragraph. In addition, a fundamental issue relative to managing either a public or hybrid cloud computing service is that the service has at least three separate management domains: the enterprise, the WAN service provider(s) and the various cloud computing service providers. As a result, IT organizations need to adopt an approach to IT management that is based on gathering management data across myriad data centers, including ones that are owned and operated by a third party.

The initial set of Network Performance Management Systems (NPMS) worked acceptably well for traditional client/server applications and other centrally hosted applications. One of the limitations of these systems is that they measured performance across the entire path, but did not isolate which network segments had performance issues. One of the challenges associated with the traditional approach to application performance management is that it was typically performed separately from network performance management. Since these tasks are typically done by different parts of the IT organization using different tool sets and management frameworks, it is quite common that conflicting answers are given for the source of application performance issues.

In the traditional approach to IT management, one set of tools is used to manage enterprise data applications and a different set of tools is used to manage voice and video traffic. That approach is expensive and leads to a further hardening of the technology domains, which then leads to a lengthening of the time it takes to resolve problems. The reality for most IT

organizations is that voice and video traffic is becoming an increasing percentage of the overall traffic on their networks. This reality is one of the reasons why IT organizations need to adopt an approach to management in which one set of tools is used to manage enterprise data applications as well as voice, video and complex interrelated applications.

According to the survey data, the majority of IT organizations believe that getting better at managing the inter-related applications that comprise a business service is either very or extremely important. In order to successfully respond to this pressure, IT organizations need to adopt an approach to service management that enables them to holistically manage the four primary components of a service:

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

In addition, IT organizations should adopt an approach to service delivery management that is unified across the various IT domains so that IT organizations have visibility across all of the applications, services, locations, end users and devices. Among other advantages, this approach will enable IT organizations to overcome the previously mentioned limitations of the traditional approach to application performance management.

There are a number of services and technologies that IT organizations can use to manage the applications and services that they get from a CCSP. One such class of service was previously mentioned – a cloud-based network management service. Another technology that can help IT organizations to manage the applications and services that they get from a CCSP is a highly scalable and integrated DNS/DHCP/IPAM solution, which is also well integrated with the virtual server management system.

An increasingly popular approach to building cloud data centers is based on pre-integrated and certified infrastructure packages from a broadly-based IT equipment vendor, a group of partners or a joint venture formed by a group of complementary vendors. These packages typically are offered as turn-key solutions and include compute, server virtualization, storage, network, and management capabilities. Management systems for converged infrastructure typically support APIs for integration with other management systems that may be currently deployed in order to manage the end-to-end data center. These APIs can provide integration with enterprise management systems, automated service provisioning systems, fault and performance management systems and orchestration engines.

Service orchestration is another 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.

A key component of application performance management is the ability to perform root cause analysis. 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. For example, if IT organizations can effectively identify which components of the infrastructure

support a particular application or service, monitoring can much more easily identify when services are about to degrade due to problems in the infrastructure. As part of this approach, predictive techniques such as heuristic-based trending of software issues and infrastructure key performance indicators can be employed to identify and alert management of problems before they impact end users.

Ideally the issue of application performance would be addressed at all stages of an application's lifecycle, including multiple iterations through the design, implement, test, and operate phases as the application versions are evolved to meet changing requirements. However, the vast majority of IT organizations don't have any insight into the performance of an application until after the application is fully developed and deployed. In addition, the vast majority of IT organizations have little to no insight into how a change in the infrastructure, such as implementing server virtualization, will impact application performance prior to implementing the change. To overcome these issues, IT organizations need to develop more of a focus on Application Performance Engineering, which is the practice of first designing for acceptable application performance and then testing, measuring and tuning performance throughout the application lifecycle.

Over the last several years the sophistication of hackers has increased by an order of magnitude. Many of the new generation of sophisticated attacks are focusing on vulnerabilities in mobile devices, social media and cloud computing. In order to respond to these attacks, IT organizations have on average implemented 4.8 network security systems. That said, almost half of all IT organizations either don't have a data classification policy or they have one that isn't used or enforced. In addition, just over half of all IT organizations don't use full disk encryption on PCs and in the majority of instances, network security and application security are architected, designed and operated separately.

According to the survey data, over a quarter of IT organizations either currently acquires security functionality from a CCSP or they expect that they will within the next year. A cloud-based security service needs to be able to allow access to a social media site such as Facebook, but block specific activities within the site, such as gaming or posting. Analogously, the service needs to have the granular controls to be able to allow users to send and receive mail using a provider such as Yahoo, but block email attachments.

