The 2010 Cloud Networking Report

By Dr. Jim Metzler
Ashton Metzler & Associates
Co-Founder, Webtorials Analyst Division

Table of Contents

1. Executive Summary

The majority of IT organizations have either already adopted, or are in the process of adopting cloud computing. The broad interest in cloud computing is understandable given that, as explained in this report, 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 primarily to the fact that on a going forward basis, IT organizations will continue to be required to provide the highest levels of availability and performance for a small number of key applications and services. However, an ever-increasing number of applications and services will be provided on a best effort basis. The phrase *good enough* refers secondarily to the fact that the SLAs from both traditional network service providers as well as from public cloud computing providers are often weak or non-existent. As such, these services are currently provided on a *good enough* basis, whether or not that is explicitly acknowledged.

The adoption of cloud computing creates some very significant networking challenges. In recognition of those challenges, the phrase ***cloud networking*** refers to the LAN, WAN and management functionality that IT organizations must put in place in order to enable cloud computing.

The Emerging Data Center LAN

The majority of IT organizations either recently has, or intends to redesign their data center LANs in the near term. The broad interest that IT organizations have in redesigning their data center LANs is driven primarily by the desire to reduce cost while simultaneously implementing the ability to support an increasingly virtualized and dynamic data center.

One of the most important characteristics of the contemporary data center is that an ever-increasing amount of the traffic is between servers. As such, a critical goal of the next generation data center LAN is to facilitate server-to-server communications. One approach for improving server-to-server communications is to flatten the data center LAN from the current norm that is either a three or four tier design, to a two tier LAN design consisting of access layer and aggregation/core layer switches.

One of the factors that has driven many IT organizations to implement a four-tier data center LAN is the fact that once an IT organization has implemented server virtualization there is a virtual switch (vSwitch) inside the server. The vSwitch presents a number of concerns related to management, security, functionality and organizational responsibilities. For example, an IT organization that has a thousand virtual servers in one of their data centers also has a thousand vSwitches that must be managed and configured. An emerging approach that potentially eliminates most of the issues caused by vSwitches is Edge Virtual Bridging (EVB). With EVB, all the

traffic from VMs is sent to the network access switch. If the traffic is destined for a VM on the same physical server, the access switch returns the packets to the server over the same port on which it was received; e.g., a “hair pin turn”.

One of the challenges associated with the redesign of data center LANs is that the combination of server consolidation and virtualization creates an “all in one basket” phenomenon that drives the need for highly available server configurations and highly available data center LANs. One approach to increasing the availability of a data center LAN is to use a combination of redundant subsystems within network devices such as LAN switches, in conjunction with redundant network designs that feature multiple links between devices. One of the new technologies that enables IT organizations to design data center LANs that are both faster and more highly available is switch virtualization. In this context, switch virtualization means that two or more physical switches are made to appear to other network elements as a single logical switch (e.g., virtual switch) with a single control plane.

If the data center LAN is designed with multiple links between devices, the connections between the end systems and the virtual access switches and between the virtual access switches and the virtual aggregation switches can be based on multi-chassis (MC) link aggregation group (LAG) technology. The combination of switch virtualization and multi-chassis LAG (MC LAG) can be used to create a logically loop-free topology without the need for the spanning tree protocol. This is important in part because the spanning tree protocol (STP) prevents all available forwarding resources in a redundant network design from being simultaneously utilized. As a result, the elimination of STP increases the link resource utilization and hence the scalability of the data center LAN. The elimination of STP also enhances the availability of the data center LAN because it eliminates the relatively long convergence times that are associated with STP. Unfortunately, the hashing algorithms that are associated with MC LAG are not standardized. As a result, each vendor’s implementation of MC LAG is proprietary.

A key characteristic of the emerging generation of data center LANs will be the convergence of block-level storage and data traffic over a common high-speed Ethernet data center switching fabric. This unified fabric offers significant cost savings in multiple areas including converged network adapters on servers and reductions in rack space, power and cooling capacity, cabling, and network management overhead. Traditional Ethernet, however, only provides a best effort service. In order to emulate the lossless behavior of a Fibre Channel (FC) SAN, Ethernet must be enhanced in such as way that it exhibits lossless behavior. Lossless Ethernet will be based on a number of emerging standards, which are commonly referred to as IEEE Data Center bridging (DCB). All data center LAN switching vendors are planning to support the DCB standards when they are available. In some cases the timing of the availability of that support may differ between the vendor’s access and core switches. In addition, some vendors are currently offering pre-standard support for DCB capabilities.

DCB will play a key role in supporting the Fibre Channel over Ethernet (FCoE) protocol specification that maps Fibre Channel’s upper layer protocols directly over a

bridged Ethernet network. FCoE provides an evolutionary approach to the migration of FC SANs to an Ethernet switching fabric while preserving Fibre Channel constructs and providing reliability, latency, security, and traffic management attributes similar to those of native FC. There is broad interest in FCoE on the part of the data center LAN switch vendors. However, since FCoE can be implemented in a variety of ways, there are several different levels of support that data center switch vendors can provide and still claim to support FCoE.

Wide Area Networking

The twenty-year period that began in 1985 saw the deployment of four distinct generations of enterprise WAN technologies[1]. For example, in the mid to late 1980s, it became common for enterprise IT organizations to deploy integrated TDM-based WANs to carry both voice and data traffic. In the early 1990s, IT organizations began to deploy Frame Relay-based WANs. In the mid to late 1990s, some IT organizations replaced their Frame Relay-based WANs with WANs based on ATM (Asynchronous Transfer Mode) technology. In the early 2000s, most IT organizations replaced their Frame Relay and ATM-based WANs with WANs based on MPLS. However, in contrast to the volatility of this twenty-year period, today there is not a fundamentally new generation of WAN technology in development. What often happens in this environment is that a new WAN service is created by modifying, and sometimes combining, existing WAN technologies and services.

The typical IT organization currently utilizes a wide range of WAN services with the primary WAN services used by IT organizations being MPLS and the Internet. It is common for the volume of WAN traffic to increase at an annual rate of thirty percent of more. One of the side effects of the movement to adopt cloud is that it will result in more WAN traffic. Unfortunately, the price/performance of MPLS tends to improve by only a couple of percentage points per year and few IT organizations are experiencing a significant increase in their WAN budget. Pulling these factors together yields the conclusion that IT organizations will not be able to support the added WAN traffic that results from the adoption of cloud computing unless they make changes that enable them to make more cost effective use of WAN services.

One relatively new WAN service that is generating a lot of interest on the part of IT organizations is Virtual Private LAN Service (VPLS). VPLS is an example of creating a new WAN service by combining existing WAN services and technologies. In particular, VPLS represents the combination of Ethernet and MPLS whereby an Ethernet frame is encapsulated inside of MPLS. As is typically the case with WAN services, the viability of using VPLS vs. alternative services will hinge largely on the relative cost of the services. This will vary by service provider and by geography.

Another WAN service that is created by combining existing WAN services and technologies is a hybrid WAN based on Policy Based Routing (PBR). When a router

[1] An enterprise WAN is designed to provide for connectivity primarily within the enterprise and between the enterprise and key contacts such as partners. This is in contrast to the Internet that is designed to provide universal connectivity.

receives a packet it normally decides where to forward it based on the destination address in the packet, which is then used to look up an entry in a routing table. Instead of routing by the destination address, PBR allows network administrators to create routing policies to allow or deny paths based on factors such as the identity of a particular end system, the protocol or the application. The advantage of PBR is that it enables IT organizations to leverage lower cost Internet services. The biggest limitation of this simple approach to hybrid networking is it that it creates a static allocation of traffic to multiple links and it doesn't have the ability to reallocate the traffic when the quality of one of the links degrades.

