The 2010 Cloud Networking Report

Dr. Jim Metzler, Ashton Metzler & Associates

Cloud Networking & Cloud Computing

The Emergence of Cloud Networking

The majority of IT organizations have either already adopted, or are in the process of evaluating the adoption of one or more classes of cloud computing. The broad interest in cloud computing is understandable given that, as explained in the next section of this report, the goal of cloud computing is to enable IT organizations to be dramatically more agile and cost effective. The adoption of cloud computing, however, creates some very significant networking challenges. Below are some examples of these challenges.

Manual Network Reconfiguration to Support Virtual Machine (VM) Migration

Many of the benefits of cloud computing depend on the ability to migrate VMs among physical servers located in the same data center or in geographically separated data centers. The task of moving a VM is a relatively simple function of the virtual server management system. There can, however, be significant challenges in assuring that the VM's network configuration state (including QoS settings, access control lists (ACLs), and firewall settings) is also transferred to the new location. In the vast majority of instances today, making these modifications to complete the VM transfer involves the time-consuming manual configuration of multiple devices.

Regulatory compliance requirements can further complicate this task. For example, assume that the VM to be transferred is supporting an application that is subject to PCI compliance. Further assume that because the application is subject to PCI compliance that the IT organization has implemented logging and auditing functionality. In addition to the VM's network configuration state, this logging and auditing capability also has to be transferred to the new physical server.

The Performance and Cost of the Wide Area Network (WAN)

As explained in the next section, one of the primary characteristics of a cloud computing solution is the centralization of IT resources; e.g., desktops, applications, servers, data and storage resources. That centralization drives notably more traffic over the WAN and creates a number of challenges relative to the performance and cost of the WAN. Relative to the performance of the WAN, the centralization of IT resources will not be regarded as being successful if the users' experience with accessing those resources is significantly degraded compared to what it was when they accessed those resources locally. Relative to the cost of the WAN, the combination of technological advances, Moore's Law and a competitive marketplace have continually reduced the unit cost that IT organizations pay for most of the major components of IT including processing, storage, and LAN bandwidth. Hence it is often possible for an IT organization to support a significant increase in storage requirements without a dramatic increase in cost. Unfortunately those factors have had little impact on the unit cost of traditional private WAN services such as Multi-Protocol Label Switching (MPLS). As a result, the cost of a significant increase in the use of the WAN will not be offset by a reduction in the unit cost of traditional WAN services. This could potentially lead to a situation in

which IT organizations adopt cloud computing in part to reduce cost, but end up significantly increasing the cost of the WAN.

The Focus on Services

Part of the shift that is occurring as part of the adoption of cloud computing is the growing emphasis on everything as a service (XaaS). In many cases an application and a service are the same thing. However, in a growing number of instances a service is comprised of multiple inter-related applications. Enterprise Resource Planning (ERP) is an example of a complex, business critical IT service that is comprised of multiple inter-related applications that support functions such as product lifecycle management, supply chain management and customer relationship management (CRM). A service, however, doesn't have to be an application or a combination of applications. A service can also be one of the key components of IT such as storage or computing.

On a going forward basis, a service will increasingly be supported by an infrastructure that is virtual and that is ***dynamic***. By ***dynamic*** it is meant that the service can be provisioned or moved in a matter of seconds or minutes. One example of this phenomenon is that compute services have already become virtual and dynamic. Unfortunately, the dynamic nature of creating and moving services creates significant management challenges. For example, the first generation of virtual switches (vSwitches) that resides inside of virtualized servers don't have the same traffic monitoring features as do physical access switches. This limits the IT organization's ability to do security filtering, performance monitoring and troubleshooting within virtualized servers.

Another example of the management challenges associated with the growing deployment of dynamic services is that due to the combination of the dynamic nature of IP and the meshed nature of enterprise networks, it is often not possible in a traditional IT environment to know what path the traffic took from origin to destination. This lack of knowledge complicates many critical tasks, such as troubleshooting, and this results in network organizations not being able to ensure acceptable application performance. The difficulty of knowing the path from origin to destination is greatly increased in a cloud environment because services can be dynamically moved between servers both within and between data centers.

This report will describe the challenges and solutions that are associated with cloud networking.

> ***The phrase cloud networking refers to the LAN, WAN and management functionality that must be in place to enable cloud computing.***

Included in this report are the results of surveys given to the subscribers of Webtorials.com and to the attendees of the Interop conferences. Throughout this report, those two groups of respondents will be respectively referred to as The Webtorials Respondents and The Interop Respondents.

Cloud Computing

As noted, cloud networking is heavily influenced by cloud computing. This section will identify some of the key concepts that are associated with cloud computing and will then use those concepts to characterize a cloud network. More information on cloud computing can be found in two reports: ***A Guide for Understanding Cloud Computing*** and ***Cloud Computing: A Guide to Risk Mitigation***.

Goal of Cloud Computing

Within the IT industry there is considerable confusion and disagreement relative to exactly what is meant by the phrase ***cloud computing***. An example of that confusion is the fact that the January 2009 edition of The Cloud Computing Journal published an article[1] that had twenty one definitions of cloud computing. The position taken in this report is that creating yet one more definition of cloud computing would only add to the confusion. This report also takes the position that it is notably less important to define exactly what is meant by the phrase *cloud computing* than it is to identify the goal of cloud computing.

> ***The goal of cloud computing is to enable IT organizations to achieve a dramatic improvement in the cost effective, elastic provisioning of IT services that are good enough.***

The phrase *good enough* refers in part to the fact that the SLAs that are associated with public cloud computing services such as Salesforce.com or Amazon's Simple Storage System are generally weak. For example, most of the SLAs don't contain a goal for the performance of the service. In addition, it is common to access these services over the Internet and nobody provides an SLA for the availability or performance of the Internet. As such, organizations that use these services do so with the implicit understanding that if the level of service they experience is not sufficient, their only recourse is to change providers. It may seem counter-intuitive that a company would utilize public cloud computing services for which SLAs are essentially non-existent. However, as described in section 5.0, two thirds of The Webtorials Respondents indicated that the SLAs that they receive from their network service providers are either not worth the paper they are written on, or that the SLAs they receive are not much better than nothing.

> ***SLAs from both traditional network service providers as well as public cloud computing providers are a work in progress.***

However, the phrase *good enough* should not be construed as meaning that a lower level of service is always acceptable. In some instances, *good enough* refers to a service that provides the highest levels of availability and performance.

> ***On a going forward basis, IT organizations will continue to need to provide the highest levels of availability and performance for a small number of key***

[1] Twenty-One Experts Define Cloud Computing, http://cloudcomputing.sys-con.com/node/612375

services. However, an ever-increasing number of services will be provided on a best effort basis.

Characteristics of Cloud Computing Solutions

The following set of bullets identifies the primary characteristics of cloud computing solutions. There is not, however, a litmus test to determine if a particular service is or is not a cloud computing service.

- **Centralization** of applications, servers, data and storage resources.
- Extensive **virtualization** of every component of IT, including servers, desktops, applications, storage, switches, routers and appliances such as WAN optimization controllers, application delivery controllers and firewalls.
- **Automation and Orchestration** of as many tasks as possible; e.g., provisioning, troubleshooting, change and configuration management.
- The **dynamic creation and movement of resources** such as virtual machines and the associated storage.
- Heavy reliance on the **network**.
- **Self-service** to allow end users to select and modify their use of IT resources without the IT organization being an intermediary.
- **Usage sensitive chargeback** that is often referred to as pay-as-you-go. An alternative is for IT organizations to show the consumption of IT resources by certain individuals or organizations; a.k.a., showback.
- **Simplification** of the applications and services provided by IT.
- **Standardization** of the IT infrastructure.
- **Technology convergence** such as the convergence of LAN and SAN and of switch and server.
- The development of **standards** that enable, among other things, the federation of disparate cloud computing infrastructures with one another (see below).
- The **federation** of disparate cloud computing infrastructures with one another.

Classes of Cloud Computing Solutions

Cloud Computing Service Providers (CCSPs) that provide their services either over the public Internet or over other WAN services are offering a class of solution that is often referred to as the *public cloud* or *public cloud computing.*

As described in the report entitled ***A Guide for Understanding Cloud Computing***, the primary types of services provided by CCSPs are:

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

Some IT organizations have decided to implement the characteristics of cloud computing solutions (e.g., virtualization, automation) within their internal IT environment. This approach is usually referred to as a ***Private Cloud***. In those instances in which an enterprise IT department uses a mixture of public and private cloud services, the result is often referred to as a ***Hybrid Cloud***.

Characterizing Cloud Networking

The following are the key characteristics of a cloud network.

1. **A cloud network has the same goal as cloud computing**. The goal of a cloud network is to be notably less expensive and dramatically more agile than traditional networks.

2. **A cloud network supports the characteristics of a cloud computing solution**. For example, a cloud network should support the dynamic creation and movement of IT resources.

3. As a minimum, **a cloud networking should do no harm**. It should not, for example, make the dynamic movement of resources or automated troubleshooting more difficult.

4. **A cloud network provides solutions that are good enough**. As was the case with cloud computing, in some cases *good enough* is a WAN with two 9's of availability. In other cases, it is a WAN with five 9's of availability.

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.Webtorias.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.

Leveraging Software-based Networking for End-to-End Cloud Infrastructures

VYATTA Open Networking

Industry Praise

Gartner

"Vyatta is certainly the headline name behind open-source networking"
Mark Fabbi
Gartner Inc. Analyst

Tolly

"Vyatta's open system running on standard hardware not only can scale better in enterprise and service provider edge deployments, but it also delivers enough headroom for expansion"
The Tolly Group

DELL

"Vyatta is able to provide the network services and secure connectivity that Dell GIS Cloud requires in a package that addresses the virtualization, commoditization and cost-benefit requirements of cloud computing."
Sanjay Basu
Dell Services

ZDNet

"Vyatta, has taught Cisco and the market that a networking box is really nothing but a computer with software."
Dana Blankenhorn
ZDNet

TechRepublic

"..anything you want to do with a standard Cisco router, you can do with Vyatta for the most part, and you don't have to worry about the various Cisco IOS licenses."
David Davis(CCIE, CCNA, CCNP)
TechRepublic

The Vyatta Network OS

The Vyatta network operating system is a scalable, integrated, enterprise-class networking solution that delivers advanced routing and network security functionality for physical, virtual and cloud networking environments. The Vyatta network OS includes dynamic routing, stateful firewall, VPN support, threat protection, traffic management and more in a package that is optimized to take advantage of multicore x86 processing power, common hypervisor platforms and emerging cloud architectures. All features are configured through Vyatta's familiar, networking-centric CLI, web-based GUI or third party management systems using the Vyatta Remote Access API.

The Power of Open Networking

Open and flexible networking is a requirement for today's evolving network. For the first time in two decades the industry is experiencing platform shifts that are dictating that networking be delivered as a software solution.

- **Datacenter Shifts:** Infrastructure shifting to the cloud requires flexible networking and security.
- **Virtualization:** Server and application consolidation requires virtualization-ready, platform independent application protection.
- **Edge Consolidation:** Special-purpose devices are giving way to multi-function, best-of-breed, multi-vendor integrated solutions.

The New Network Requirements

Features	Vyatta Network OS	Cisco IOS
Multifunction Layer 3+ (Routing, Firewall, VPN, IPS, Web Filter +)	Yes	Yes
Hardware Scalability	Seamless across x86 Cores	Cisco Limited
Software Performance	Unlimited	Platform Limited
Virtual Machine Availability	Yes (VMware, Xen, XenServer, KVM)	No
Open Management API	Yes	No
Integration into Custom Edge Devices	Yes	No
Cloud Readiness	Yes	No

The Vyatta Advantage

- **Network Right-Sizing:** As a single network OS that scales up and down to meet your requirements, Vyatta puts the freedom in your hands to right-size your network as needed. Using readily available off-the-shelf systems and components, Vyatta breaks the "box lock" model of proprietary hardware vendors and allows you to drive as little or as much performance as your network requires.

- **Hardware Price/Performance:** Standards have turned networking into a server workload. Today x86 hardware can easily outperform proprietary network devices at a small fraction of the cost. And the x86 universe means that faster systems at lower price are always on the horizon.

- **Virtualization:** Vyatta gives you the optional power of running networking functions as a virtual machine. Whether it's VMs at the network edge or VMs in the cloud datacenter, Vyatta radically increases your infrastructure flexibility and produces a substantially higher ROI than proprietary solutions.

Deploying Vyatta in the Cloud: Common Use Cases:

As cloud moves from vision to reality, networking quickly moves to the front as a major impediment to meeting these major requirements. The reason is simple: traditional networking infrastructure has not been modernized the way server and storage infrastructure has been over the past decade. While the business promise of cloud computing is broad, there are a few basic enabling themes underlying an effective cloud design:

- Highly dynamic, on-demand infrastructure
- Granular service control levels
- High infrastructure utilization (multi-tenancy)
- Elastic pricing

CLOUD INFRASTRUCTURE

Designing a network infrastructure for cloud computing should deliver the same benefits as the rest of the cloud computing infrastructure in terms of lowered cost, flexibility, scalability and high utilization. Choosing a software-based network OS allows cloud providers to standardize entire infrastructures on x86 server hardware, leverage investments in hypervisor platforms and utilize a single network OS from the network edge to the customer for everything from high-performance BGP routing to per customer firewalling and LAN bridging.