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

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

In order to be effective, a cloud-based security service that provides Web application firewall functionality needs to be deployed as broadly as possible, preferably in tens of thousands of locations. A cloud-based security service that provides Web application firewall functionality is

complimentary to a premise-based Web application firewall. That follows because while the cloud-based Web application firewall service can perform many security functions that cannot be performed by an on premise Web application firewall, there are some security functions that are best performed by an on premise Web application firewall.

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.

Customer Driven Innovation

Cost Effective Cloud Networking | Virtualization as the Enabler

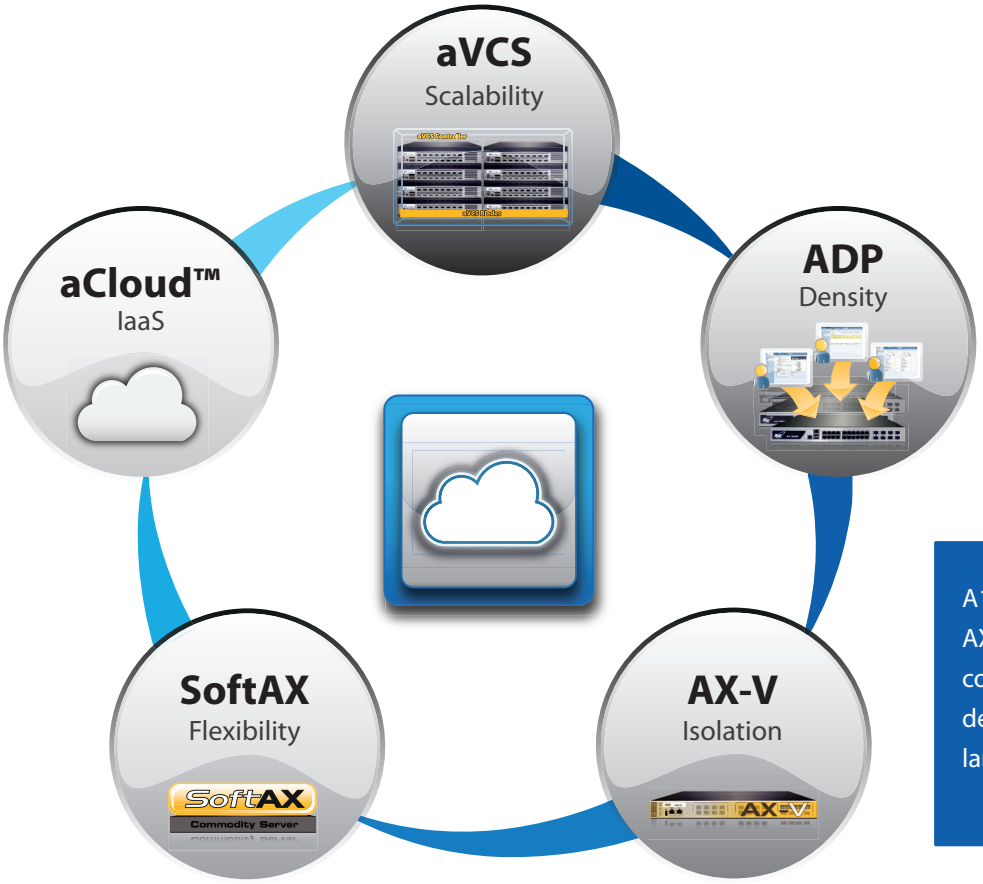

A10 also offers a powerful choice of AX Series ADC form factors with comprehensive management options, delivering flexibility and efficiency for large scale deployments.

AX Series Virtualization Products & Solutions

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

SoftAX™

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

AX-V Appliance

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

AX Virtual Chassis System (aVCS™)

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

Application Delivery Partitions (ADPs)

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

The Application Fluent Data Center Fabric

Introduction

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

Application Fluency for the Data Center

Resilient Architecture

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

Automatic Controls

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

Streamlined Operations

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

The Alcatel-Lucent Mesh

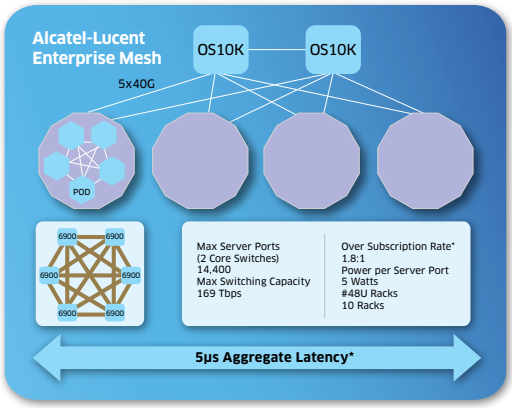

Alcatel-Lucent provides a unique Application Fluent approach to maximize the benefit from virtualization technologies for servers, the desktop, as well as the network. Alcatel-Lucent's application fluent data center fabric can scale from several hundred to over 14,000 server facing ports while keeping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

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

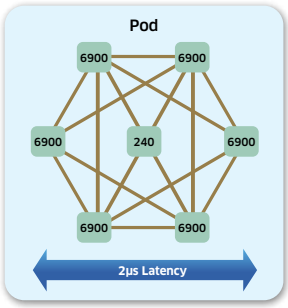

Application fluency in the corporate data center includes its transformation into a multi-site private cloud by extending layer 2 connectivity between data center sites and allowing for seamless delivery of public cloud-based services on the corporate network.