In order to be truly cost effective, a hybrid WAN has to be able to perform adaptive path selection across two or more WAN links in a dynamic, intelligent fashion. One of the principal advantages of a dynamic hybrid WAN (vs. a static PBR-based hybrid WAN) is that it allows IT organizations to add significant amounts of additional bandwidth to an existing MPLS-based WAN at a relatively low incremental cost. WAN Virtualization can be thought of as a variation of a dynamic hybrid WAN. In addition to enabling the augmentation of an MPLS WAN with inexpensive Internet connectivity, WAN Virtualization also gives IT organizations the option to reduce its monthly ongoing expense by either eliminating or reducing its MPLS connections while simultaneously providing more bandwidth than the original design did. This is accomplished in part by being able to leverage lower cost Internet access services including DSL, cable and on a going forward basis by leveraging 4G services.

A hybrid cloud relies on a WAN to provide the connection between the enterprise locations, including the enterprise data center(s) and remote sites, and the public cloud data center providing the IaaS or other cloud service. Ideally, the resulting hybrid cloud would appear to both users and administrators as a single cloud data center, with the physical location of application resources as transparent as possible. One of the challenges associated with hybrid cloud computing is that hybrid clouds depend heavily on virtual machine (VM) migration among geographically dispersed servers connected by a WAN. This is necessary in order to ensure high availability and dynamic response to changes in user demand for services. The desire to have transparency relative to the location of the applications has a number of networking implications including:

- **VLAN Extension**
  The VLANs within which VMs are migrated must be extended over the WAN between the private and public data centers.

- **Secure Tunnel**
  These tunnels must provide an adequate level of security for all the required data flows over the Internet.

- **Universal Access to Central Services**
  All application services, such as load balancing, DNS, and LDAP, should be available and function transparently throughout the hybrid cloud.

- **Application Performance Optimization**
  Application performance must meet user expectations regardless of user location within the enterprise network and the server location within the hybrid cloud.

Cloud bridging solutions that provide the functionality listed above are just now becoming commercially available.

The traditional approach to providing Internet access to branch office employees is to carry the Internet traffic on the organization's enterprise network (e.g., their MPLS network) to a central site where the traffic was handed off to the Internet. The advantage of this approach is that it enables IT organizations to exert more control over the Internet traffic and to simplify management in part because it centralizes the complexity of implementing and managing the organization's security policy. One disadvantage of this approach is that it results in extra traffic transiting the WAN, which adds to the cost of the WAN. Another disadvantage of this approach is that it adds additional delay to the Internet traffic. IT organizations are in the process of making increased use of local access to the Internet.

As previously noted, one of the side effects of adopting cloud computing is that it results in more traffic transiting the WAN. This has the potential to increase the cost of the WAN and to cause performance problems. One way to respond to those challenges is to implement network and application optimization techniques such as those provided by application delivery controllers (ADCs) and WAN optimization controllers (WOCs). ADCs evolved from server load balancers (SLBs) and offer a range of functionality including TCP offload, SLB and global SLB, SSL offload, XML offload, scripting and application firewalls. Until recently, ADCs were always hardware-based appliances. While that is still an option, today software-based ADCs are available from multiple vendors. One of the advantages of a software-based appliance is that it enables the type of agility that is associated with cloud computing and cloud networking. IT organizations can either provide ADCs themselves or acquire them from a third party as part of a managed service.

The goal of a WOC is to improve the performance of applications delivered across the WAN from the data center either to the branch office or directly to the end user, typically over a network such as MPLS. WOCs provide a wide range of functionality including compression, caching, de-duplication, protocol and application acceleration, spoofing and forward error correction. One of the primary reasons that have driven the existing deployment of WOCs is the consolidation of servers into centralized data centers. The movement to cloud will drive further server consolidation and hence further increase the need for WOCs. Other factors that are driving the increased need for WOCs is the need to support cloud bridging, VM migration, desktop virtualization and mobile workers. As was the case with ADCs, deployment options include both hardware and software-based WOCs, whether they are provided by the IT organization itself or by a third party as part of a managed service. Another

alternative is to acquire WOC functionality from a Software-as-a-Service (SaaS) provider.

Management

One of the primary characteristics of a cloud computing solution is virtualization, and the most commonly deployed form of virtualization is server virtualization. Unfortunately, server virtualization creates a number of management challenges including:

- **Breakdown of Network Design and Management Tools**
  The workload for the operational staff can spiral out of control due to the constant stream of configuration changes that must be made to the static date center network devices in order to support the dynamic provisioning and movement of VMs.

- **Limited VM-to-VM Traffic Visibility**
  The first generation of vSwitches doesn't have the same traffic monitoring features as does physical access switches. This limits the IT organization's ability to do security filtering, performance monitoring and troubleshooting within virtualized server domains.

- **Poor Management Scalability**
  Many IT organizations have experienced VM proliferation sometimes called VM sprawl. In addition, the normal best practices for virtual server configuration call for creating separate VLANs for the different types of traffic to and from the VMs. The combined proliferation of VMs, and VLANs places a significant strain on the manual processes traditionally used to manage servers and the supporting infrastructure.

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

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

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. A service can also be one of the key components of IT such as storage or computing. Historically IT organizations focused their management efforts on individual technology domains;

e.g., LAN, WAN, servers, firewalls. While that is still the most common approach to management, in the current environment a significant and growing percentage of IT organizations focus their management activities on the performance of applications and/or services.

As recently as two years ago, few IT organizations offered an SLA to the company's business and functional managers; a.k.a., an internal SLA. However, that situation has changed and now it is common for IT organizations to offer internal SLAs. In particular, over two thirds of IT organizations provide an internal SLA for at least some of their applications. However, the growing interest in offering internal SLAs for key applications is an impediment to the use of SaaS. In particular, few if any SaaS providers provide a meaningful end-to-end SLA for the performance of the applications that they provide. This lack of meaningful SLAs from SaaS providers is a deterrent to the Global 2000 adopting these solutions for delay-sensitive, business-critical applications.

The task of dynamically creating or 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 VLAN memberships, QoS settings, and ACLs) is established or transferred in a timely fashion. In many instances today, these network configuration or reconfigurations requires a time-consuming manual process that involves multiple devices.

In the current environment, the most common approach to automating the manual processes involved in the dynamic provisioning and migration of VMs is based on communication between the Hypervisor Management system and the switch element management system (EMS) via APIs supported by both vendors. A somewhat different approach to automating data center configuration, including the provisioning and migration of VMs is based on Orchestration engines. Service orchestration is a centralized server function that can automate many of the manual tasks involved in provisioning and controlling the capacity of dynamic virtualized services. In the case of VM provisioning and migration, the Orchestration engine would function as the point of integration between the network device EMS and hypervisor management system. Orchestration solutions are available from a number of network management vendors and hypervisor vendors. In addition, a dynamic virtualized environment can also benefit greatly from a highly scalable and integrated DNS/DHCP/IPAM solution, which is also well integrated with the virtual server management system.

In addition to the challenges listed above, the adoption of cloud computing makes troubleshooting application performance an order of magnitude more difficult. One of the reasons for this is that particularly in the case of either a public or hybrid cloud computing solution, the network topology becomes even more complex and hence understanding the end-to-end path becomes notably more difficult. For example, consider a branch office that is now using a WAN to access multiple internal data centers as well as multiple cloud computing service providers. There are typically multiple paths that the traffic can take from the branch office to each destination. The complexity of managing this is greatly complicated by the fact that applications and

services can dynamically move between servers, both within a given data center as well as between disparate data centers. Route analytics enables IT organizations and service providers to rapidly troubleshoot the complex logical problems that can occur in any large meshed network, and which are more likely to occur in public and hybrid cloud solutions. The value that route analytics offers is that it provides visibility, analysis, and diagnosis of the issues that occur at the routing layer in complex, meshed networks.

Another fundamental challenge relative to managing either a public or hybrid cloud computing solution is that the service has at least three separate management domains: the enterprise, the WAN service provider(s) and the various cloud computing service providers. In order to effectively manage, monitor and troubleshoot a public or hybrid cloud computing solution, detailed management data has to be gathered from all three domains. While some providers provide an API to enable that to happen, for the most part, effectively managing a public or hybrid cloud computing solution is still largely a work in progress.