SECURE CONNECTIVITY

Cloud users access their applications and data over the Internet, requiring every user's connection to be encrypted for security. Software-based networking is an exceptionally clean solution for this requirement. Within the cloud a new Vyatta VPN virtual machine can be started in moments, using a small fraction of an existing server. The high cost associated with acquiring and installing a unique physical device is completely eliminated, as is the requirement for more space, power and cooling. The customer can deploy the same software or virtual machine at each access location rapidly and with minimal expense, as a "secure cloud connector."

CLOUD ON-BOARDING - SECURE LAYER 2 BRIDGING

An often overlooked requirement in cloud computing is the need to enable customers to securely migrate data to the cloud from the enterprise datacenter. The Vyatta Network OS combines Layer 2 bridging and IPSec/GRE Tunneling functionality to deliver a cloud bridging solution which allows physically separate networks to securely communicate with each other over the internet as if they were on a single Ethernet network. This capability simplifies the migration of applications and physical servers between data centers, ensures continuity during a phased migration, and enables the moving of virtual machines between physical servers on physically separate networks.

VIRTUAL FIREWALLING

For IT architectures within a customer's own datacenters, it's common for firewalls to be deployed at various places to ensure data security for sensitive databases and transaction systems. Issues related to both internal security (HR databases, financial systems) and external compliance (credit cards, health care, etc) must be clearly addressed. Deploying these IT systems in a cloud environment increases this firewall requirement. The customer not only must firewall its sensitive systems as it had before, but also to ensure security in a multi-tenant environment using a shared connection to the public Internet. Using traditional networking would require a lot of traditional hardware firewalls at a high cost, slow deployment, and with deep inflexibility. Software-based networking allows firewalls to be instantly deployed as virtual machines with no operating cost.

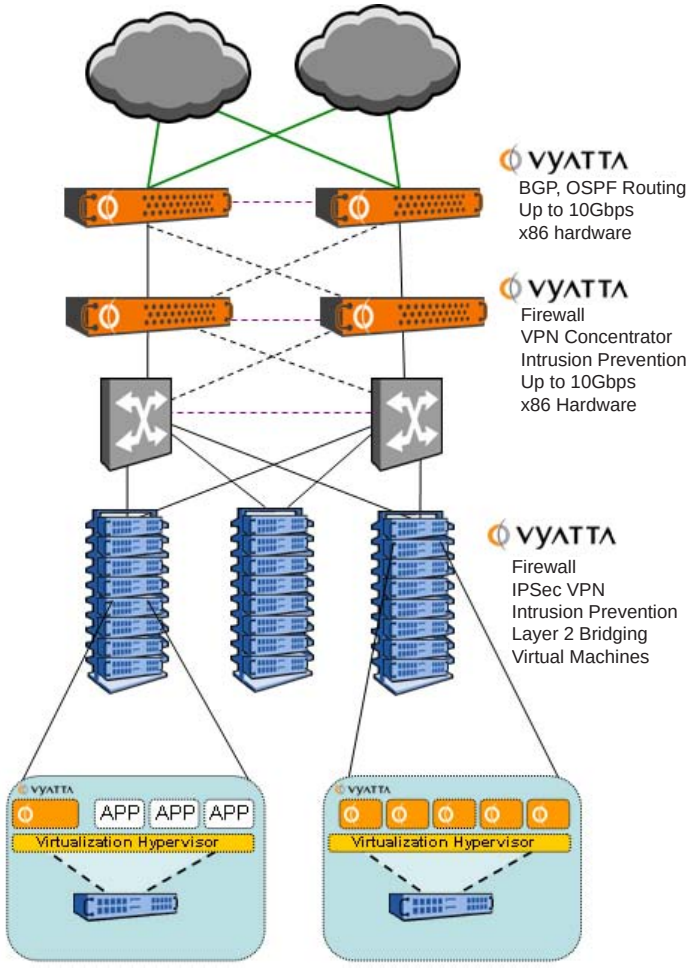

The Vyatta Network OS

The Vyatta network operating system is a scalable, integrated, enterprise-class networking solution that delivers advanced routing and network security functionality for physical, virtual and cloud networking environments. The Vyatta network OS includes dynamic routing, stateful firewall, VPN support, threat protection, traffic management and more in a package that is optimized to take advantage of multicore x86 processing power, common hypervisor platforms and emerging cloud architectures. All features are configured through Vyatta's familiar, networking-centric CLI, web-based GUI or third party management systems using the Vyatta Remote Access API.

Vyatta Software Highlights:

Network Connectivity

At the core of the Vyatta system is a complex routing engine with full support of IPv4 and IPv6 dynamic routing protocols (BGP, OSPF, RIP). Vyatta systems include support for 802.11 wireless, Serial WAN Interfaces and a wide variety of 10/100 thru 10Gb Ethernet NICs.

Firewall Protection

The Vyatta firewall features IPv4/IPv6 stateful packet inspection to intercept and inspect network activity and protect your critical data. Vyatta advanced firewall capabilities include stateful failover, zone and time-based firewalling, P2P filtering and more.

Content and Threat Protection

Vyatta systems offer an additional level of proactive threat protection with integrated secure web filtering and advanced intrusion prevention rules available as subscription-based Vyatta PLUS services.

Secure Connectivity

Establish secure site-to-site VPN tunnels with standards-based IPSec VPN between two or more Vyatta systems or any IPSec VPN device. Or provide secure network access to remote users via Vyatta's SSL-based OpenVPN functionality.

Traffic Management

The Vyatta system provides a wide variety of QoS queuing mechanisms that can be applied to inbound traffic and outbound traffic for identifying and prioritizing applications and traffic flows.

High Availability

Mission critical networks can deploy Vyatta with the confidence that high availability and system redundancy can be achieved through a number of industry standard failover and configuration synchronization mechanisms.

IPv6 Compatibility

Vyatta Subscription Edition software is the only software-based routing and security solution with proven IPv6 functionality and interoperability, ensuring a future-proof investment in a solution that offers a simplified migration path from IPv4 to IPv6.

Administration & Authentication

Vyatta systems can be managed through our familiar network-centric command line interface, web-based GUI or through external management systems using Vyatta's Remote Access API. All network management sessions can be securely managed using SSHv2, RADIUS or TACACS+.

Monitoring and Reporting

Vyatta systems present complete logging and diagnostics information that can be monitored using in industry standard toolsets such as SNMP, Netflow, Syslog, Wireshark and more.

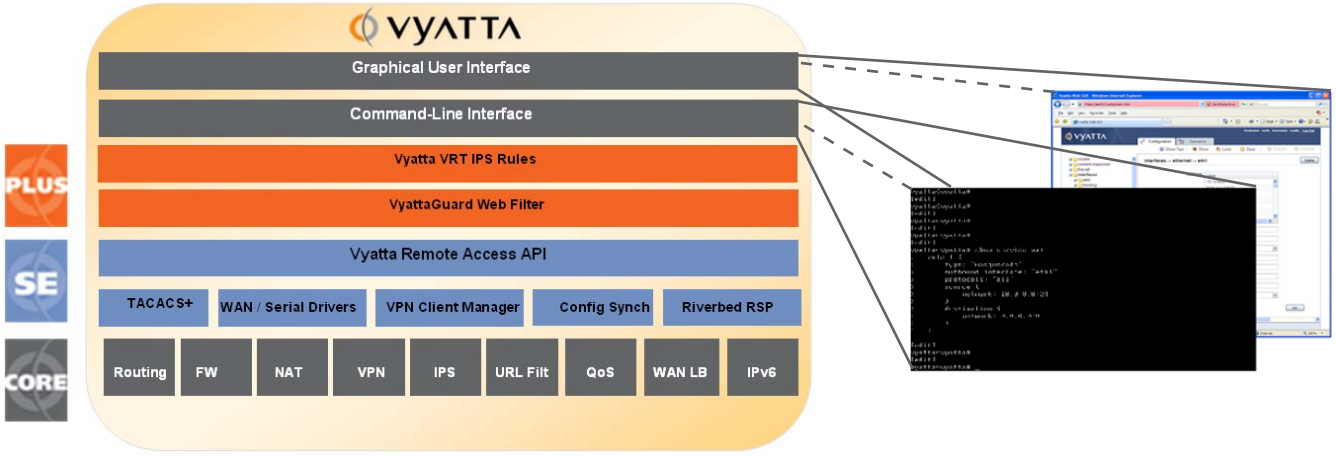

About Vyatta

Vyatta is disrupting the networking industry by delivering a software-based, open-source, network operating system that is portable to standard x86 hardware as well as common virtualization and cloud computing platforms. Vyatta software provides a complete enterprise-class routing and security feature set capable of scaling from DSL to 20Gbps performance at a fraction of the cost of proprietary solutions. Thousands of physical and virtual infrastructures around the world, from small enterprise to Fortune 500 customers, are connected and protected by Vyatta. For more information, please visit http://www.vyatta.com.

Managing End-user Experience, Application Performance Across The Cloud Infrastructure

Overview

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

While Cloud Computing is a significant technology turn that impact how future business services will be delivered to the enterprise, there are a number of challenges to overcome, before broad adoption in the enterprise. One of the characteristics of Cloud Computing is the ability to provision services on demand, and the flexibility to increase capacity on demand. This means a very dynamic environment with changes taking place. While change management is not a new concept within IT, dealing with the nature and volume of change are new.

As stated in the Cloud Networking Report, the key challenges created for the network to support Cloud Computing include:

- Manual network re-configuration to support Virtual Machine (VM) migration
- Maintaining the performance and controlling the cost of the wide area network (WAN)
- Services supported by a virtual and dynamic infrastructure

Each of these challenges creates potential application performance issues impacting the end users and the business. While IT organizations are faced with managing these issues, they still have to deliver consistent application and business services.

Best practices for managing change includes:

- Establishing what is normal so one can easily tell whether there is a change
- Measuring and identifying the impact of infrastructure changes on application performance and end user experience

If there are indeed performance issues, the support teams needs to:

- Identify performance degradation incidents
- Determine the root cause of degradation
- Resolve the problem

Actionable Intelligence is Key

As the report had identified, the task of moving a VM is simple in a virtual server management system. The challenge is in making sure the VM's network configuration state is also transferred. A best practice in managing change is to ensure end users are minimally impacted includes establishing a baseline of expected normal performance of the applications. When changes to the VM or network configurations are made in error, the support team can be alerted to the deviation in performance before the phone rings at the helpdesk.

When performance degradation occurs, the support team needs to be able to assess and identify the user, locations and maybe even impact to the business. Cloud Services, which could be a single application or group of applications, are supported by a virtual and dynamic infrastructure that can expand and move to accommodate changing business capacity. It is even more important to establish performance baseline so that the impact of the changes to the underlying infrastructure can be properly managed to maintain service levels.

Since one of the characteristics of Cloud Computing is the centralization of resources, it will drive more traffic over the WAN and have an impact on performance and cost. Having visibility of the WAN traffic and application performance, in relationship to the remotes sites supported provides the intelligence to make informed decisions about whether additional bandwidth is required to maintain the prior levels of end user experience.

A Unified Performance Management System

Application and network performance management systems are not new to IT support teams. The current breed of products is designed for silo used. This means the network team focus on network performance, and the application support team focuses on application performance. Managing application performance and end user experience in a Cloud infrastructure requires the information to be combined and presented in the context of the Enterprise and not in terms of the technology.

A unified performance management system is a solution that brings together data sources that are useful in tackling performance problems. The architecture of the unified solution includes a common data model and the ability to correlate data from multiple sources. These two aspects are crucial to providing a view of the various domains (application, database, server and network) in context with each other during the time when the performance problem occurred.

Increasing the business value of IT

For many CIOs, one key concern is not only increasing the business value of IT, but also quantifying the positive impact. While many organizations view IT as a cost center with a focus on reducing the expenses, many leading companies view IT as a strategic asset. These organizations focus on how IT can improve overall business value to the organization.

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

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

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

www.visualnetworksystems.com/ESI.

IT organizations have begun implementing business service dashboards and automating service desk workflow. A business service dashboard provides the line of business owners a clear view of the availability and performance of critical application and services that impact the bottom line. For IT, this dashboard increases the visibility of impact of infrastructure performance and changes on the business and users. This is especially important as the adoption of Cloud Computing expands.

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

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.