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

Open Ecosystems and Market Success

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

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

For More Information

Alcatel-Lucent Data Center Switching Solution
Alcatel-Lucent Application Fluent Networks
Alcatel-Lucent Enterprise

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

Alcatel·Lucent
Enterprise

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

How to Lower Networking Costs and Safely Improve Performance

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

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

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

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

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

WebPulse is based on sound analysis-system design principles:

- Massive input: WebPulse analyzes up to 1 billion web requests per day.
- In-depth analysis: 16 layers of analysis support over 80 categories in 55 languages.
- Granular policy: Up to 4 categories can be applied to each web request for multi-dimensional ratings.
- Speed: Automated systems process inputs – in most cases, in real time.
- Results: This collective intelligence allows WebPulse to categorize new Internet applications and websites quickly to PacketShaper without software updates/upgrades.

BYOD BANDWIDTH CONSUMPTION - JUNE 2011 TO JUNE 2012

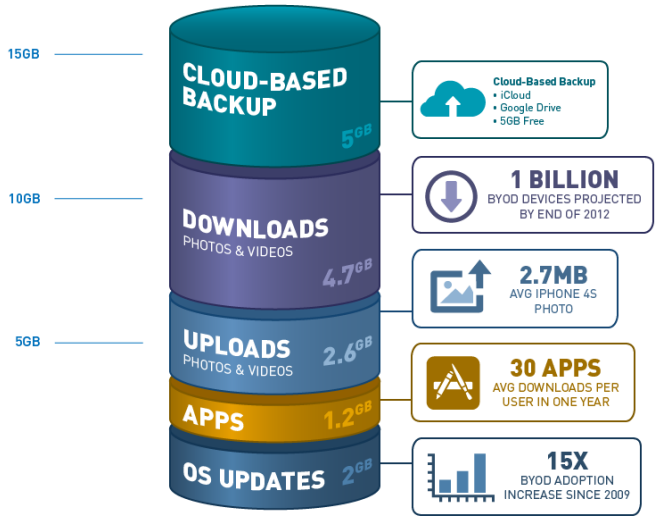

The graphic details the impact of BYOD and Recreational video traffic can have on a network if left unchecked.

Second: Optimize Performance

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

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

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

On The Road To The Cloud?

agility made possible™

ca technologies

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

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

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

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

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

Speeding proactive triage and remediation with less effort

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

Meeting massive scalability demands cost-effectively

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

Shifting operations costs to innovation

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

Improving revenue streams

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

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

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

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

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

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

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

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

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

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

The Power of Cloud for the Enterprise

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

Defining Cloud

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

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

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

CISCO VIRTUAL NETWORK PORTFOLIO

Routing and Switching

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

Security and VPN

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

WAN Optimization

- Cisco Virtual Wide Area Application Services (vWAAS)

Network Analysis and Monitoring

- Cisco Prime Virtual Network Analysis Module (NAM)

Application Delivery Controllers

- Citrix NetScaler VPX virtual application delivery controller

Virtual Services Deployment Platform

- Cisco Nexus 1100 Series Virtual Services Appliance

Cloud Orchestration and Management

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

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

Barriers to Adoption

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

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

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

Practical Solutions for Cloud-ready Virtual Networks and Infrastructure

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

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

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

Where to Begin Your Cloud Journey

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

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

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

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

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

For More Information

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

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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

Application Performance for Business Efficiency

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

82%* of organizations suffer application performance problems.

63%* of organizations don't know the number of apps using the network.

72%* of organizations use very occasionally their network to its full data transmission capacity.

Business and IT performance are tightly coupled…

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

**Ipanema Killer Apps survey 2012*

IT departments are witnessing change at a pace never seen before

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

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

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

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

Business efficiency requires guaranteed application performance

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

With Ipanema, control all your IT transformations!