2. 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.

3. 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[2] 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.

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

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 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.

4. The Emerging Data Center LAN

First and Second Generation Data Center LANs

As recently as the mid 1990s Local Area Networks (LANs) were based on shared media. Throughout this report these shared media LANs will be referred to as First Generation LANs. In the mid 1990s, companies such as Grand Junction introduced Ethernet LAN switches to the marketplace. The two primary factors that drove the deployment of Second Generation LANs based on switched Ethernet were performance and cost. For example, performance drove the deployment of switched Ethernet LANs in data centers because FDDI was limited to 100 Mbps whereas there was a clear path for Ethernet to evolve to continually higher speeds. Cost was also a factor that drove the deployment of Ethernet LANs in data centers because FDDI, in addition to being the only viable, high-speed First Generation LAN technology for interconnecting servers, was fundamentally a very expensive technology.

A key characteristic of Second Generation data center LANs is that they are usually designed around a three-tier switched architecture comprised of access, distribution and core switches. The deployment of Second Generation LANs is also characterized by:

- The use of the spanning tree protocol at the link layer to ensure a loop-free topology.
- Relatively unintelligent access switches that did not support tight centralized control.
- The use of Ethernet on a best-effort basis by which packets may be dropped when the network is busy.
- The need to support applications that are neither bandwidth intensive nor sensitive to latency.
- Switches with relatively low port densities.
- The over-subscription of uplinks.
- The separation of the data network from the storage network.
- VLANs to control broadcast domains and to implement policy.
- Redundant links and fast failover protocols to increase availability.
- Access Control Lists (ACLs) for rudimentary security.
- The application of policy (QoS settings, ACLs) based on physical ports.

Drivers of Change

The Webtorials Respondents were asked “Has your IT organization already redesigned, or within the next year will it redesign, its data center LAN in order to support cloud computing in general, and virtualized servers in particular?” Their responses are shown in Table 4.1.

	Already Have	Will Within the Next Year	No Plans
Cloud Computing in General	28.6%	42.9%	28.6%
Virtualized Servers in Particular	50.5%	30.7%	18.8%

Table 4.1: Redesign of the Data Center LAN

The data in Table 4.1 indicates that one of the key factors driving IT organizations to redesign their data center LANs is the deployment of virtual servers. The Webtorials Respondents were asked to indicate the percentage of their company's data center servers that have either already been virtualized or that they expected would be virtualized within the next year. The responses of The Webtorials Respondents are shown in Table 4.2.

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

Table 4.2: Deployment of Virtualized Servers

As pointed out in ***Virtualization: Benefits, Challenges and Solutions***, server virtualization creates a number of challenges for the data center LAN. As previously discussed, one of these challenges is the requirement to manually configure parameters such as QoS settings and ACLs in order to support the dynamic movement of VMs. Other challenges include:

- **Contentious Management of the vSwitch**
  Each virtualized server includes at least one software-based virtual switch (vSwitch). This adds yet another layer to the existing data center LAN architecture. It also creates organizational stress and leads to inconsistent policy implementation.

- **Inconsistent Network Policy Enforcement**
  Traditional vSwitches can lack some of the advanced features that are required to provide the degree of traffic control and isolation required in the data center. This includes features such as private VLANs, quality of service (QoS), and sophisticated ACLs.

- **Layer 2 Network Support for VM Migration**
  When VMs are migrated, the network has to accommodate the constraints imposed by the VM migration utility; e.g., VMotion. Typically the source and

destination servers have to be on the same VM migration VLAN, the same VM management VLAN, and the same data VLAN.

Server virtualization, however, is not the only factor causing IT organizations to redesign their data center LANs. Table 4.3 contains a list of the factors and the percentage of The Interop Respondents who indicated that it was the primary factor that is driving their organization to redesign their data center LAN.

Factor	Percentage
To reduce the overall cost	22.4%
To support more scalability	11.6%
To create a more dynamic data center	11.6%
To support server virtualization	11.2%
To reduce complexity	9.9%
To make it easier to manage and orchestrate the data center	9.2%
To support our storage strategy	7.5%
To reduce the energy requirements	6.5%
Other (please specify)	6.1%
To make the data center more secure	4.1%

Table 4.3: Factors Driving Data Center LAN Redesign

The data in Table 4.3 indicates that a broad range of factors are driving IT organizations to re-design their data center LANs. However, as is so often the case: The primary factor driving IT organizations to re-design their data center LAN is the desire to reduce cost.

Third Generation Data Center LAN Architecture and Technology Options

During the transition from First Generation LANs to Second Generation LANs there was considerable debate over the underlying physical and data link technologies. Alternatives included Ethernet, Token Ring, FDDI/CDDI and ATM. One of the few aspects of Third Generation Data Center LANs that is not up for debate is that they will be based on Ethernet. In fact, the Third Generation LAN will provide the possibility of exploiting Ethernet as the single data center switching fabric, eventually

displacing special purpose fabrics, such as Fibre Channel for storage networking and InfiniBand for ultra low latency HPC cluster interconnect.

Below is a listing of some of the primary objectives of a Third Generation Data Center LAN and an analysis of the various alternatives that IT organizations have relative to achieving those objectives.

Improved Topologies for Server-to-Server Communications

Many of the on-going IT initiatives that are aimed at improving the cost-efficiency of the enterprise data center (e.g., server virtualization, SOA, Web 2.0, access to shared network storage, and the exploitation of HPC and cluster computing) are placing a premium on IT organizations being able to provide a highly reliable, low latency, high bandwidth communications among both physical and virtual servers. Whereas the hub and spoke topology of the traditional three tier Second Generation LAN was optimized for client-to-server communications (sometimes referred to as 'north-south traffic'), it is decidedly sub-optimal for server-to-server communications, which is sometimes referred to as "east-west" traffic.

One approach for improving server-to-server communications is to flatten the network from three tiers to two tiers consisting of access layer and aggregation/core layer switches.

A two-tier network reduces the number of hops between servers, reducing latency and potentially improving reliability. The typical two-tier network is also better aligned with server virtualization topologies where VLANs may be extended throughout the data center in order to support dynamic VM migration at Layer 2.

Two tier networks require switches that have very high densities of high-speed ports and a higher level of reliability to protect the soaring volumes of traffic flowing through each switch. Reliability and availability also require redundant switch configurations in both tiers of the network. Some of the common characteristics of two tier networks are described in the following subsections.

High Port Density and Port Speed

The network I/O requirements of multi-core physical servers that have been virtualized are beginning to transcend the capacity of GbE and multi-GbE aggregated links. As the number of cores per server increases, the number of VMs per physical server can increase well beyond the 5-10 VMs/server that are typical today. With more VMs per server, I/O requirements increase proportionally. Thankfully the traditional economics of Ethernet performance improvement[3] is falling into place for 10 Gigabit Ethernet (10 GbE). As a result, Third Generation data center LAN

[3] Ethernet typically provides a 10x higher performance for a 3-4x increase in cost. This is an example of how Moore's Law impacts the LAN.

switches will need to support high densities of 10 GbE ports to provide connectivity for high performance virtualized servers, as well as an adequate number of 10 GbE ports (and 40 GBE and 100 GbE ports when these are available) for uplinks connecting the access switches to the core tier.