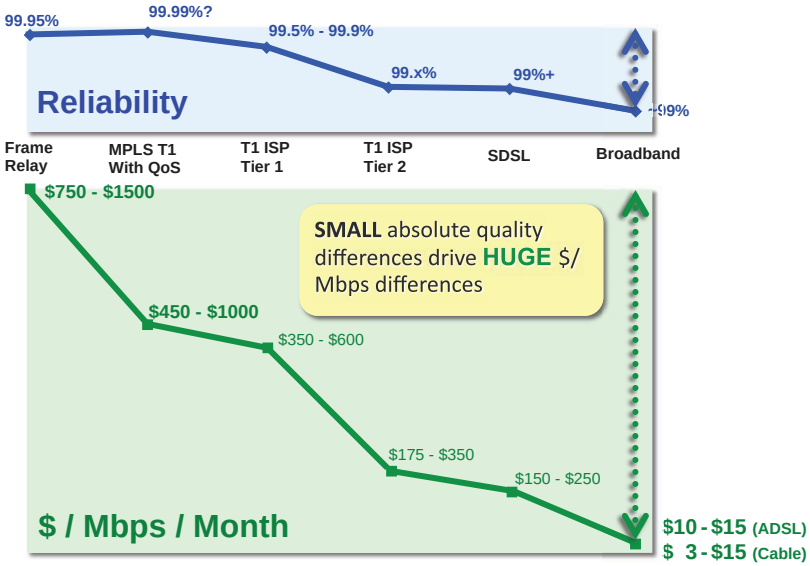

Figure 1: Private / Public WAN Pricing Disparity

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

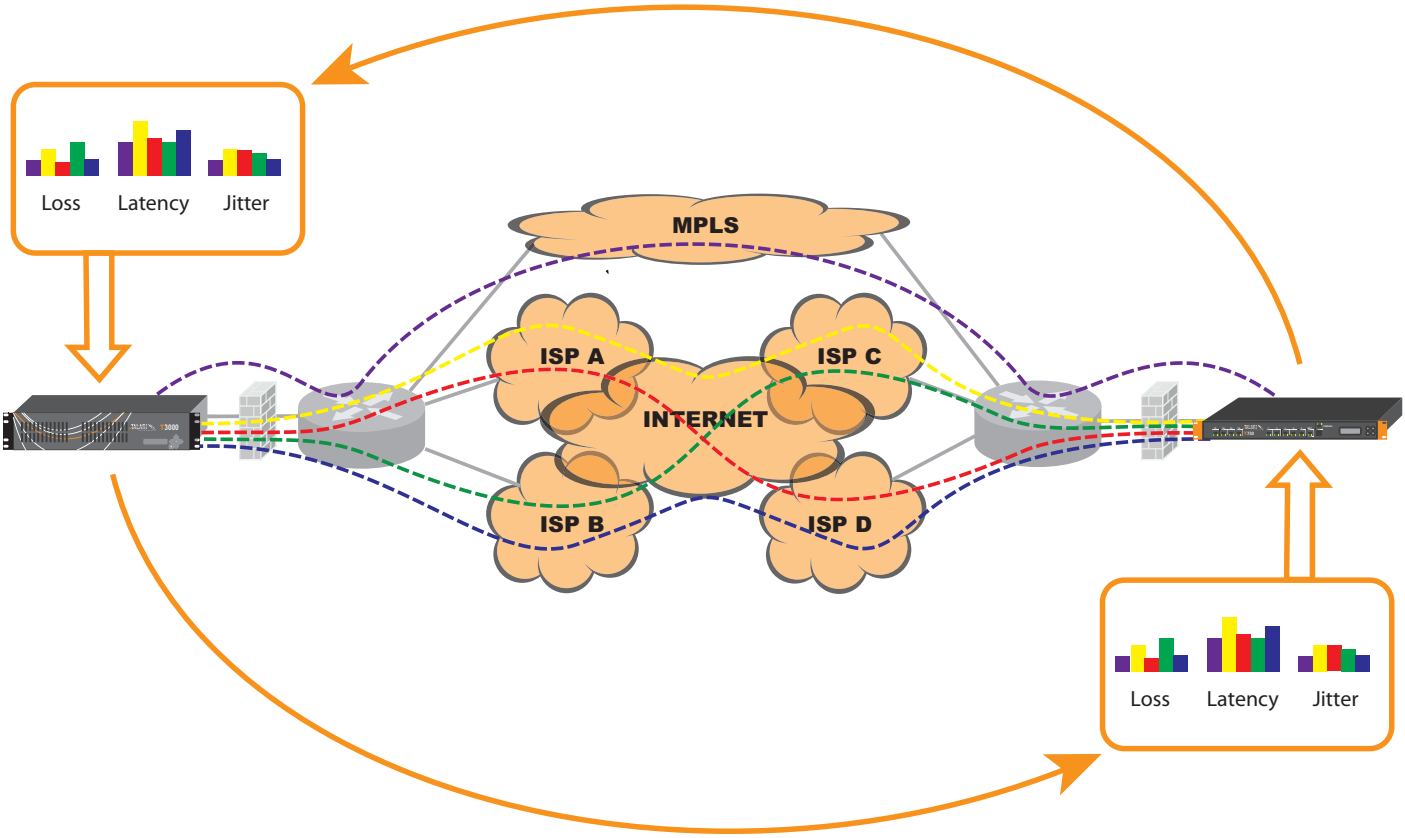

Figure 2: Continuous Measurment and Adaptation to Network Conditions

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**.

ENSURE THE BEST NETWORK PERFORMANCE FOR PUBLIC CLOUD COMPUTING

STREAMCORE
MAKE YOUR NETWORK CONSCIOUS

Increasingly enterprises are using cloud computing to improve agility, efficiency and cost-effectiveness of IT operations. However, some enterprises fear the risks of migrating critical, time sensitive business applications to the public cloud because guaranteeing network performance over the Internet is very difficult. By providing visibility and performance control over centralized corporate Internet access links or sites with direct-to-branch Internet connectivity, Streamcore solutions ensure that the network does not negatively impact the performance of public cloud services.

The use of enterprise software-as-a-service (SaaS) applications, such as Webex, GoToMeeting, Salesforce, Google Apps and Microsoft Online Services, is on the rise. Aside from security and regulatory compliance issues, maintaining acceptable service levels is the biggest concern that enterprises have when considering public cloud services. These interactive or real-time SaaS applications are accessed by employees through corporate Internet access links, whether centralized or not, and compete with bandwidth intensive traffic such as recreational Web surfing, emails and software updates. Consequentially, network congestion can severely degrade the performance of SaaS traffic, hindering all the benefits of public cloud services.

WHAT IS NEEDED:
DEEP PACKET INSPECTION + AUTOMATED QOS + ADVANCED VISIBILITY

The adoption of cloud computing services results in the need for both controlled network performance and better WAN traffic visibility, two of Streamcore's core competencies. In order to apply visibility and control for cloud computing traffic, a third key feature is required, the capability to identify these cloud computing services on the network.

DPI engine for cloud traffic

Public cloud computing traffic is always encrypted and exchanged over HTTPS for obvious security reasons, making useless traditional classification processes based on TCP/UDP ports or even on URL for HTTP traffic.
Streamcore has developed a powerful Deep Packet Inspection (DPI) engine focused on business traffic, such as VoIP, videoconferencing and Web business applications, whether encrypted or not. The Streamcore solutions allows automatic identification and classification of encrypted Webex, Salesforce and other public cloud computing traffic in specific classes for monitoring and prioritization.

Automated advanced QoS

Streamcore dynamically applies traffic shaping and prioritization based on the DPI classification process. It eliminates network congestion on corporate Internet access links by prioritizing cloud-based traffic. A single business criticality parameter is required, making provisioning extremely simple.
Another unique Streamcore feature is the ability to automatically manage competition between users of the same application, based on each session's behavior. For example, if different users access a SaaS application, Streamcore's patented QoS engine will analyze the behaviour of each encrypted HTTPS session, and perform appropriate automated prioritization for interactive flows.

Advanced visibility

Streamcore also provides visibility of traffic usage and performance, with application response measurements and quality indicators for voice and video communications. These measurements help IT staffs continually monitor traffic performance and ensure that all cloud-based applications and communications are performing at acceptable levels for end users.
Visibility is provided in true real-time (over the last 10 seconds) and over the long-term for up to two years. Different set of tools are available, either through a Web portal or PDF email report, to share information with stakeholders or within the IT team.

STREAMCORE SOLUTIONS: FOR ANY INTERNET ACCESS ARCHITECTURE

Streamcore provides monitoring and traffic shaping with asymmetrical deployment. Therefore, any type of network architecture for Internet and public cloud computing access is supported.

Centralized Internet Access

Today, most enterprises centralize their gateway and accompanying demilitarized zones (DMZ) toward the public Internet through major data center hubs. This type of architecture is often required by the IT security team, in order to minimize risk and costs, and to ease management of security products. In this case, enterprises can deploy StreamGroomers, Streamcore traffic management appliances, in front of the centralized Internet access link, in order to manage all public cloud computing traffic and guarantee its performance.

If SaaS and cloud computing traffic has to be delivered to remote branch offices from the data centers via the centralized Internet access, additional StreamGroomers can be deployed in front of the data center's private WAN access links. The StreamGroomers can manage cloud computing traffic delivery to remote branch offices over the WAN.

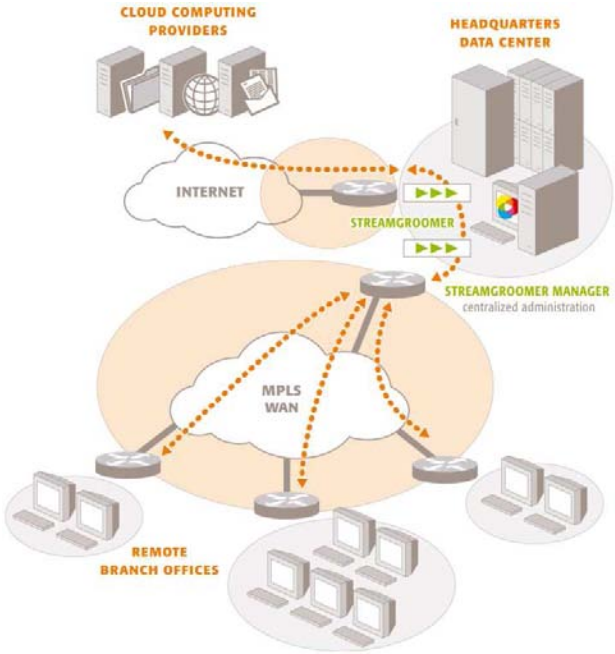

Fig. 1:
Centralized Internet access with branches over a private WAN

Branches with Direct Connections to the Internet

Backhauling Internet traffic to a centralized data center gateway adds latency and load on the private WAN.

Therefore, some enterprises prefer to provide Internet access directly into the VPN core, especially when they begin to rely heavily on public computing resources:

- If the private WAN uses IPSec technology over the Internet, this type of architecture is quite relevant. However, it can be challenging in terms of security because firewalls and security solutions (secure web gateways, antivirus...) must be fully distributed.
- Companies using a MPLS WAN sometimes have the option to migrate their Internet gateways and DMZs to the MPLS provider core. But, carriers may only offer a limited number of Internet gateways hubs around the world.

In such cases of branches with direct connections to the Internet, enterprises can deploy StreamGroomers in each branch office in order to manage and guarantee public cloud computing traffic performance over the branch WAN access link.

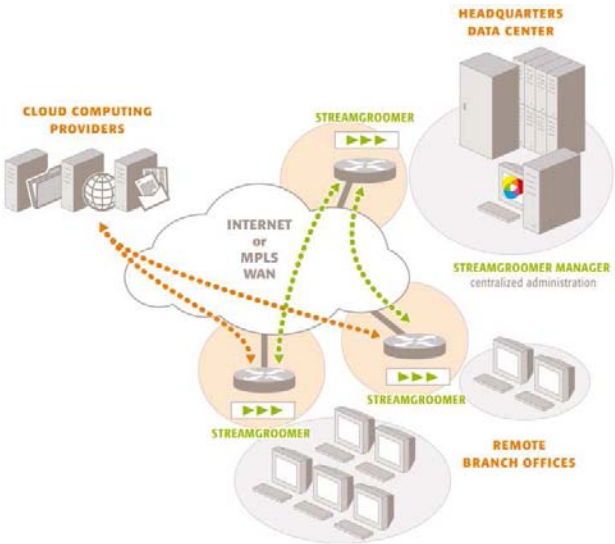

Fig. 2:
Branches with direct connections to the Internet

Hybrid Networks

On rare occasions, enterprises select an architecture in which there are two types of connectivity for each branch: an access link connected to a private MPLS network and another access link connected to a private IPSec network with direct-to-branch Internet access. This hybrid architecture combines the disadvantages of the two previous architectures: the high cost associated with MPLS, the complexity of securing distributed Internet gateways and the additional burden of managing traffic routed between the MPLS and the IPSec networks. However, this hybrid architecture can present advantages as well, such as extreme high availability, for enterprises with the budget and the right network/security team to manage it.

The full benefits of this architecture can be achieved by adding StreamGroomers at the branch: in addition to providing visibility and control for public cloud computing traffic, the Streamcore appliances can offer advanced load balancing per application. Bandwidth intensive applications can be automatically offloaded from the MPLS network to the IPsec network, and the MPLS access links can be dedicated to time-sensitive, real-time and business critical traffic.

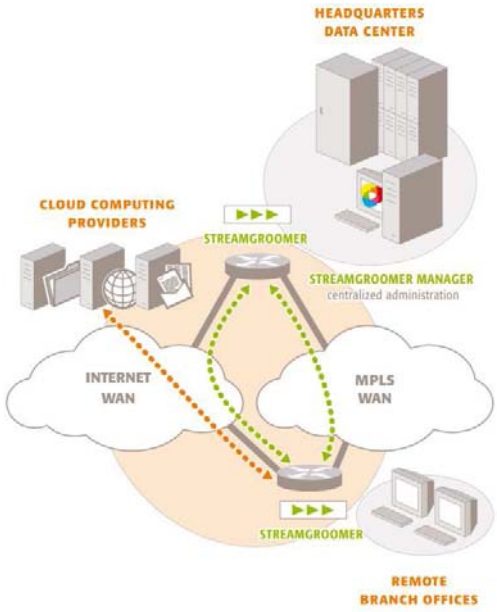

Fig. 3:
Hybrid network with load balancing per application performed by StreamGroomers

SUMMARY

By providing DPI, automated QoS, and advanced visibility for public cloud computing traffic, Streamcore provides the best solutions to monitor and ensure the best performance for SaaS applications. Streamcore products are suitable for all types of architectures that provide access to public cloud computing applications including centralized Internet access, direct-to-branch Internet access, and even hybrid networks that combine MPLS and IPSec technologies.