What our customers say about us:

Do more with less

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

Guarantee Unified Communications and increase network capacity

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

Reduce costs in a cloud environment

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

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

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

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

Enabling the cloud:

Award-winning NEC ProgrammableFlow® Open Software Defined Networking…

…delivering automated, efficient, and agile networks for the cloud

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

Specific functions customers prize include:

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

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

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

Empowered by Innovation NEC

® ProgrammableFlow is a registered trademark of NEC Corporation

Expand Your Cloud Offering with Advanced Cloud ADC Solutions

Challenges in the Cloud Provider Business

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

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

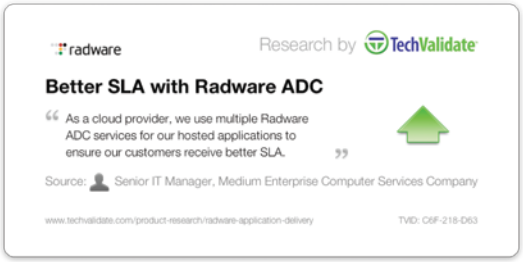

Above all, cloud providers are pressed to build solutions with minimal capital expenditure, maintain low operational costs and rapidly meet spikes in customer demand. Flexible procurement models by vendors and platforms that are easily scalable and centrally managed support the overall operational constraints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

Radware Solutions for Cloud Service Providers

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

- **Radware ADC-VX™** – highly scalable ADC virtualization and consolidation solution offering high speed global and local load balancing, application acceleration and SSL offloading that supports dynamic availability requirements of cloud customers. ADC-VX can host multiple fully isolated, fully featured vADC instances.

- **Radware Alteon VA®** – flexible virtual ADC instance running atop most commercial, general purpose x86 server hypervisors.

- **Radware VADI®** – comprehensive virtual application delivery infrastructure solution including Alteon VA and ADC-VX-based virtual ADCs (vADC) and vDirect, an ADC service automation plug-in that simplifies ADC service deployment in cloud environments.

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

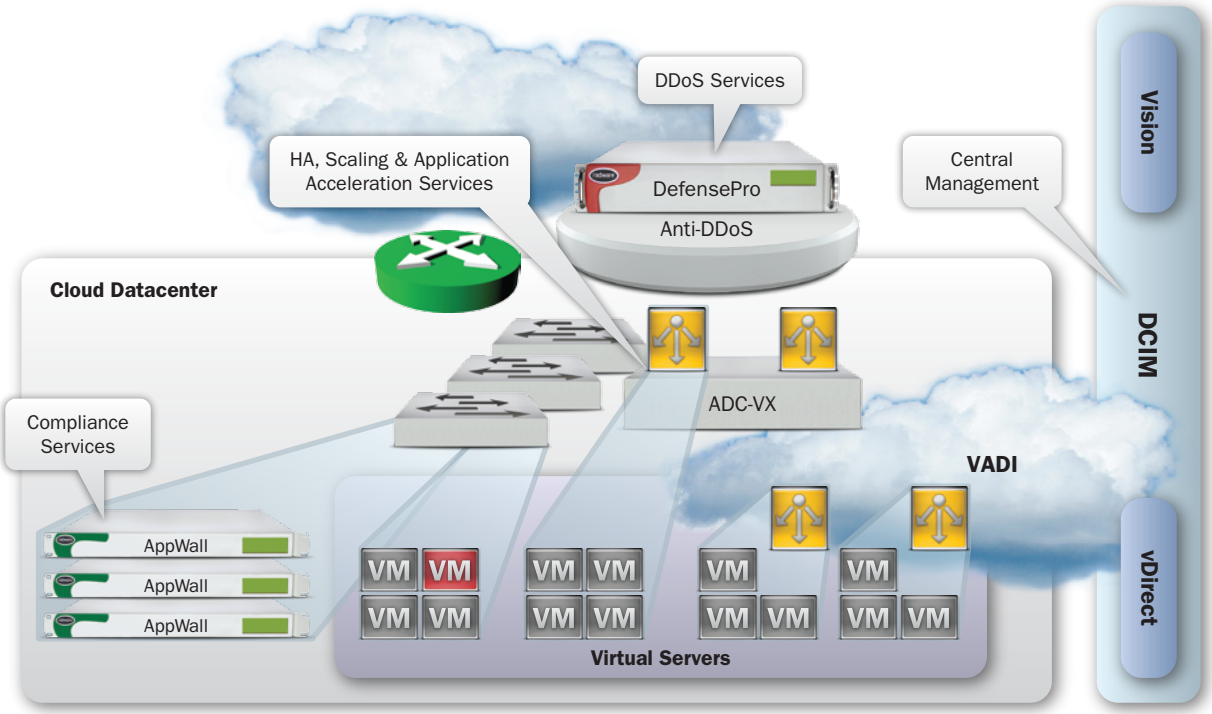

Figure 1 - Radware Service Architecture for Cloud

Benefits of Radware Solutions for Cloud Service Providers

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

Summary

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

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

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