As noted, second generation LAN switches had fairly low port density. In contrast:

> ***The current generation of switches has exploited advances in switch fabric technology and merchant silicon switch-on-a-chip integrated circuits (ICs) to dramatically increase port densities.***

Modular data center switches are currently available with up to 256 non-blocking 10 GbE ports, while the typical maximum port density for stackable switches (generally based on merchant silicon) is 48 10 GbE ports. Today, high-speed uplinks are comprised of multiple 10 GbE links that leverage Link Aggregation (LAG)[4]. However, a 40 GbE uplink typically offers superior performance compared to a 4 link 10 GbE LAG. This is because the hashing algorithms that load balance traffic across the LAG links can easily yield sub-optimal load distribution where a majority of traffic is concentrated in a small number of flows. Most high performance modular switches already have a switch fabric that provide 100 Gbps of bandwidth to each line card, which means that when 40 GbE and 100 GbE line cards are available, these can be installed on existing modular switches, preserving the investment in these devices. Most vendors of modular switches expect to be shipping 40 GbE line cards by the middle of 2011, while 100 GbE line cards will take until 2012 or 2013.

In the case of stackable Top of Rack (ToR) switches, adding 40 or 100 GbE uplinks will generally require new switch silicon, which means that existing ToR switches will probably need to be swapped out in order to support the next generation of uplink speeds. The next generation of ToR switches with 40-48 10 GbE ports and 2-4 40 GbE ports are expected to be available in the second half of 2011.

High Availability

As previously noted, IT organizations will be implementing a growing number of VMs on high performance multi-core servers.

> ***The combination of server consolidation and virtualization creates an "all in one basket" phenomenon that drives the need for highly available server configurations and highly available data center LANs.***

One approach to increasing the availability of a data center LAN is to use a combination of redundant subsystems within network devices such as LAN switches in conjunction with redundant network designs. A high availability modular switch can provide redundancy in the switching fabric modules, the route processor modules, as well as the cooling fans and power supplies. In contrast, ToR switches are generally

[4] www.ieee802.org/3/hssg/public/apr07/frazier_01_0407.pdf

limited to redundant power supplies and fans. Extensive hardware redundancy is complemented by a variety of switch software features, such as non-stop forwarding, that ensure minimal disruption of traffic flow during failovers among redundant elements or during software upgrades. Modular switch operating systems also improve availability by preventing faults in one software module from affecting the operation of other modules.

The bandwidth efficiency of Layer 2 networks with redundant links can be greatly improved by assuring that the parallel links from the servers to the access layer and from the access layer to the core layer are always in an active-active forwarding state. This can be accomplished by eliminating loops in the logical topology without resorting to the Spanning Tree Protocol (STP). In the current state of evolution toward a Third Generation data center LAN, loops can be eliminated using switch virtualization and multi-chassis LAG (MC LAG) technologies, which are described below.

Switch Virtualization

With switch virtualization, two or more physical switches are made to appear to other network elements as a single logical switch or virtual switch, with a single control plane.

In order for multiple physical switches to form a virtual switch, they need a virtual switch link (VSL) or interconnect (VSI) that supports a common control plane and data flows between the members of the virtual switch. In redundant configurations, connections between end systems and virtual access switches and between virtual access switches and virtual aggregation switches are based on multi-chassis (MC) link aggregation group (LAG) technology[5], as shown in Figure 4.1. MC LAG allows the links of the LAG to span the multiple physical switches that comprise a virtual switch. From the server perspective, links to each of the physical members of a virtual access switch appear as a conventional LAG or teamed links, which means that switches can be virtualized without requiring any changes in the server domain.

The combination of switch virtualization and multi-chassis LAG can be used to create a logically loop-free topology.

This means that data center LANs can be built without using the spanning tree protocol (STP) and first hop router redundancy protocols (e.g., VRRP). This is important because these protocols prevent all available forwarding resources in a redundant network design from being simultaneously utilized.

In Figure 4.1, loops are eliminated because, from a logical perspective, there are only two switches with a single LAG from the server to the access switch and a single LAG from the access switch to the aggregation switch. The traffic load to and from each server is load balanced across the two links participating in the multi-chassis

[5] http://en.wikipedia.org/wiki/Link_aggregation

LAG connecting each server to the virtual access switch. Therefore, both server connections are actively carrying traffic in both directions rather than being in an active state for some VLANs and in an inactive state for others. In the same fashion, traffic between the access virtual switch and the aggregation virtual switch is load balanced across all four physical links connecting these devices. Both physical switches participating in the aggregation layer virtual switch are actively forwarding traffic to the network core (not shown in the figure). The traffic is load balanced via the LAG hashing algorithms rather than being based on VLAN membership, as is the case with more traditional redundant LAN designs. The virtual switch not only improves resource utilization but also enhances availability because the relatively long convergence times of STP topology calculations are circumvented. Virtual switch technology also simplifies management because multiple physical switches can be managed as a single entity.

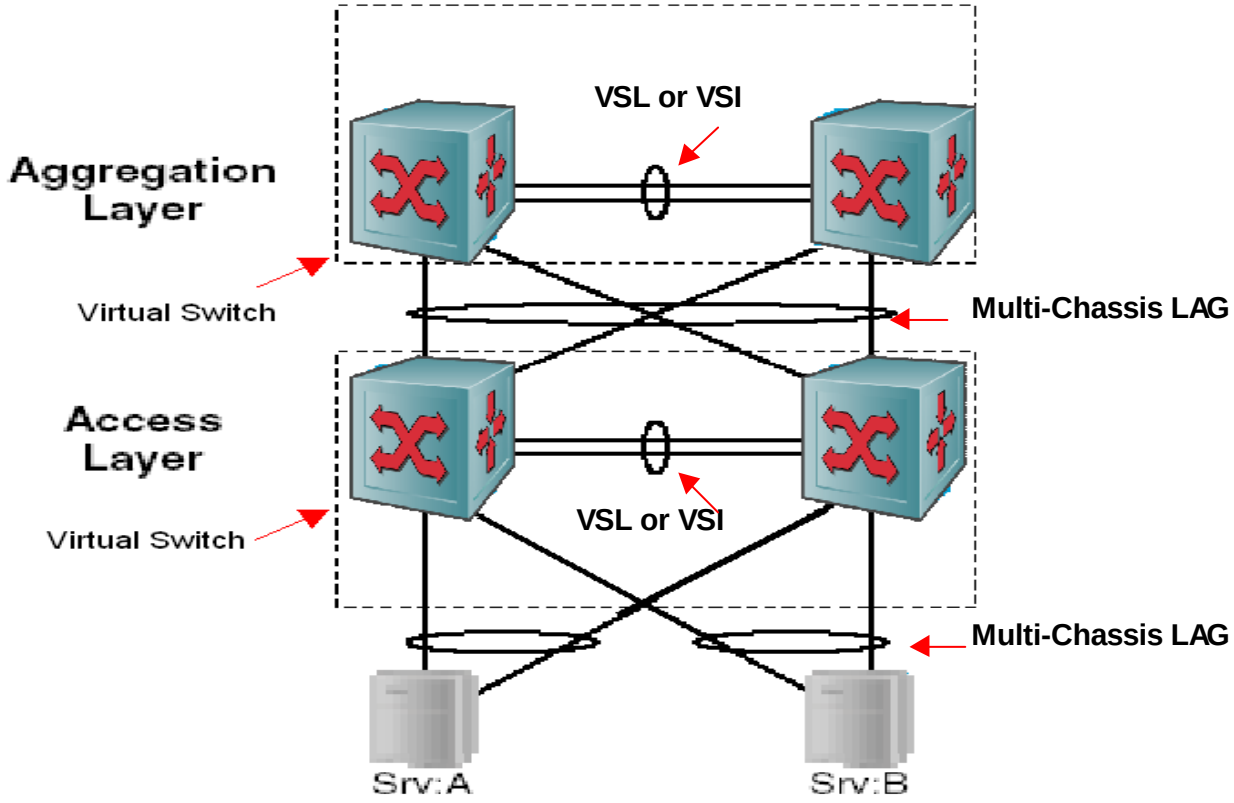

Figure 4.1: Switch Virtualization and Multi-Chassis LAG

Most vendors of data center switches support switch virtualization and MC LAG in their ToR and modular switches, and these technologies are fully utilized in the two-tier LAN designs that they are currently recommending to enterprise customers. As a result, most two tier LAN designs being proposed will not be based on STP for loop control. There are some differences among vendors in the VSL/VSI technology and in the LAG hashing algorithms. For example, some vendors of stackable ToR switches take advantage of the stacking interconnect as the VSL/VSI link, while other vendors will use 10 GbE or 40 GbE ports (when available) for VSL/VSI. Most LAG

implementations conform to the IEEE 802.3ad standard. However, LAG hashing algorithms are outside the 802.3ad standard and more sophisticated hashing algorithms can provide differentiation by improving load balancing across LAG links. In addition, there are some differences in the number of ports or link that can participate in a LAG. Some vendors support up to 32 links per LAG, while 8 links per LAG is the most common implementation.