For more information, visit www.streamcore.com.

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

Axxia™ Communication Processor Accelerates Cloud Networking

APPLICATION INTELLIGENCE COMPONENTS

Application Visibility

- Who is accessing what?
- Top N applications
- Bandwidth consumed per application

Application Profiling and Control

- Network readiness for applications
- Troubleshoot application performance
- Application access control and QoS

Application Acceleration

- Application caching
- Application proxies
- WAN acceleration

Axxia Communication Processor

Cloud computing is all the rage. By 2014, Gartner expects worldwide spending on cloud computing to reach almost $150 billion. The goal of cloud computing is to enable IT organizations to achieve an order of magnitude improvement in the cost effective, elastic provisioning of IT services, including Software as a Service (SaaS), Infrastructure as a Service (IaaS), and Platform as a Service (PaaS).

Significant networking needs challenge this goal: dynamic scalability, lower latency, real-time resource management, self-healing reliability, lock-tight security, and guaranteed application performance. Existing networking infrastructure is already stressed to the breaking point. Faster ports, greater bandwidth, and flatter networking topologies can mitigate some of the challenges, but ultimately we need solutions that will scale with unforeseen demands. We need new cloud networks.

At LSI, we believe application intelligence is the essential ingredient, the key, to new cloud networks. Application intelligence allows every application across the network to obtain its fair share of resources, bandwidth, quality of service (QoS), and Service Level Agreement (SLA) in the presence of all other applications. Key ingredients for delivering application intelligence include:

- Application visibility
- Application profiling and control
- Application acceleration

The next-generation data center demands a new processor. A communication processor with a highly optimized architecture that enables each task to be allocated to the right resources for the job.

Cloud Networking: Lofty challenges, down to earth solution.

Application Visibility

Traditionally, applications were classified for QoS on network port plus IP address source+destination pair. That's no longer good enough. Cloud networks need to peer into data packets for fine-grained application visibility. In addition, applications such as unified communications, IP video, and telepresence require reliable real-time performance. The LSI™ Axxia Communication Processor features hardware-based deep packet inspection (DPI) for fine-grained application visibility with reduced packet latency and increased per-flow performance versus common approaches. DPI also allows the analysis of application signatures to eliminate common security threats like viruses, worms, and denial of service (DOS) attacks.

Application Profiling & Control

The ability to view and gain insight into how applications behave while flowing through network infrastructure can lead to improved design, better user experience and improved business innovation. The LSI Axxia Communication Processor incorporates a high-performance stateful flow processing architecture that targets the right on-chip resource for the job, from classification, to data and control plane processing, to traffic management, all necessary for profiling & control. True scalability, low latency, and deterministic performance result from this unique architecture.

Application Acceleration

Application visibility, profiling, and control enable real application acceleration and WAN optimization. Dramatic improvement in response times can be achieved with compression, application caching, content proxies, and virtualized application hosting.

Axxia has impressive CPU processing power and optimized application–specific resources to allow OEMs to deliver on the promise of cloud networking. In addition to application acceleration with Axxia, LSI offers media acceleration and storage acceleration solutions targeted at cloud networking.

Axxia Communication Processor

Powered by Virtual Pipeline™ Technology

The Axxia Communication Processor (ACP) is designed to meet the increased performance and lower power demands of next-generation communication networks. Using an innovative asymmetric multicore architecture, the ACP delivers fully deterministic performance with up to 20 Gb/s of data throughput, regardless of packet size, system loading, or protocol.

At the heart of each ACP is a high-performance multicore PowerPC® processor made by IBM® capable of reaching 2GHz operating frequency. Function-specific acceleration engines deliver fast path processing without unnecessarily taxing the multicore complex. These acceleration engines are derived from silicon-proven, cores used extensively on the broad product portfolio from LSI, including deep packet inspection, security, packet processing, and traffic management abilities. The ACP architecture uses Virtual Pipeline, a patented message-passing technique, for intra-processor communication between the acceleration engines, multicore complex and system on chip (SOC) subsystem components.

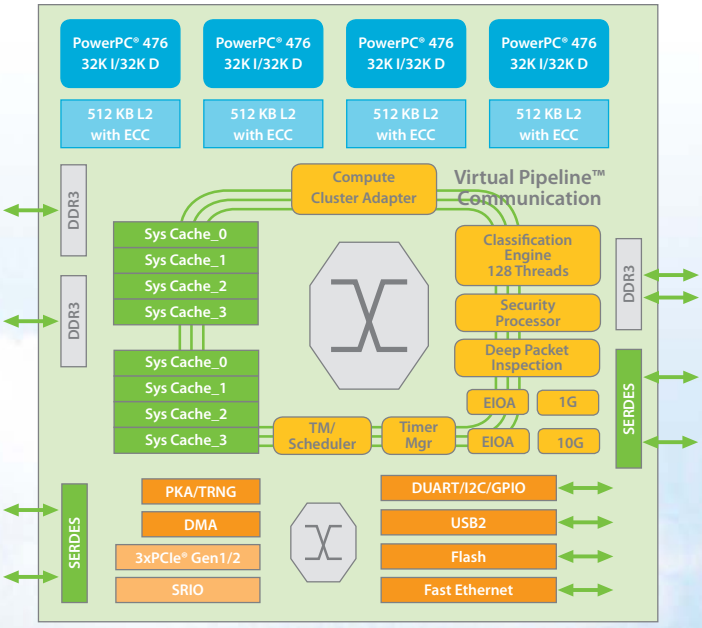

Figure 1 - ACP Block Diagram

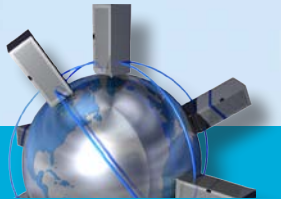

Axxia Intelligent Network Interface Card

This PCI Express® (PCIe®) NIC delivers an integrated cloud networking solution in a small footprint. Based on the Axxia Communication Processor, this turn-key solution provides application intelligence in cloud servers for security and monitoring applications, as well as server offload capabilities.

Axxia Intelligent NIC Hardware Features

- Based on Axxia Communication Processor
  - Built-in hardware accelerator engines including classification, deep packet inspection, packet integrity check, timer, packet assembly, programmable scheduler/ buffer manager, stream editor engines, and quad-core PowerPC® processor
- Dual 10GbE small form-factor pluggable (SFP) Network Interfaces
- PCIe Gen2 x4
- On board flash for board boot-up
- 10/100 Fast Ethernet port for initial development and debugging
- Optional serial port for management and debugging
- Supports fiber loopback
- Small foot print – PCIe half length card, full height
- On board system and configuration DDR3 SDRAM memory

Axxia Intelligent NIC Software Features

- Throughput up to 20 Gb/s (cut-through mode)
- Pattern recognition and replacement based on powerful classification and DPI engines
- Application recognition with ACP classification and DPI engines, and quad-core PowerPC® processor
- IEEE® 1588 support: flow classification and time-stamping, message type mapping, PTP egress processing
- IPsec: various encryption/ integrity/ authentication algorithms
- TCP proxy server offload: with ACP packet assembly and classification engine
- Packet delivery to host x86 CPU
  - Mechanism for transferring data block over PCIe
  - Large data block transfer; either per flow basis for each transfer or a mix of flows in same transfer

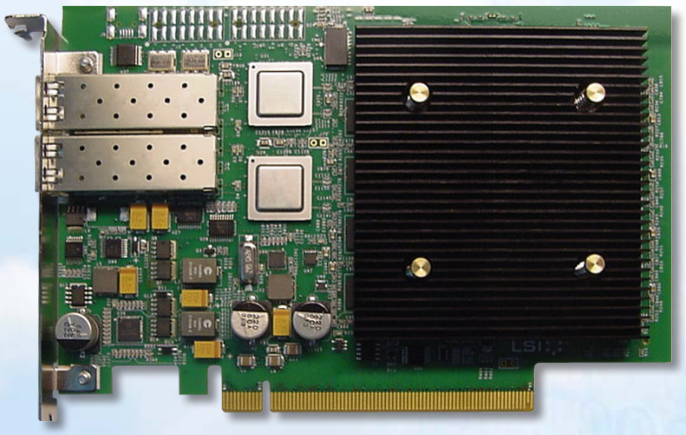

Axxia Intelligent NIC Software Architecture

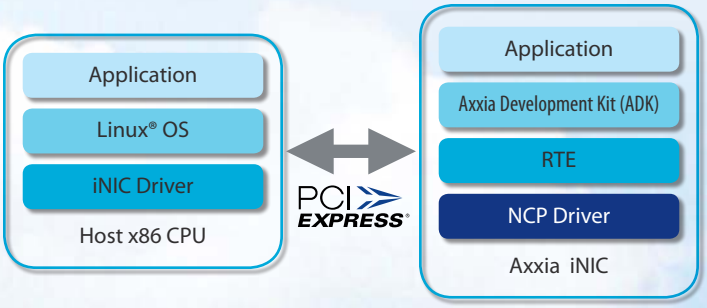

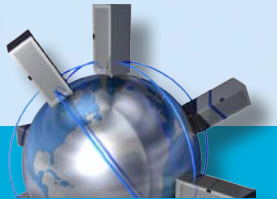

JUNIPER NETWORKS SOLUTION FOR CLOUD COMPUTING

Juniper Networks is dedicated to building simplified, scalable, agile, and secure networks that deliver the best performance and greatest efficiencies for cloud-ready data centers, while simultaneously controlling costs.

Getting Ready

Success in building a cloud-ready data center network requires three steps: (1) simplify, (2) share, (3) secure, as well as applying automation for smoother operations at each step. Whether you are running your internal IT infrastructure to be cloud-like or plan to connect with public cloud services, designing a cloud-ready data center network gives you significant advantages that can help you lower costs, increase efficiency, and keep your data center agile enough to accommodate any changes in your business or your technology infrastructure.

Key Components

Juniper recommends three steps to make your network infrastructure cloud-ready, reducing the cost and complexity of networking while improving application and business performance:

- **Simplify the architecture** — Consolidate siloed systems and collapse inefficient tiers using innovative fabric technology and a single network operating system. This results in fewer devices, a smaller operational footprint, reduced complexity, and easier management from a "single pane of glass."
- **Share the resources** — Segment the network into simple, logical, and scalable partitions for your various applications and services with privacy, flexibility, high performance, and quality of service (QoS) as primary goals. This sharing enables agility for multiple users, applications, and services.
- **Secure the data flows** — Integrated and dynamic security services resident in the network can provide benefits to users and applications sharing the infrastructure. Comprehensive protection secures data flows between external, internal, and inter-data center endpoints. Implement centralized orchestration and enforcement of dynamic, application- and identity-aware policies.

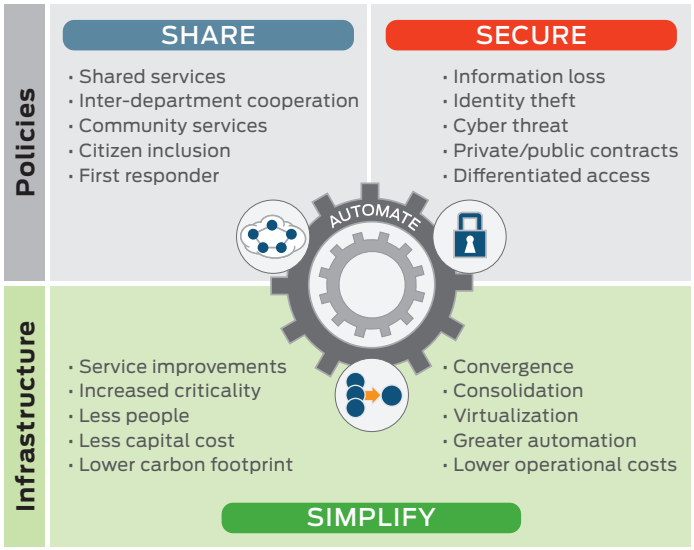

SIMPLIFY

The network design that used to work for the business might not be capable of supporting new demands on IT infrastructure and, most importantly, new business requirements. Networks built on fragmented and oversubscribed tree structures have problems with scaling and consistent performance. Design and management complexity and costs increase exponentially as more devices are added.

3-2-1 Data Center Network Architecture

Juniper simplifies the data center network and eliminates layers of cost and complexity with a "3-2-1 Data Center Network Architecture." Using fabric technologies such as Virtual Chassis technology, Juniper helps flatten data center networks, reducing them from three layers to two or even one layer. In the future, Juniper's Project Stratus will manage a 10GbE network at scale, as a single logical switch.

In addition, to help further simplify operations, Juniper consolidates multiple services into single high-performance platforms such as Juniper Networks® SRX Series Services Gateways, and utilizes the Juniper Networks Junos® operating system as the single OS across routing, switching, and security platforms.

SHARE

The cloud-ready data center requires network resources to be elastic, so that they can be allocated on-demand and at scale. Juniper's uniquely architected platforms deliver the agility and scaling required by virtualizing network configurations, segmenting services into logical domains, and using industry-leading hardware designs to scale without complexity. With a large pool of resources to draw from, customers can efficiently partition those resources to meet service requirements, remain flexible, and ensure operational performance, security, and control.

Edge Service Consolidation and Management

Juniper accomplishes this by building an intelligent network where these high-level policies can be enforced at the port level, and even at the date center's edge where connections to other data centers and networks occur over the WAN, the Internet, or a partner's network—effectively creating an even larger pool of resources to share across the organization. The Juniper Networks M Series Multiservice Edge Routers and MX Series 3D Universal Edge Routers are powerful, reliable, and the industry's most scalable solutions for the intelligent edge and inter-data center mobility.

SECURE

Security administrators must protect client-to-server traffic as well as traffic between physical and virtual servers, applications, and systems in other data centers. Security solutions need to be flexible to adapt to the changes in traffic volumes and data flows that occur because of virtualization, Web 2.0 applications, and cloud services. The increasing user access and the rising sophistication of security threats in a cloud-ready data center require expanded protection. Appropriate policies affect availability of business critical applications and operations.

To address these challenges, security services must be consolidated and

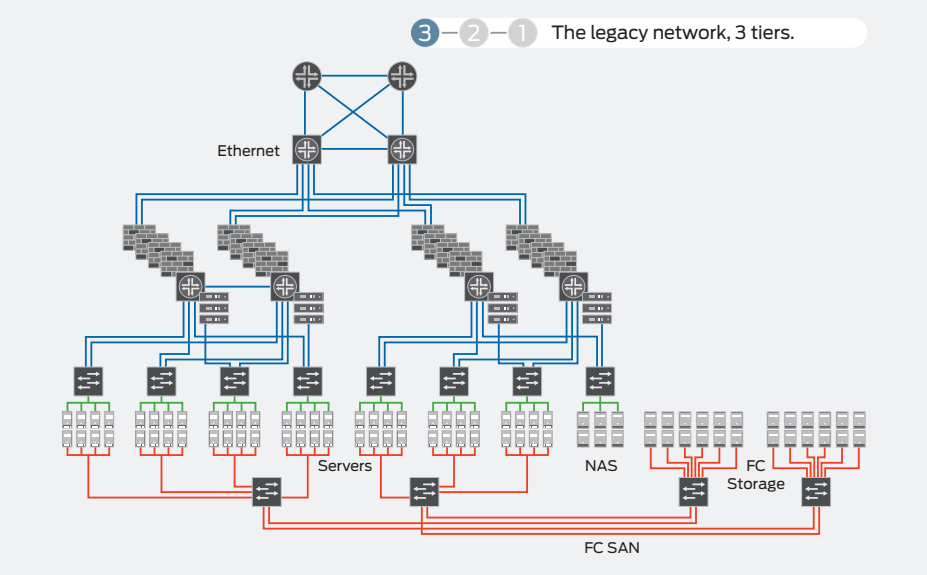

Figure 1: The legacy network

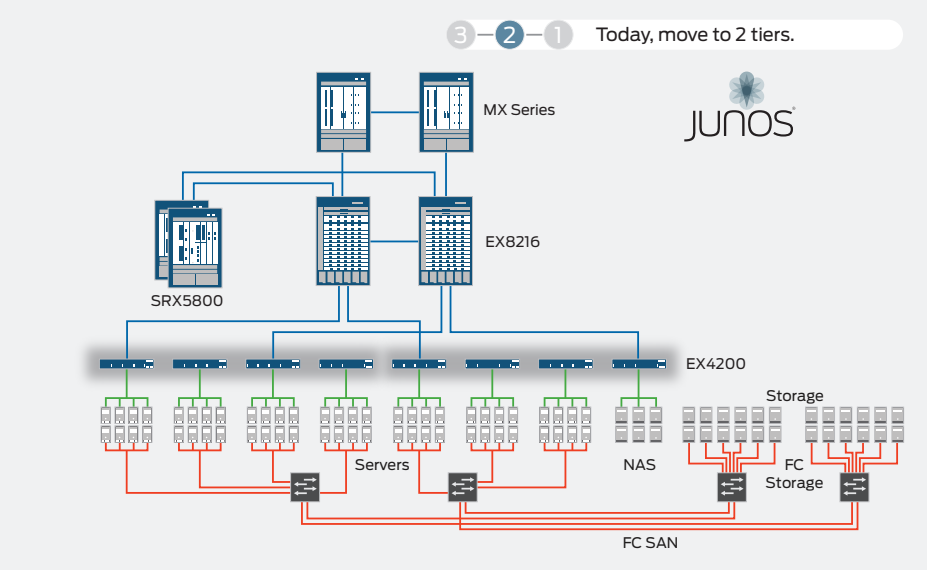

Figure 2: Juniper delivers a simplified two-tier network today with Virtual Chassis fabric technology.

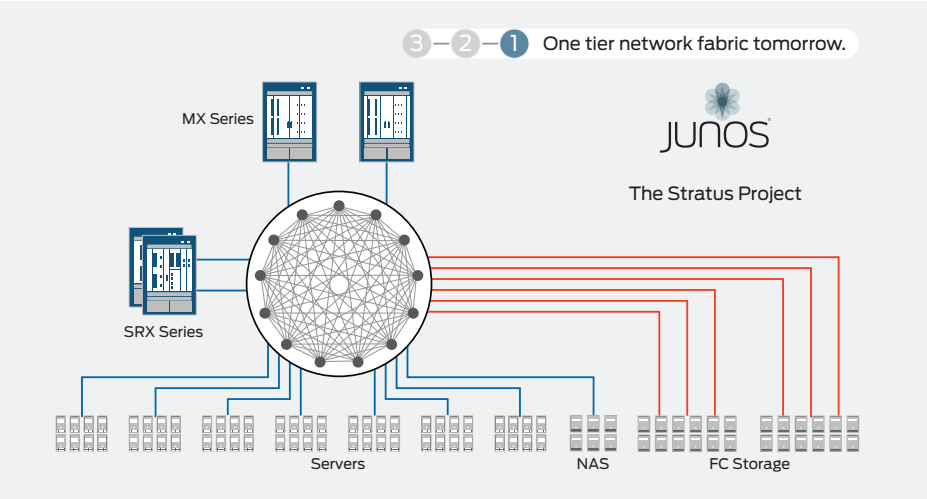

Figure 3: Juniper's vision for the ultimate simplification of the data center is Project Stratus, delivering a single fabric that unites Ethernet, Fibre Channel, and Infiniband networks.

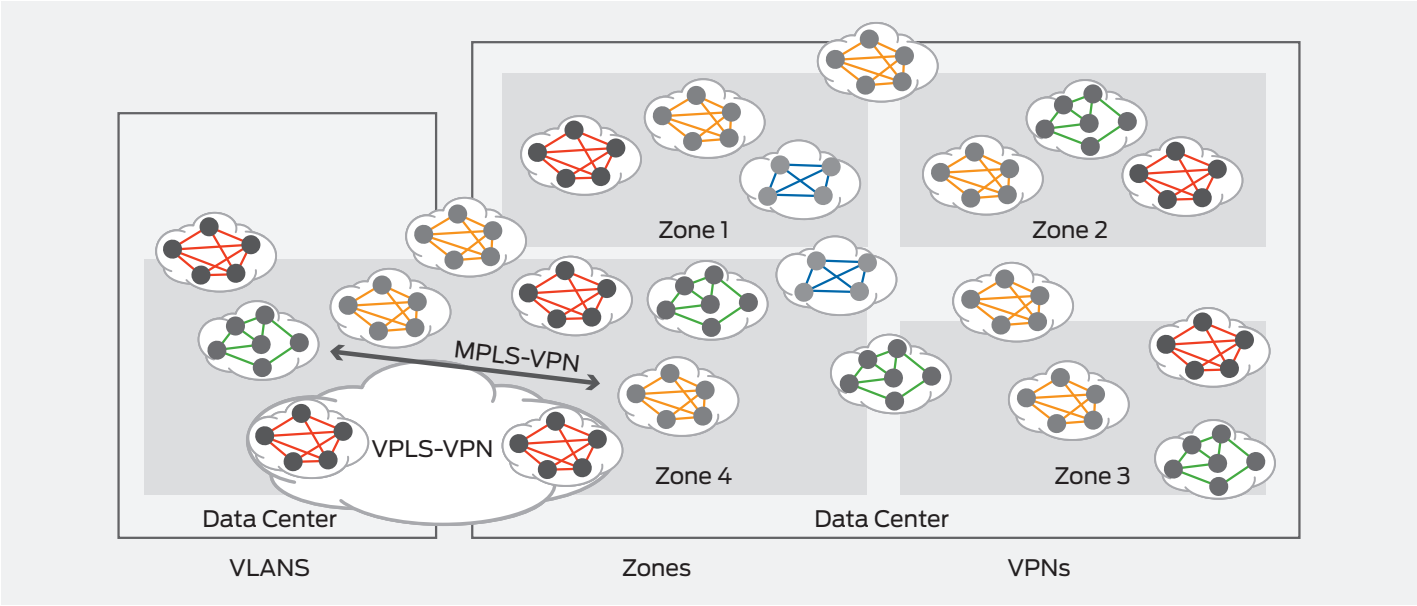

Figure 4: Scalable network virtualization technologies

pooled in a coordinated fashion to complement the simplification and sharing of the network. This approach enhances the flexibility and efficiency of the entire solution.

Juniper Networks has developed high-performance, cloud-enabled dynamic security services to meet today's security and performance requirements while accommodating future on-demand growth. Services such as application monitoring, stateful firewall, intrusion detection and prevention, and VPNs are consolidated on an expandable platform that flexibly and dynamically assigns resources as needed. Security services must be application- and identity-aware, while providing secure access for the mobile workforce to data center applications. Juniper provides best practices implementation guides to minimize risk and speed time to deployment when configuring security solutions for cloud-ready data centers.

AUTOMATE

Juniper's open, extensible network automation software makes it easier to manage and administer the data center by simplifying repetitive and complex tasks, defining and implementing policies within the network, and orchestrating implementation across multiple systems using network-based software. This greatly lowers operational expenses by reducing configuration errors, measurably improving reliability, and freeing up labor resources to innovate rather than administer.

The Juniper Networks Junos Space network application platform was designed to provide end-to-end visibility and control to enable network resources to be orchestrated in response to business needs. Operators can significantly simplify the network life cycle, including configuration, provisioning, and troubleshooting with an open automation platform.

Improve the Economics and Experience of Information Technology to Deliver Greater Business Value

Many organizations can benefit from cloud-ready data center networks, whether building a cloud-like infrastructure for internal purposes, connecting to public cloud services, or preparing to connect to public cloud services in the future. Juniper Networks, as a partner with wide-ranging experience, can help organizations reduce complexity and overall IT costs while accelerating delivery of IT services to users over a secure, simplified network.

For more information, please visit: **www.juniper.net/us/en/solutions/enterprise/data-center/**

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

Oct 2010

WAN Governance in a Cloud environment

Perform today, take control of tomorrow

Ipanema enables any large enterprise to have full control and optimization of their global networks; private Cloud, public Cloud or both. Moreover, Ipanema is the only system with a central management and reporting platform that scales to the levels required by service providers and large enterprises.

Leading the service providers market for application-centric network services, Ipanema has been proven in large enterprise global networks.

Enterprise infrastructure and WAN, are under constant transformation

Enterprises are on their way to the Cloud…

- They deploy private and hosted datacenters
- They use more and more SaaS applications (Salesforce, Googleapps…)
- Social media (LinkedIn, Twitter) and recreational applications (YouTube, Facebook) are popular
- Employees work not only from branch offices, also from home, hotels, airports…

… and yet to perform today they require:

- Guaranteed application performance
- Total business continuity
- Business process agility
- IT cost savings

WAN Governance aligns the network to IT priorities

WAN Governance is a unique Top-Down approach enabling enterprises to align their global network to IT and business priorities.

It fully controls and optimizes the global network, private Cloud, public Cloud or both. It guarantees that enterprises are always in control of critical applications. It unifies application performance across disparate networks. It dynamically adapts to whatever is happening in the network.

WAN Governance is the answer to all these challenges:

How to get full visibility of your global network:

- Discover which applications use your network resources
- Understand what is the root cause of slow applications
- Communicate clear data about application performance

How to deliver business applications:

- Guarantee voice, tele-presence and data applications over a converged network
- Ensure excellent application performance to your distributed workforce
- Manage social media and recreational applications

How to cost optimize your WAN:

- Reduce your WAN bandwidth requirements now and plan for tomorrow
- Use the Internet as a business network
- Get global control without deploying extra technology everywhere

ANS™, the Autonomic Networking System is the way to deliver WAN Governance

The Ipanema Autonomic Networking System is unique in many aspects:

- Its **central management** based on application performance objectives provides unmatched operation simplicity and automation
- It tightly couples key features in an **All-in-One** approach to ensure the best possible user experience
- Based on a **fully automated** "sense-and-respond" architecture, it adapts to any traffic situation and any network topology
- Its **collaborative agents** deliver full control with physical deployment in only 10-20% of locations
- It **scales** up to 10M users, 100K sites and 10K networks and can match any enterprise and large Service Provider deployment

Key features for an All-in-One system

Application Visibility provides full transparency for application traffic using a true L7 deep packet inspection, topology and performance. Its unique end-to-end metrics (like one-way-delay) easily differentiates network and IT problems. Embedded data consolidation and reporting provides all needed reports from C-level KPIs to technical information for the helpdesk team.