It must be noted that two-tier LANs and switch virtualization are far from the final word in the design of data center networks. Standards bodies have been working on technologies that will allow active-active traffic flow and load balancing of Layer 2 traffic in networks of arbitrary switch topologies. TRILL (Transparent Interconnection of Lots of Links) is an Internet Engineering Task Force (IETF) project to develop a Layer 2 shortest-path first (SPF) routing protocol for Ethernet. A similar effort is being pursued by the IEEE 802.1aq working group which is defining a standard for shortest path bridging of unicast and multicast frames (based on the IS-IS protocol) and which supports multiple active topologies. With TRILL or 802.1aq, it would be possible to achieve load-balanced, active-active link redundancy without having to resort entirely to switch virtualization and VSL/VSI interconnects.

SPF bridging should have major implications for data center LAN designs and most of the larger switch vendors are well along in developing switches that can support SPF bridging and network designs based on these technologies. While the TRILL standard is possibly still a year or more away, some vendors are preparing pre-standard and proprietary enhancements for introduction well before the standard is ratified. It may well turn out that two-tier networks based on switch virtualization and MC LAG are just a mid-way point in the evolution of the Third Generation LAN.

With technologies like TRILL, the difference between access switches and core switches may shrink significantly.

As a result, the switch topology may shift from a two-tier hub and spoke, such as the one in Figure 4.1, to a highly meshed or even fully meshed array of switches that appears to the attached devices as a single switch. SPF bridging can support a variety of other topologies, including the fat tree switch topologies[6] that are popular in cluster computing approaches to HPC. Fat trees are also used by Ethernet switch vendors to build high density, non-blocking 10 GbE switches using merchant silicon switch chips. Figure 4.2a shows how a 48 port 10 GbE TOR switch can be constructed using six 24-port 10 GbE switch chips. By increasing the number of leaf and spine switches, larger switches can be constructed, The maximum density switch that can be built with a two-tier fat tree architecture based on 24 port switch chips has 288 ports. A number of high density 10 GbE switches currently on the market use this design approach.

[6] www.mellanox.com/pdf/.../IB_vs_Ethernet_Clustering_WP_100.pdf

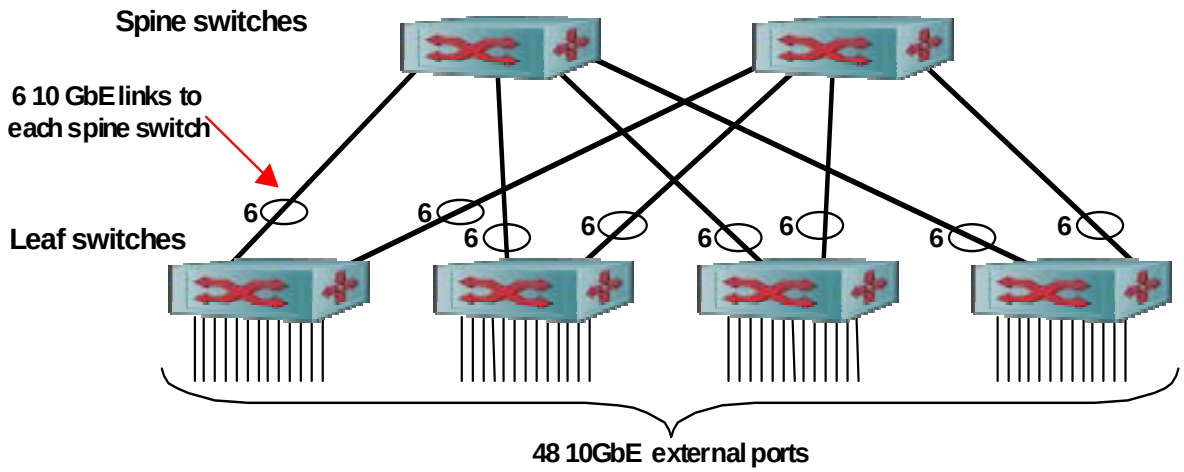

Figure 4.2a: TOR Switch Fat Tree Internal Architecture

With the emergence of TRILL/SPF bridging, a number of Ethernet switch vendors are adopting the fat tree as the topology for the 10 GbE data center LAN. With core/aggregation switches playing the role of spine switch and access switches playing the role of leaf switches. TRILL/SPF bridging allows traffic between two access switches to be load balanced across the numerous parallel paths through the core/aggregation layer switches (Equal Cost Multi-Path bridging). Using 72 48-port 10 GbE TOR switches as the common building blocks, a two- tier fat tree data center network can be built with 1,152 10 GbE ports. With 48 256 port 10 GbE switches, a two-tier data center fat tree network with 8,192 10 GbE ports would be possible, as shown in Figure 4.2b, The upper limit for two-tier fat trees based on 256 port switches is 32,768 ports using 384 switches

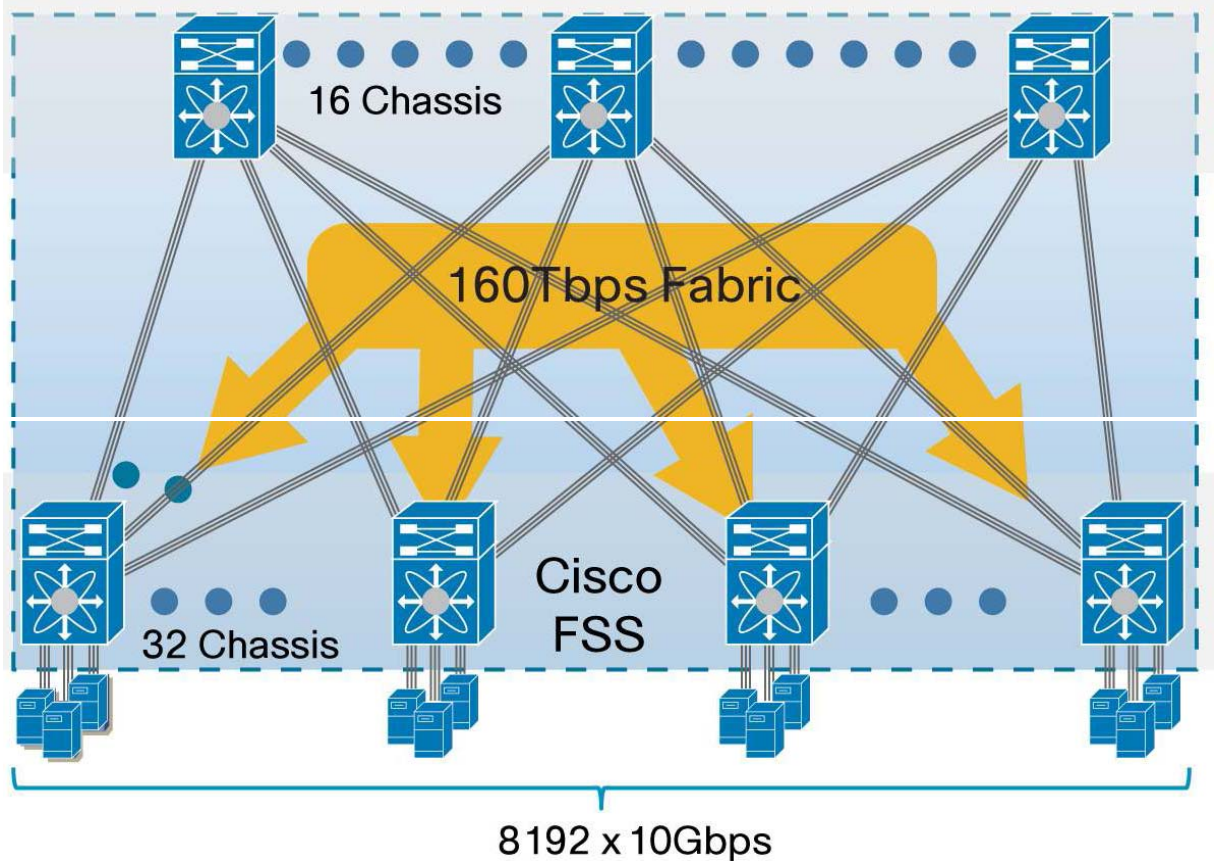

Figure 4.2b Two-Tier Fat Tree Data Center LAN with 8,192 10 GbE Ports
Source: Cisco Systems