QoS and Control dynamically allocates network resources and combines all type of traffic (voice and tele-presence, Citrix, file transfer, CIFS…) fluidly – based on user behavior, application technical requirements and business criticality. It automatically takes into account complex situations like some-to-many and any-to-any traffic mesh and Cloud-based application delivery over private and hosted datacenters as well as SaaS.

WAN Optimization accelerates application response time and reduces bandwidth requirements by using all up-to-date techniques like byte caching, CIFS acceleration, TCP acceleration, etc.

Dynamic WAN Selection (DWS)automatically selects the best network for each new communication according to their availability, load and performance. Taking full advantage of Autonomic Networking System, DWS delivers many benefits to enterprises including:

- Unify application performance across hybrid networks
- Improve business communication continuity
- Seamlessly integrate Cloud based applications
- Exploit large network capacity at low cost
- Turn back-up lines into business lines

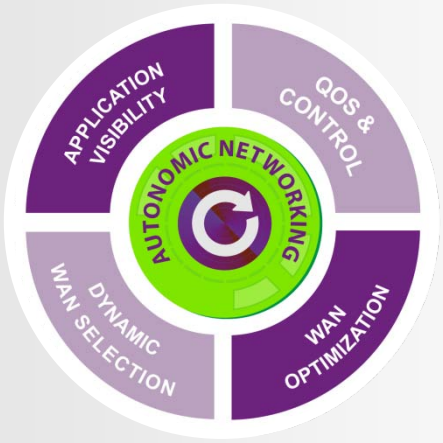

Powered by ANS™, WAN Governance brings tangible results

Get full visibility over your global network

- Eliminate 90% of network application performance issues
- Reduce problem identification and time-to-repair by 80%
- Ensure performance SLAs for all critical applications for 99,9% of the time

Deliver business applications

- Improve response time by 20x
- Reduce document download times from 5 minutes down to 15 seconds
- 0 business application brownouts during Olympic Games and Tour de France

Cost optimize your WAN

- Delay bandwidth upgrades by 24 months
- ÷3 the cost to transfer a Gbyte of data across the network
- Get full control with only 20% of technology expenses

Open Network Automation is Critical to the Virtual Data Center

Authored by Stephen Garrison,
Vice President Marketing,
Force10 Networks, Inc.

The Evolving Data Center

The data center has undergone several significant transformations since the birth of computing. The data center has evolved from mainframe computing to client server to Internet computing to SOA. Now we sit on the precipice of another major technology shift – the move to a fully virtualized data center (Figure 1). With each transition, the cost of computing was driven down by orders of magnitude and organizations were able to increase the efficiency of data center operations, software development, and most importantly, corporate workers.

Figure 1. Computing through the ages

The shift to a virtual data center will be the most significant IT transformation since the invention of the mainframe as it promises to bring together the network stack, storage and the computing layer to optimize application performance. In a fully virtualized data center, compute resources exist as VMs (virtual machines), storage becomes virtualized "pools" that can exist anywhere, and the network fabric connects these virtual elements to form a flexible, scalable computing environment (Figure 2).

The use of virtualization technology is widespread. A recent enterprise survey revealed that 82% of organizations today are using virtualization technology[1]. The primary driver for almost all companies using virtualization is to consolidate the number of servers. Obviously, this can have a huge impact on TCO since the number of servers can be dramatically reduced, sometimes by a factor of 10.

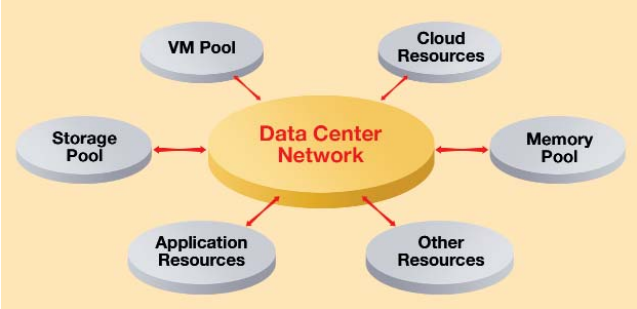

Figure 2. The virtualized data center is connected by the network

However, there are many other reasons for deploying virtualization technology, including:

- It allows software developers or other users to provision their own virtual machines. This will allow developers, engineers or others to have near instantaneous access to compute resources without having to involve several other departments.
- It ensures application performance is maintained when the workload is increased by provisioning additional computing resources.
- It increases the uptime of applications by mobilizing virtual workloads. In the event of an outage, the VM can move across a rack, across the data center or across the network whenever required.
- It acts as the bridge between physical and cloud-based data centers. Resources must be virtualized if they are to easily migrate between private data centers and cloud-based data centers.

The adoption of new technology always creates new challenges for data center managers, and virtualization is no exception. While server consolidation can dramatically reduce the number of physical servers, an unfortunate side-effect is that it results in an explosion in the number of virtual machines. Managing this so-called "sprawl" of virtual machines is much more difficult than managing physical resources. As organizations move from hundreds of VMs to thousands, questions such as "Where is that VM?", "Who created that VM?", "Who owns that VM",

[1] *Yankee Group Survey 2010*

"Why did it migrate?" and "Where is the data?" become more common. This new complexity results in additional work for server administrators as they shift their workload from managing tens or hundreds of physical servers to managing hundreds to thousands of virtual machines.

But the challenge does not stop there. With virtual machines, data center managers must also provision virtual storage pools and virtual network resources. In earlier times, managing the computing environment, which consisted of a static stack of compute, network and storage resources, was much simpler. But with virtual compute, storage and network resources, complexity has dramatically increased, resulting in more work for system, network and storage administrators.

The Role of Automation

The solution to the additional complexity caused by the extensive use of virtualization in the data center is automation. Automation will play an important role in helping data center engineers better manage virtual resources. Without automation, data center managers need to manually re-provision and optimize server, storage and network resources every time the smallest change in the environment is made. Keeping all of the virtual resources in sync is a near-impossible task for any data center of significant size. In fact, only 17% of respondents polled in Yankee's recent survey[2] feel that the tools to virtualize mission critical applications exist today. This leaves a big gap between the vision of the fully virtualized data center and the current market reality.

The challenge associated with managing a virtual environment is not limited to just deploying new technology, as data center operations and organizational structure are also impacted in a significant way. Today, most large data centers have administrative staff for supporting server, network and storage resources (Figure 3), and each of these groups have expertise in managing their respective technology. Prior to the adoption of virtualization technology, these groups could successfully operate in what were essentially independent groups. But the adoption of virtualization, combined with the need to quickly shift resources as demanded by the business, is now requiring these groups to work closely with each other.

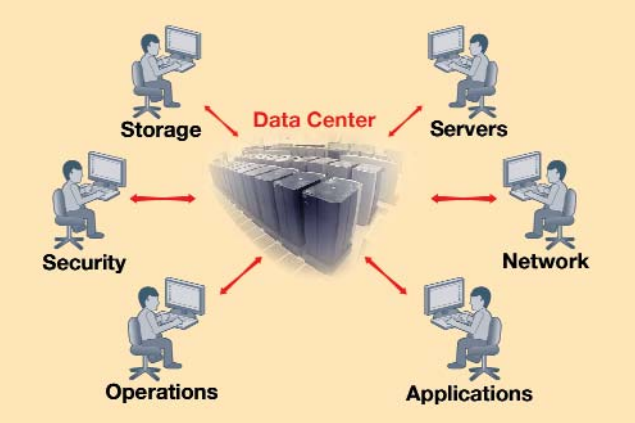

Figure 3. The virtual data center is everyone's responsibility

Automation

The additional complexity caused by the explosion of VMs, the need to tightly coordinate the provisioning of virtual resources, and the organizational challenges of managing this new virtual environment are best solved by automation. Automating the monitoring, management and provisioning of common tasks can greatly reduce the additional workload caused by virtual environments. Automation can also help standardize data center configurations, enforce best practices and increase availability.

For the network, automation can improve data center operations in the following ways:

- Instantly adjusts to changes in data flows, without manual reconfiguration, to optimize application performance. Virtualization, cloud computing, web 2.0 and other trends have given rise to bursty and unpredictable traffic flows. A congestion free network that provides non-blocking switching and routing performance can reduce the end to end latency of the transaction. This will also lead to the flat, layer 2 network that VMotion requires.
- Delivers an "always on" data center fabric. A high capacity, modular, fully redundant network can shift resources almost instantly to withstand any outage. Additionally, the network architecture can be simplified by increasing the density of the ports in the network devices. This means less hardware, a simpler architecture and increased uptime.
- Provides on demand resource allocation through automated network reconfiguration. The network can adhere to any business SLA (service level agreement) to automate tasks such as reallocating resources by moving VLANs, changing priorities through QoS policies, reallocation of bandwidth or reducing power consumption by shutting off underutilized resources.

[2] *Yankee Group Survey 2010*

Because the network is at the heart of the virtual data center, it is unique in its ability to enable organizations to maximize their investments in virtualization and cloud architectures.

The Different Approaches to Network Automation

The goal of network automation is to provide a self-optimized network that is capable of dynamically allocating virtual resources to where they are needed in a timely fashion. Several approaches to network automation have emerged, and data center architects, CIOs and others involved in designing virtual data centers need to be aware of the differences. The network vendors can be seen as falling into one of three high-level approaches:

Approach 1: Integrated Network Automation

This approach involves the vendor adopting a highly integrated, proprietary architecture that requires the customer to source all elements in the stack from a single vendor, or closed system of vendors. The upside of this "vertically integrated" approach is that it delivers a solution that works "out of the box", so there is some short-term benefit. Long-term however, this approach means vendor lock-in, which deprives customers of the power to choose the best technology for their specific environment. To date, Cisco has chosen to adopt this approach.

Approach 2: Network Controlled Automation

In this environment, the monitoring, management and provisioning of virtual environments is controlled from, or by, the network. When a new virtual environment is required, or if an existing virtual environment needs more resources, network management tools provision the network, compute and storage resources. This is a network-centric strategy that requires all of the data center functions to fall under the control of the network rather than working in a cooperative manner. This requires a huge cultural and operational shift by data center managers. This approach has been adopted by Brocade and Extreme Networks.

Approach 3: Open Network Automation

The third approach toward network automation is one that leverages open standards that allow the data center network fabric to be controlled by existing automation or middleware tools. Because this approach is server and application centric, it is consistent with current data center operations, allowing an organization to adopt network automation more seamlessly because current best practices can remain in place. With an open strategy, the network infrastructure aids the operations of the virtual data center but doesn't take on the role of managing the virtual environment. Managing the virtual environment is done by existing virtualization management or system management tools designed for this express purpose. Additionally, standards based protocols are used for exchanging information between the network fabric and hypervisors or virtual switches to manage network configurations. This allows companies to choose best-of-breed technologies and still have the assurance that the solution will work. The open, standards based approach to network automation provides the best long-term benefits for the customer, as it retains the current data center operational structure but still provides a path to the future. Force10 Networks is an example of a vendor that utilizes this approach

What to Look for in a Solutions Provider

As network automation continues to evolve, more and more vendors will claim to have solutions that can help an organization make the transition to a virtual data center. Considering the important role the network will play in the evolution of the data center, it is critical that the following be considered when making a purchase decision:

- An open, standards based approach. There are many solution providers that claim to be open and many that claim to be standards-based. However, it is crucial that the network truly be both. Some vendors that claim to be both will actually be including a number of proprietary features that are "based on standards".
- Hypervisor, virtual switch and server agnostic. If this isn't the case, the organization may lose its choice in compute platforms. Considering the rate of innovation and the reach of virtualization, it's important the network be able to support any of the hypervisor vendors.
- Non-blocking, congestion free architecture. This will minimize the end-to-end latency of traffic flowing across the network. Solutions that are "near non-blocking" or over-subscribed could lead to congestion problems that impair the performance of applications.

- Future proofed technology – high density, 40 GbE and 100 GbE ready. The network infrastructure being purchased today should be thought of as a five year investment. So, the hardware being procured needs to provide sufficient density to allow simplification of the network and upgradability to both 40 and 100 GbE. This will avoid a rip and replace event in the future.
- A vendor with a history of data center innovation. Networking in the data center has many demands that are unique. Choose a vendor that understands the demands placed on the data center network. Vendors who grew in the wiring closet may not have the right culture to meet the challenges of a data center.
- A broad ecosystem of partners. No single vendor can deliver on the vision of the virtual data center. The network solution provider used should have solutions that work with all of the major compute, virtualization, storage and management vendors.
- A solution provider that utilizes common scripting languages. Data center operations today are driven by scripts written in perl, python and UNIX. A network vendor that utilizes the de facto standard scripting tools can help bridge the gap between networking and computing more efficiently and more quickly.

Conclusions

The data center is on the verge of another major transition – the shift to a fully virtualized data center This will lower the cost of computing, improve uptime and application performance and raise corporate productivity to new heights. However, along the way, data center managers will encounter new challenges in managing a data center built on pools of virtual resources instead of physical ones.

Open network automation can help meet many of these challenges by delivering a network that works with the compute infrastructure to automate many of the mission critical, time sensitive tasks needed to run a virtual data center. Open network automation will:

- Enable a virtual infrastructure that can scale to handle unpredictable traffic demands.
- Create an elastic environment where virtual resources can be allocated where and when they are needed based on business policy.
- Improve application uptime by instantly adapting and applying network configuration changes that arise due to changes in the compute environment.
- Provide a bridge to cloud computing by allowing companies to coordinate the movement of resources to the cloud at their own pace.
- Help move customers towards the vision of a virtual data center much faster than solutions that use vertically integrated technology.