Controlling and Managing Inter-VM Traffic

With server virtualization, each physical server is equipped with a hypervisor-based virtual switching capability that allows connectivity among VMs on the same physical platform. Traffic to external destinations also traverses this software switch. From the network perspective, the hypervisor vSwitch poses a number of potential problems:

1. The vSwitch represents another tier of switching that needs to be configured and managed, possibly requiring an additional management interface. This can partially defeat an effort to flatten the network to two–tiers.

2. The vSwitch adds considerable complexity, because there is an additional vSwitch for every virtualized server.

3. vSwitch control plane functionality is typically quite limited compared to network switches, preventing a consistent level of control over all data center traffic

4. As more VMs per server are deployed, the software switch can place high loads on the CPU, possibly starving VMs for compute cycles and becoming an I/O bottleneck.

5. VM-VM traffic on the same physical server is isolated from the rest of the network, making these flows difficult to monitor and control in the same fashion as external flows.

6. The vSwitch functionality and management capabilities will vary by hypervisor vendor and IT organizations are increasingly deploying hypervisors from multiple vendors.

The vSwitch presents a number of concerns related to management, security, functionality and organizational responsibilities.

There are two approaches to the problems posed by the early generation vSwitch: Distributed Virtual Switching (DVS) and Edge Virtual Bridging (EVB). With DVS, the control and data planes of the embedded vSwitch are decoupled. This allows the data planes of multiple vSwitches to be controlled by an external centralized management system that implements the control plane functionality. Decoupling the data plane from the control plane makes it easier to tightly integrate the vSwitch control plane with the control planes of physical access and/or aggregation switches and/or the virtual server management system. Therefore, DVS can simplify the task of managing a large number of vSwitches, and improve control plane consistency, but it doesn't address the other issues listed above.

With EVB, all the traffic from VMs is sent to the network access switch. If the traffic is destined for a VM on the same physical server, the access switch returns the packets to the server over the same port on which it was received (e.g., a "hair pin turn"). With Edge Virtual Bridging, the hypervisor is relieved from all switching functions, which are now performed by the physical access network. With EVB, the vSwitch now performs the simpler function of aggregating hypervisor virtual NICs to a physical NIC. Basic EVB can be supported by most existing access switches via a relatively simple firmware upgrade. The IEEE 802.1bg Working Group is creating an EVB standard based on a technology known as Virtual Ethernet Port Aggregator (VEPA) dealing with hair-pin turns and a definition of a multi-channel service for remote ports to access local VMs. A companion effort, IEEE 802.1bh Port Extension is defining a tagged approach to deal with frame replication issues in the EVB. EVB/VEPA standards supported in switches and hypervisors will address all of the issues listed above.

Essentially all vendors of data center switches support the IEEE's EVB standards efforts. Some vendors are waiting until the standard is finalized and are supporting hypervisor vSwitches in the interim. Other vendors have pre-standard implementations of basic EVB/VEPA already available or under development.

Network Convergence/Fabric Unification

In contrast to Second Generation Data Center LANs:

> ***A key characteristic of Third Generation Data Center LANs will be the convergence of block-level storage and data traffic over a common high-speed Ethernet data center switching fabric.***

This unified fabric offers significant cost savings in multiple areas including converged network adapters on servers and reductions in rack space, power and cooling capacity, cabling, and network management overhead.

Traditional Ethernet, however, only provides a best effort service that allows buffers to overflow during periods of congestion and which relies on upper level protocols (e.g., TCP) to manage congestion and recover lost packets through re-transmissions. In order to emulate the lossless behavior of a Fibre Channel (FC) SAN, Ethernet needs enhanced flow control mechanisms that eliminate buffer overflows for high priority traffic flows, such as storage access flows. Lossless Ethernet will be based on the following standards, which are commonly referred to as IEEE Data Center bridging (DCB):

- **IEEE 802.1Qbb Priority-based Flow Control (PFC)** allows the creation of eight distinct virtual link types on a physical link, with each virtual link mapped to an 802.1p traffic class. Each virtual link can be allocated a minimum percentage of the physical link's bandwidth. Flow is controlled on each virtual link via the pause mechanism which can applied on a per priority basis to prevent buffer overflow, eliminating packet loss due to congestion at the link level. In particular, block-level or file-level storage traffic on one of the virtual lanes can be protected from loss by pausing traffic on one or more of the remaining lanes.

- **IEEE 802.1Qau Congestion Notification (CN)** is a traffic management technique that eliminates congestion by applying rate limiting or back pressure at the edge of the network in order to protect the upper network layers from buffer overflow. CN is intended to provide lossless operation in end-to-end networks that consist of multiple tiers of cascaded Layer 2 switches, such as those typically found in larger data centers for server interconnect, cluster interconnect and extensive SAN fabrics.

- **IEEE 802.1Qaz Enhanced Transmission Selection (ETS)** will specify advanced algorithms for allocation of bandwidth among traffic classes including the priority classes supported by 802.1Qbb and 802.1Qau. While the queue scheduling algorithm for 802.1p is based on strict priority, ETS will extend this by specifying more flexible drop-free scheduling algorithms. ETS will therefore provide uniform management for the sharing of bandwidth between congestion managed classes and traditional classes on a single bridged network. Priorities using ETS will coexist with priorities using 802.1Qav queuing for time-sensitive streams. ***Data***

Center Bridging Exchange (DCBX) protocol is also defined in the 802.1Qaz standard. DCBX protocol is an extension of the Link Layer Discovery Protocol (LLDP) that allows neighboring network elements to exchange request and acknowledgment messages to ensure consistent DCB configurations. DCBX is also used to negotiate capabilities between the access switch and the adapter and to send configuration values to the adapter.

- **IEEE 802.1aq Shortest Path Bridging (SPF)** is a standard for shortest path bridging of unicast and multicast frames (based on the IS-IS protocol) supporting multiple active topologies. SPF is part of the IEEE standards efforts relative to the data center, but is not strictly required for standards-compliant lossless Ethernet.

DCB Lossless Ethernet will play a key role in supporting Fibre Channel over Ethernet (FCoE) technology that will allow the installed base of Fibre Channel storage devices and SANs to be accessed by Ethernet-attached servers (with converged FCoE network adapters) over the unified data center switching fabric. DCB will benefit not only block-level storage, but also all other types of loss and delay sensitive traffic. In the storage arena, DCB will improve NAS performance and will make iSCSI SANs based on 10/40/100 GbE a more competitive alternative to Fibre Channel SANs at 2/4/8 Gbps. In order to take full advantage of 10 GbE and higher Ethernet bandwidth, servers accessing iSCSI storage resources may also need intelligent converged NICs that offload iSCSI and TCP/IP processing from the host.

All data center vendors are planning to support the DCB standards when they are available. In some cases the timing of the availability of that support may differ between access and core switches. In addition, some vendors are offering pre-standard support for DCB capabilities, including PCF, CN, ETS, and DCBX.

Fibre Channel over Ethernet (FCoE) is an industry standard that is being developed by the International Committee for Information Technology Standards (INCITS) T11 committee.

The FCoE protocol specification maps Fibre Channel upper layer protocols directly over a bridged Ethernet network. FCoE provides an evolutionary approach to the migration of FC SANs to an Ethernet switching fabric while preserving Fibre Channel constructs and providing reliability, latency, security, and traffic management attributes similar to those of native FC. FCoE also preserve investments in FC tools, training, and SAN devices (e.g., FC switches and FC attached storage). Implementing FCoE over a lossless Ethernet fabric requires converged server network adapters (CNAs with support for both FCoE and IP) and some form of FC Forwarding Function (FCF) to provide attachment to native FC devices (FC SAN switches or FC disk arrays). FCF functionality can be provided by a FCoE switch with both Ethernet and FC ports or by a stand alone gateway device attached to a FCoE passthrough switch, as shown in Figure 4.3.

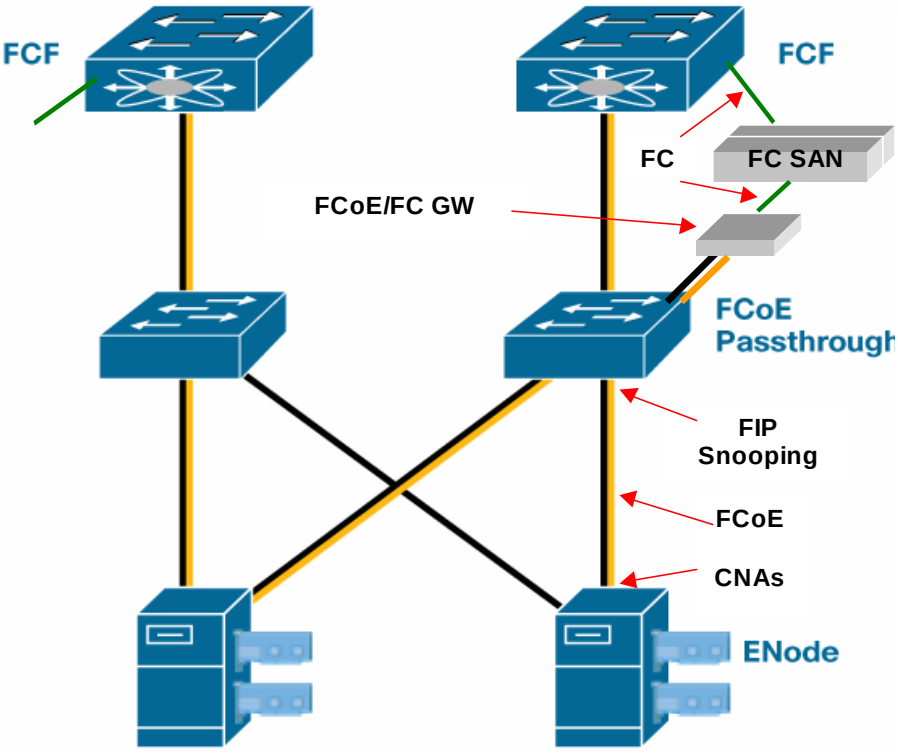

Figure 4.3: FCoE Converged LAN
Source: Cisco Systems

As shown in the figure, End Nodes (servers) do not need to connect directly to a FCF capable switch. Instead the FCoE traffic can pass through one or more intermediate FCoE passthrough switches. The minimal requirements for a simple FCoE passthrough switch is support for lossless Ethernet or DCB. The FCoE Initialization Protocol (FIP) supports handshaking between a FCoE End Node and an FCF in order to establish and maintain a secure virtual FC link between these devices, even if the end-to-end path traverses FCoE passthrough switches. For DCB passthrough switches that support FIP Snooping, the passthrough switches can inspect the FIP frames and apply policies based on frame content. FIP Snooping can be used to enhance FCoE security by preventing FCoE MAC spoofing and allowing auto-configuration of ACLs.

As this discussion illustrates:

> ***There are several levels of support that data center switch vendors can provide for FCoE.***

For example:

1. The lowest level of support is FCoE passthrough via lossless Ethernet or DCB alone.

2. The next step up is to add FIP Snooping to FCoE passthrough switches

3. A third level of support is to add a standalone FCF bridges/gateways to front end FC SAN switches or disk arrays.

4. The highest level of support is to provide DCB and FIP Snooping for FCoE passthrough switches and also to provide FCoE switches that incorporate FCF ports, creating hybrid switches with both DCB Ethernet and native FC ports.

Most vendors of Ethernet data center switches that do not also have FC SAN switches among their products are planning FCoE support at levels 1, 2, or 3 described above. In fact, most of these Ethernet-omly vendors are considerably more enthusiastic about iSCSI SANs over 10/40/100 GbE than they are about FCoE.

The primary drivers of FCoE are the vendors that offer both Ethernet and FC products.

These are the vendors that are already shipping pre-standard lossless 10 GbE Ethernet switches and hybrid lossless 10 GbE/FCF switches. Even among the vendors providing early support for FCF there are some significant differences, as shown in Figure 4.4.

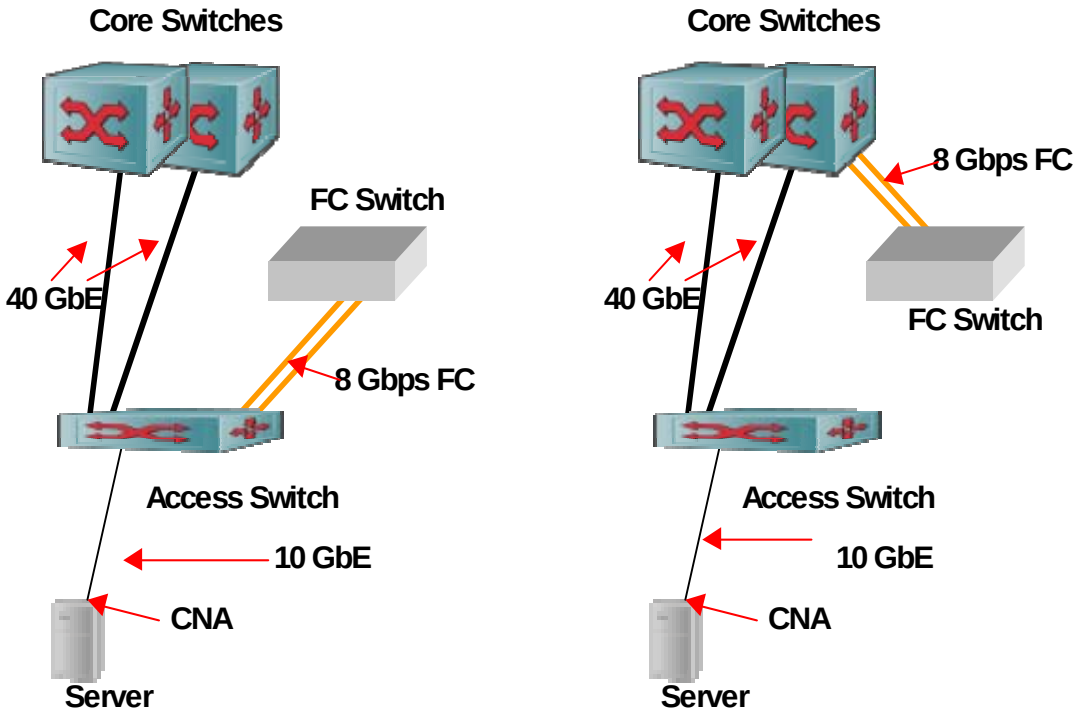

Figure 4.4: FCF Support Options

The left side of the figure shows single hop FCoE with the FCF function integrated into the access switch. It would also be possible to use intervening FCoE/FCF gateways (either standalone or incorporated in the FC switch) which would be connected to the access switch via 10 GbE, making the access switch an FCoE passthrough switch, as shown in the previous figure. The advantage of single hop FCoE is that the storage traffic doesn't compete for bandwidth in the uplinks or the core switches and the core switches aren't required to support DCB or FIP Snooping. The right side of the figure shows multihop FCoE with the FCF function integrated into the core switch, and the access switch in FCoE passthrough mode. Again it would be possible to use FCoE/FCF gateways (either standalone or incorporated in the FC switch) connected to the core switch via 10 GbE. FC SANs and disk arrays connected at the core offer the advantage of a more centralized pool of storage resources that can be shared across the data center LAN.