Force10 Networks, Inc.
350 Holger Way
San Jose, CA 95134 USA
www.force10networks.com

408-571-3500 PHONE
408-571-3550 FACSIMILE

WP39SG 1010 v1.2

The Network Platform for Cloud

Introduction

From Cisco's perspective, 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.

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

- "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 implementation.

Role of the Network

With users, devices and partners accessing virtualized resources and applications within the data center, the network is the essential platform for accessing and delivering cloud computing models. This includes the network in the cloud data center, the network between data centers, and the network connecting users from around the world.

This ubiquity creates a unique opportunity to build and take advantage of capabilities that can be delivered from the network to drive greater value out of cloud infrastructures.

Today's networks are already adopting key innovations for cloud computing: 10Gb Ethernet, WAN and application acceleration, Virtual Machine (VM) level traffic awareness, and enablement of VM mobility within and across data centers.

Cisco's networking capabilities align with three technology pillars: Unified Computing, Unified Fabric, and Unified Network Services. Together, these pillars are woven into Cisco's new Data Center Business Advantage architectural framework enabling enterprises to go from simple system consolidation and virtualization through to enabling infrastructure automation and secure private cloud deployment.

Figure 1. Network Capabilities for Cloud from Cisco's Data Center Business Advantage Architectural Framework

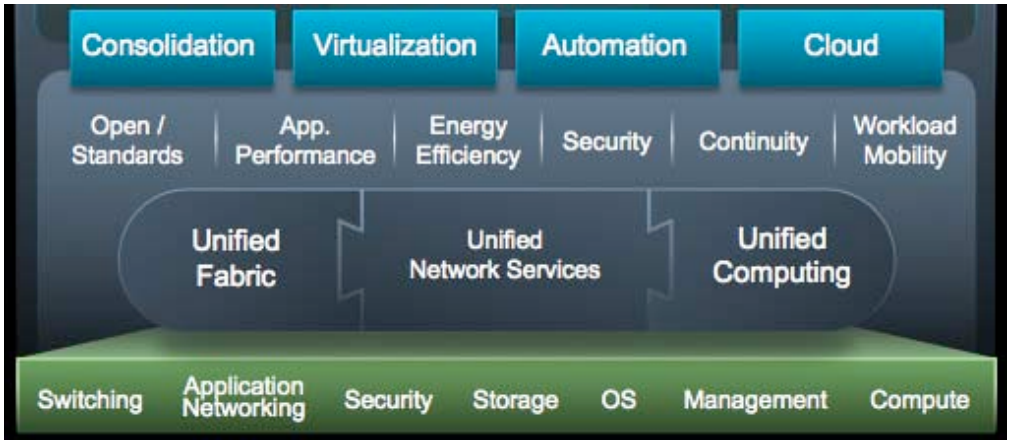

Unified Computing

Cisco's Unified Computing System (UCS) aims to provide scalable, dynamic compute resources for open, physical and virtualized environments. It does this by bringing together compute, network and storage access with virtualization to deliver better resource utilization, operational simplicity and workload mobility. It leverages the network intelligence and scale of Unified Fabric and the service readiness of the Unified Network Services.

UCS brings several innovative capabilities to data center servers, including:

- Extended memory technology allowing very dense VM hosting with up to 384GB of RAM per blade.
- Complete hardware abstraction through server profiles that allow mapping of configurations to the stateless compute blades in minutes.
- Native 10Gb Fiber Channel over Ethernet (FCoE) support.
- High Performance Virtual I/O (Ethernet NIC and FC HBA).
- Open, XML-Based API to provision, orchestrate and manage the UCS system.

Unified Fabric

Unified Fabric provides a simplified and integrated physical network for *all* I/O and communications in the cloud, including data, storage, voice and video. The fabric provides a converged network at scale with embedded intelligent capabilities that enable cloud.

With the widespread deployment of 10Gb Ethernet technology today, a roadmap to 40 and 100 Gb speeds, and the ratification of FCoE standards, Cisco views Ethernet as the fundamental layer for a unified fabric that can support multiple types of storage and data traffic simultaneously.

In addition to traffic within a data center, the unified fabric concept includes the extension of networks across facilities or geographic locations and the capabilities required to enable workload mobility.

Cisco delivers Unified Fabric across the breadth of its data center portfolio, including but not limited to the following:

- Unified Fabric in data center switching, from the hypervisor to the core with the Nexus 1000v, 2000, 5000 and 7000 series, interconnected with storage networks on MDS switches—all leveraging the consistent data center class operating system NX-OS.
- Cisco FabricPath Switching System (FSS) enabling broad Layer 2 data center networks, expanded VM mobility and efficient use of all available network bandwidth.
- Cisco Overlay Transport Virtualization (OTV) allowing Layer 2 continuity between geographically dispersed networks over any transport that supports IP, which in turn enables live migration of VMs between networks, data centers, and clouds.

Unified Network Services

Unified Network Services (UNS) is architected deliver services such as WAN optimization, firewalls, and load balancing in a concerted way across traditional appliances, inside other network devices such as switches and routers, and as virtualized appliances delivered on a hypervisor. This pillar executes a simple vision: to deliver any network service (security, WAN optimization, application delivery and load balancing, etc.), in any form factor (physical, virtual, appliance, integrated), in any environment (network, compute) and with any delivery mechanism (hardware-coupled or dynamic on-demand).

In addition to the industry leading physical appliances and network services that are embedded in different switches and routers either in software or via network modules, Cisco is now tapping into a new inflection point in the data center with the introduction of virtualized network services as part of Unified Network Services.

Cisco VSG works with the Cisco Nexus 1000v virtual switch's vPath capability and the Cisco Virtual Network Management Center (VNMC) to:

- Secure segmentation with zone-based firewall.
- Provide VM-level traffic visibility and granularity with context-aware rules.

Figure 2. Cisco's Unified Network Services Vision

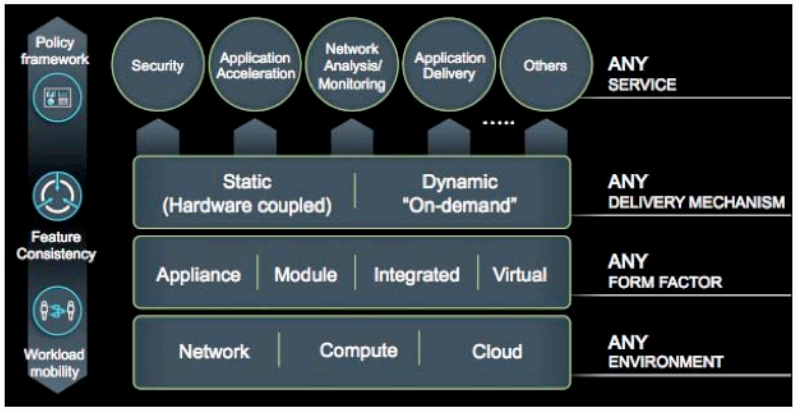

Policy-based centralized management. vWAAS is the industry's first cloud-ready WAN optimization solution. vWAAS works with the Cisco Nexus 1000v virtual switch's vPath capability to:

- Enable on-demand orchestration and policy-based application of rules down to the level of specific VMs.
- Provide separation of compute and storage with cache stored on SAN.
- Support multi-tenancy for cloud providers.
- Designed for optimizing traffic between and to clouds, both within the enterprise and from service providers.

Open Ecosystem and Market Success

Cisco's Data Center Business Advantage architecture is committed to delivering best-of-breed, open-standard networking solutions for cloud. Leveraging technology innovation and new delivery models, Cisco is giving customers greater choice than they've ever had within the Data Center.

- 11x World Record performance – Cisco Unified Computing System.
- 3x "Best of VMworld" winner (Cisco UCS, Nexus 1000v, Cisco OTV).
- Over 1.5 million 10Gb Ethernet ports shipped on Nexus switches.
- Over 40 ISV partners leveraging the UCS-API.
- VCE Coalition (VMware, Cisco, EMC) Vblock Infrastructure Packages.
- IVA Alliance (VMware, Cisco, NetApp) SMT Architecture.
- Cisco, Citrix and NetApp VDI Architecture.
- Application partnerships with Microsoft, Oracle, SAP and many others.
- Management partnerships with BMC, CA and many others.

For More Information

As you begin your own journey to the cloud, we invite you to discuss the right approach for your organization with Cisco. For additional information about:

Cloud: http://www.cisco.com/go/cloud

Data Center Business Advantage: http://www.cisco.com/go/dcba

Unified Computing: http://www.cisco.com/go/unifiedcomputing

Unified Fabric: http://www.cisco.com/go/unifiedfabric

Unified Network Services: http://www.cisco.com/go/unifiednetworkservices

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.**

Software WAN Optimization

Accelerate Your Business

certeon®

Transform computing, storage and networking resources *into* an integrated, agile and scalable cloud infrastructure *with* aCelera WAN Optimization software

certeon
THE Application Performance Company

"‥ application performance ‥ one of the top three inhibitors of cloud adoption"

Clouds and Beyond: Positioning for the Next 20 Years of Enterprise IT, Frank Gens, IDC

"Deploying virtual WAN optimization software has been as simple and inexpensive as remotely connecting to the server over the WAN"

"Virtual WAN Optimization software gives much more flexibility, which is imperative"

Ernest Ostro: Director of Information Services, Pathfinder International

Certeon Inc.
4 Van de Graaff Drive
Burlington, MA 01803
781 425 5200
http://www.certeon.com

Cloud Promise and Challenge

Cloud services look like a $100 billion-plus opportunity by mid decade, but is cloud computing worth this level of excitement? Think, Internet 1997. Companies were excited about the technology potential and worried about *security, privacy, bandwidth*, standards and more. In spite of these questions, what transformed communication and commerce? The ability to deliver business value!

In 2010 and beyond Cloud successes will be measured in business value. The units of measure will be the ability to increase business agility, decrease cost through on-demand provisioning and teardown of infrastructure and services, speed development, and improved reliability. It must be utility-based, self-service, secure and most importantly, have levels of application performance that improve productivity. User adoption is the linchpin of any business value equation.

Leveraging cloud computing and maximizing its value business value requires full featured, secure, scalable, high performance WAN Optimization software that allows applications to perform as expected, and can be part of any on demand architecture, rather than part of a farm of tactical hardware or limited virtual appliance solutions.

Cloud success requires integrating network services that are very far away and often owned by strangers

Business information and resources are increasingly being accessed at global scale distances, from enterprise and cloud sources using Internet, VPN or MPLS connections. At the same time, expectations for application performance are rising.

Enterprises embracing the cost and scalability benefits of cloud computing and service providers delivering consumption and utility-based models, balance the need for security and user expectations for access and application performance. Users don't care if the resource is in a cloud or on the moon, they expect their applications to work quickly and flawlessly.

Bottom line: the success of cloud computing is irreversibly linked to software based WAN Optimization and Application Acceleration technologies as the result of distance induced latency and the need to

provide ad-hoc secure and multi-tenant access. aCelera software WAN Optimization's ability to provide secure access, application performance and global scale make it the ideal cornerstone of cloud environments, from Private to Public to Hybrid.

Certeon

Certeon is the leading supplier of 21st century WAN optimization software for agile, elastic, and multi-tenant deployment. Certeon aCelera solves application performance challenges for cloud-based networks as effectively as it does for corporate networks. aCelera software and virtual appliances enable automated, secure and optimized performance for any application, on any device, across any network reducing response time by up to 95% while reducing the bandwidth used from 65 to 95 percent.

aCelera's creates global web of data that will enable businesses to leverage corporate and cloud provider networks to create new services or revenue streams. Certeon aCelera enables cloud service providers to offer on demand WAN Optimization to their catalogs as a one click value-added service.

Enterprise heterogeneous and decentralized needs

Enterprises today are a heterogeneous mix of hardware and virtualization platforms, custom and off the shelf applications, storage technologies, networking equipment and service providers all strung together in a web around the globe.

Decentralization of information sources, delivery workloads and productive users takes this heterogeneous infrastructure and explodes it's management and access problems across the globe. Clouds, company datacenters, branch offices, home offices, coffee shops are all part of the new enterprise.

The effort to make this mix of services and technologies useful, affordable and valuable has service providers of all types rolling out a range of cloud service models (IaaS, SaaS, PaaS, "X"aaS) and an array of deployment models (private, public, and hybrid clouds), that promise provide flexibility, scalability, cost savings that will create competitive advantage. But, even these environments are a heterogeneous mix of virtualization technologies from 3 or 4 vendors.

The combination of a heterogeneous infrastructure and decentralized enterprise with cloud services demands that WAN Optimization solutions be built to support this heterogeneous flexible infrastructure; they must use and be managed with the same building blocks as the environments they support. WAN optimization cannot just be a halo product targeting, or moving to a solution to cloud problems from outside the stack.

aCelera WAN optimization software: built for the cloud, not just moving to the cloud

Solutions "moving to" clouds do not support the dynamic, global and heterogeneous nature of enterprise or "X"aaService models. aCelera software and virtual appliances are "built" for the cloud and seamlessly integrate with all of these emerging technologies, delivering resources and services without compromising performance, scalability, or cost reduction.

The future of enterprise business success requires integrating network services that are very far away and often owned by strangers

"… 34% of survey respondents are using 2 virtualization solutions and 36% are using three or more."

"Users should plan for multiple virtualization platforms."

Fall 2010 ESG research study of 463 North American-based IT pros at organizations larger than 500 employees

"Productivity isn't everything, but in the long run it is almost everything." Paul Krugman

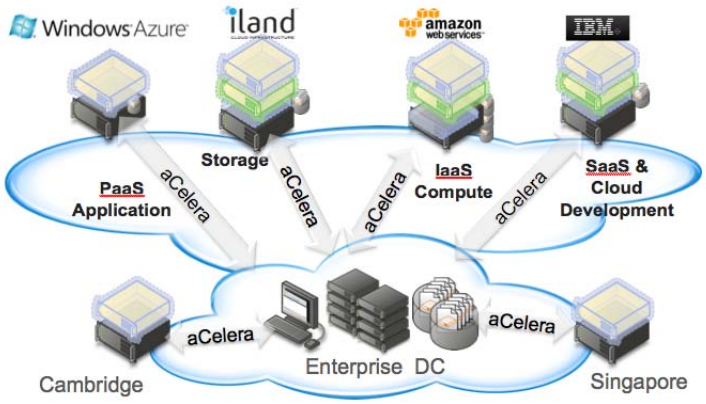

Enterprises and cloud service providers can deploy aCelera in any form factor, using any number of instances, delivering any throughput capacity, aligned with any application SLA requirement while meeting cost savings objectives and footprint limitations. This can be done in seconds on the enterprise premise, hosted, cloud sourced or in any mix of locations.

aCelera leverages enterprises' and service providers' growing heterogeneous virtualized infrastructures, in data centers, branches and on clouds. This allows organizations to turn clear TCO benefits into innovation. Saved acquisition, operations, real estate, power, cooling and maintenance/support costs create this opportunity where solutions not built for virtualized and cloud environments limit innovation.

aCelera™

Secure Automated Optimized

- Any form factor
- Any number of instances
- Any throughput capacity
- Any security requirement
- Any routing mode
- Any deployment model: enterprise, hosted, cloud sourced or combination

- Meet cost savings objectives
- Match footprint limitations

Microsoft GOLD CERTIFIED Partner

vmware READY

Virtualization was just the first step

Virtualization is a driving IT strategy and initiatives from SMBs through large enterprises, up to the very large hosting companies, carrier data centers and cloud providers. "Server virtualization provides a foundation for IT automation, dynamic workload mobility, and finally, a bridge to cloud computing."[1]

Virtualization cannot be a single vendor strategy. ISVs creating virtual appliances that support a single hypervisor platform are "moving to the cloud" with products that don't match the requirement to support heterogeneous environments. Single platform virtualization creates castaway technology - islands of virtualization capabilities that are an extension of hardware appliance platform.

aCelera: built for heterogeneous, decentralized work

Certeon's aCelera software is built to provide ALL the performance advantages of any HARDWARE WAN Optimization product along with the flexibility, scalability, manageability and cost-savings of software and virtualization. aCelera supports In-line & out-of-line deployment with software and hardware failover and any level of SSL security.

aCelera can be deployed in any virtualized private, public, and hybrid cloud computing environments and is poised to meet ANY future performance and agency demand imposed by any enterprise's heterogeneous, decentralized and cloud environments.

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

aCelera software exceeds the scalability and performance of purpose-built hardware appliances. aCelera software is built to support global enterprise scalability requirements and is ready for the Internet scale usage demands of managed services and cloud computing.

aCelera software WAN optimization - 60% better 3 year TCO and 50% better connection scalability

certeon
THE Application Performance Company

1, Enterprise Strategy Group

Virtualization: A Key Enabler for Effective Cloud Implementations

Why Use Virtualization?

As Cloud Computing adoption increases, virtualization is a key enabler, driving economies of scale and the ability to scale with hardware appliances or commodity hardware.

Virtualization solutions allow:

- Delivery of elastic, flexible and scalable solutions for changing-traffic volumes
- Enablement of a cost effective on-demand approach to reduce capital expenditure
- Efficiency for Public or Private Clouds

A10 offers a wide range of options, as one solution does not fit all requirements. Beyond the hype of Cloud generalizations is the reality of making the solution work for your unique needs. While Cloud providers take the burden off internal IT organizations, the risks of not considering the hardware used and potential issues of the wrong solution are apparent.

Organizations may no longer require owning the hardware in Cloud implementations, but they will still use similar devices to handle traffic. The advent of hypervisor solutions, or "virtual appliances" are serious options that offer an alternative to fixed hardware appliances, but each solution has its own pros and cons that must be considered, both from the feature and performance angles.

This is the reason A10 Networks' AX Series offers many solutions, from the flexible SoftAX to high performance hypervisor free AX Virtualization.

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 have the flexibility to choose the following scale-as-you-grow virtualization options.

SoftAX

- SoftADC: AX virtual machine (VM) atop a hypervisor on commodity hardware
- Rapidly scale with commodity hardware
- Reduce hardware costs and upload to compatible cloud providers
- Flexible solution leveraging an existing Cloud provider or internal virtualized infrastructure

AX Virtualization

- High performance multi-tenancy without hypervisor cost and hypervisor performance hit
- 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

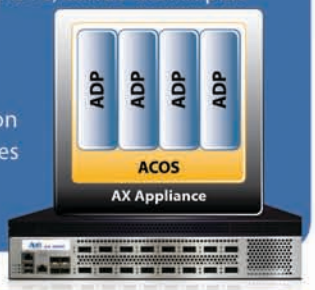

AX Virtual Chassis System (aVCS)

- Cost effective alternative to fixed ADC pairs and fixed chassis systems
- Massively increase performance to hundreds of Gbps and multiple millions of L4 connections per second
- Cluster multiple AX devices to operate as a unified single device
- Scale multiple AX devices with shared capacity, High Availability (HA) and single IP management
- Reduce cost and simplify management while adding devices as you grow

AX-V Appliance

- The first dedicated hardware platform designed specifically for hypervisor based ADCs
- Multiple SoftADCs: AX virtual machines (VMs) on dedicated AX Series hardware
- SoftAX flexibility with AX hardware performance and reliability
- Guaranteed performance, certifications, support and optimized hardware

The AX Series virtualization products and features are in addition to existing integration with leading third party virtualization vendors, such as VMware and associated solutions for vSphere acceleration, vCenter dynamic provisioning and VMotion with Global Server Load Balancing (GSLB).

vmware PARTNER TECHNOLOGY ALLIANCE

Virtualization at Work: Subaru Canada and A10 Case Study

Subaru Canada had been using the Foundry ServerIron 4G-SSL to provide server load balancing for its website (www.subaru.ca). However, when it came time to renew the support contract with Brocade Communications Systems, Inc., which had acquired Foundry Networks in 2008, Subaru Canada decided to evaluate some of the newer technologies available.

Subaru Canada's Director of eBusiness & Information Systems, George Hamin, became impressed with A10 Networks' AX Series New Generation Server Load Balancers while running a proof of concept using the AX 1000.

While Hamin and his team were impressed with the performance of the AX 1000, due to the rapid growth rate of sales at Subaru Canada, they decided they might later appreciate having the additional overhead provided by the AX 2500, with its 10 Gbps throughput capacity, as opposed to the 4 Gbps capacity of the AX 1000. With a list price of $2,500 per Gbps, the AX 2500 was a bargain, costing less than one-third of competing solutions (based on throughput-$-per-Gbps metric). Hamin said it was an easy choice, since the AX appliance cost "just a little more than the cost of renewing support on our 4G-SSL."

Hamin was originally interested in the AX's Application Acceleration features. The AX Series is optimized for SSL and L4-7 acceleration, and web caching further accelerates the user experience by reducing the time required to download each page. This, in turn, reduces the amount of bandwidth needed to serve pages and decreases the total number of requests placed to web servers. Furthermore, the AX Series offers several compression algorithms to reduce the size of each object on the page. Again, this helps reduce the amount of bandwidth being used. Hamin said he was able to leverage the compression and caching features in order to greatly accelerate the delivery of the enterprise's web content.

It was only after Subaru Canada had installed the 64-bit AX 2500 appliances that Hamin and his team learned of the additional AX virtualization feature. They were intrigued by the possibility that this feature might help them reduce the costs associated with supporting both mail and web applications. Virtualization allows customers to sub-divide an AX internally for multi-tenant purposes, whether for multiple organizations, departments, or simply, as in Subaru's case, multiple disparate applications. Each segmented area becomes an Application Delivery Partition (ADP). Within ADPs, various resources and elements are available. Layer 2/3 virtualization on a per-ADP basis was a particularly interesting enhancement to the ADP feature, as this guarantees true network segmentation between Subaru's applications.

SUBARU

Subaru Canada, Inc. markets and distributes Subaru vehicles, parts and accessories through a network of over 86 authorized dealers across Canada. This past March was their website's busiest ever, with 306,000 visitors viewing 1.97 million web pages.

"Once a potential buyer test drives one of our vehicles, the rest is easy. I feel the same way about A10's AX Series of appliances - once you try them you'll be sold… While we were originally drawn to the AX's application acceleration features, the recent enhancements to the AX Virtualization Multi-tenancy feature will allow us to consolidate our Microsoft Exchange 2010 environment and our web environment to a single pair of appliances, with high availability. This reduces the amount of Application Delivery Controllers in our network and saves us money in the process."

George Hamin
Director eBusiness & Information Systems for Subaru Canada, Inc

The AX Virtualization Multi-tenancy feature will allow Hamin to consolidate his distinct environments as if the ADCs were different platforms (i.e., a Microsoft Exchange Server 2010 environment and a web environment) onto a single pair of AX appliances. The pair of AX appliances will be set up in High Availability (HA) mode to mirror the content on the primary appliance and to act as a failover. This implementation will enable Subaru to reduce the total number of ADCs in the network, saving the company a large amount of money in the process.

"So rather than buying a pair of AX 2500s for HA web, another pair for HA Exchange, and another pair for HA SharePoint, you can virtualize a single pair and just keep throwing applications at it until you hit the limits imposed by your applications' collective peak load conditions, CPU, RAM, or ports," Hamin said.

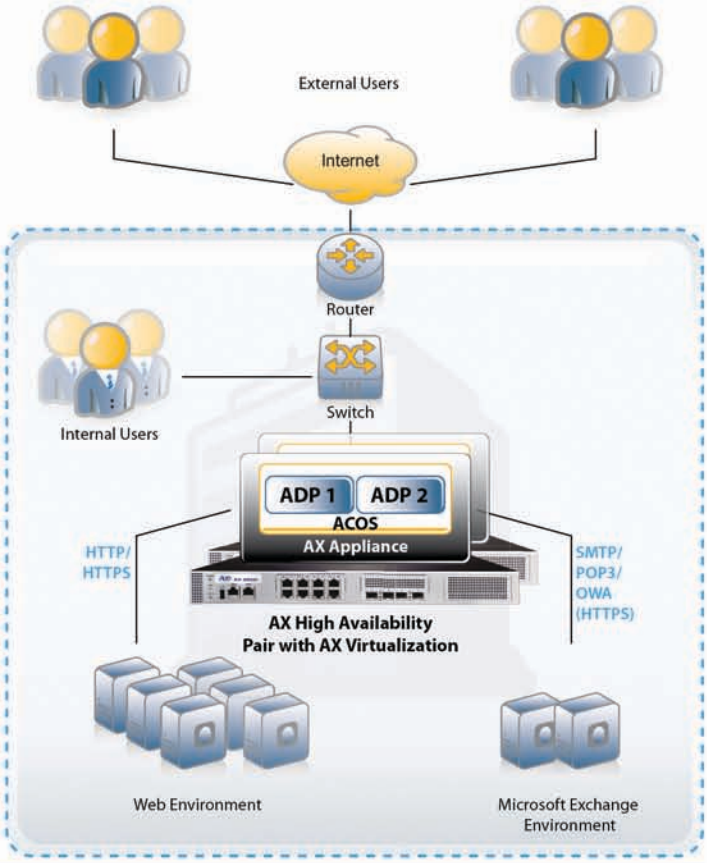

Summary

A10 Networks offers innovative virtualization solutions to enable any Public or Private Cloud deployment. With the widest range of solutions organizations can ensure they receive the right solution for their business and customers.

Please contact A10 for a free consultation of which solution would work best for your organization or to arrange a demonstration or trial at inquire@a10networks.com or www.a10networks.com

About A10 Networks

A10 Networks was founded in 2004 with a mission to provide innovative networking and security solutions. A10 Networks makes high-performance products that help organizations accelerate, optimize and secure their applications. A10 Networks is headquartered in Silicon Valley with offices in the United States, United Kingdom, France, The Netherlands, Germany, Brazil, Japan, China, Korea and Taiwan. For more information, visit: www.a10networks.com

Ask for A10

About AX Series

A10 Networks' AX Series is the industry's best price/performance advanced traffic manager – helping enterprises and ISPs maximize application availability through a high-performance and scalable web Application Delivery platform. The AX's Advanced Core Operating System (ACOS) architecture has garnered the company numerous awards and is revolutionary by market standards due to its scalable symmetrical multiprocessing (SSMP), shared memory architecture. AX includes an optimized multi-CPU architecture built from the ground up that leaps the competition in terms of performance, scalability and reliability. For more information, visit: www.a10networks.com/products/axseries