Network Support for Dynamic Creation and Migration of VMs

As noted earlier, the requirement to support the dynamic creation and movement of VMs is one of the primary factors driving IT organizations to redesign their data center LANs. As noted earlier, the requirements for VM migration with VLAN boundaries has provided a major impetus for flattening the LAN with two-tier designs featuring Layer 2 connectivity end-to-end.

Many of the benefits of cloud computing depend on the ability to dynamically provision VMs and to migrate them at will among physical servers located in the same data center or in geographically separated data centers. The task of creating or 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 VLAN memberships, QoS settings, and ACLs) is established or transferred in a timely fashion. In many instances today, these network configuration or reconfigurations involves the time-consuming manual process involving 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 most common approach to automating the manual processes involved in VM provisioning and migration is based on communication between the Hypervisor Management system and the switch element management system (EMS) via APIs supported by both vendors.***

When a Virtual Machine is created or when a virtual machine move is initiated, the Hypervisor manager signals to the EMS that the event is about to occur and provides a partial VM network profile including a virtual MAC, VLAN memberships, and the target hypervisor. Based on existing policies, the EMS extends the VM network profile to include appropriate QoS and security parameters such as ACLs. The EMS can then determine the target hypervisor's access switch and configure or reconfigure it accordingly. Where VLANs need to be created, the EMS can also create these on the uplinks and neighboring switches, as appropriate. In a similar manner, when a VM is deleted from a hypervisor, the EMS can remove the profile and then prune the VLAN as required. All of these processes can be triggered from the hypervisor.

Most data center switch vendors have already implemented some form of VM network profile software, including linking their switches to a least one brand of hypervisor. Some differences exist between the range of hypervisors supported and the APIs that were used. Distribution of VM network profiles is only one of many management processes that can benefit greatly from automation, so it would benefit IT departments to develop expertise in open APIs and powerful scripting languages that can be exploited to streamline time-consuming manual processes and thereby reduce operational expense while improving the ability of the data center to dynamically reallocate its resources in response to changes in user demand for services.

A somewhat different approach to automating data center configuration, including the provisioning and migration of VMs is based on Orchestration engines, which are discussed in more detail in the management section of this report. Service orchestration is a centralized server function that can automate many of the manual tasks involved in provisioning and controlling the capacity of dynamic virtualized services. In the case of VM provisioning and migration, the Orchestration engine would function as the point of integration between the network device's EMS and the hypervisor management system. Orchestration solutions are available from a number of network management vendors and hypervisor vendors.

The data center LAN is on the cusp of a number of quite dramatic technology developments, as summarized in Table 4.4. As shown in the table, most the items on this list are still in flux and require additional development, and/or additional work from the standards bodies[7].

[7] Exceptions to this statement are entries number 1, 2, 4 and to some extent 11.

Technology Development	Status
1: Two-tier networks with Layer 2 connectivity extending VLANs across the data center.	On-going deployment
2: Reduced role for blade switches to eliminate switch tier proliferation.	On-going
3: Changing role for the hypervisor vSwitch as a port aggregator (VEPA) for EVB, essentially eliminating the vSwitch tier.	A standard is in progress. Pre-standard implementations are occurring.
4: STP displaced by switch virtualization and multi-chassis LAG technology.	On-going deployment
5: Multi-core servers with notably more VMs per server and 10 GbE connectivity to the LAN.	Early adoption stage.
6: 40 GbE and 100 GbE uplinks and core switches.	A standard is in place: 40 GbE due in 2011 100 GbE due in 2012
7: DCB delivering lossless Ethernet for 10 GbE and higher speed Ethernet	Standards are in progress. Pre-standard implementations are occurring.
8: FCoE approach to fabric unification	FCoE standard is in place. Early implementations of FCoE are occurring over pre-standard DCB.
9: 10 GbE iSCSI approach to fabric unification	Early implementations over pre-standard DCB.
10: TRILL/SPB enabling new data center LAN topologies; e.g., fully meshed, fat tree.	Standards are in progress. Pre-standard implementations are imminent.
11: Management tools that integrate, coordinate, and automate provisioning and configuration of server, storage and network resource pools	These are proprietary and have varying levels of maturity.

Table 4.4 Status of Data Center Technology Evolution

5. The Wide Area Network (WAN)

Background

The modern WAN got its start in 1969 with the deployment of ARPANET which was the precursor to today's Internet. The technology used to build the Internet began to be commercialized in the early 1970s with the development of X.25 based packet switched networks.

In addition to the continued evolution of the Internet, the twenty-year period that began in 1985 saw the deployment of four distinct generations of enterprise WAN technologies[8]. For example, in the mid to late 1980s, it became common for enterprise IT organizations to deploy integrated TDM-based WANs to carry both voice and data traffic. In the early 1990s, IT organizations began to deploy Frame Relay-based WANs. In the mid to late 1990s, some IT organizations replaced their Frame Relay-based WANs with WANs based on ATM (Asynchronous Transfer Mode) technology. In the 2000s, many IT organizations replaced their Frame Relay or ATM-based WANs with WANs based on MPLS. Cost savings was the primary factor that drove the adoption of each of the four generations of WAN technologies. The cost savings, however, were very modest when compared to the price performance improvements that are associated with local area networking.

However, in contrast to the volatility of this twenty-year period:

Today there is not a new generation of WAN technology in development.

Relative to the deployment of new WAN services what sometimes happens in the current environment is that variations are made to existing WAN technologies and services. An example of that phenomenon is Virtual Private LAN Service (VPLS). As described below, within VPLS an Ethernet frame is encapsulated inside of MPLS. While creating variations on existing services can result in significant benefits, it does not produce fundamentally new WAN services.

The Webtorials Respondents were given a set of eleven WAN services[9] and asked to indicate the extent to which they currently utilize each WAN service. They were given a five-point scale:

1. None
2. Minimal
3. Some
4. Quite a bit
5. Extensive

[8] An enterprise WAN is designed to provide for connectivity primarily within the enterprise and between the enterprise and key contacts such as partners. This is in contrast to the Internet that is designed to provide universal connectivity.

[9] The eleven WAN services are listed in column one of Table 5.2.

The survey results indicate that The Webtorials Respondents utilize on average 4.8 WAN services.

The typical IT organization utilizes a wide range of WAN services.

The survey question included Frame Relay and ATM among the set of WAN services. In the not too distant past, these services were widely deployed. However, over half of The Survey Respondents don't have any Frame Relay in their networks and almost two thirds of The Survey Respondents don't have any ATM in their networks.

Looking just at the WAN services that The Webtorials Respondents utilize either quite a bit or extensively, they average 2.3 WAN services.

The primary WAN services used by IT organizations are MPLS and the Internet.

While IT organizations make extensive use of the Internet, quality issues in the public Internet and in consumer-class ISP services generally prevent Internet VPNs from meeting the reliability standards of enterprise IT departments. As a result, Internet VPNs are most often used only as a backup connection to a primary private WAN circuit. This is unfortunate because the shortfall in quality is fairly small when compared to the dramatic cost savings and additional bandwidth that can be realized by using broadband connections. In cases where Internet-based VPNs are deployed today, businesses typically prefer an expensive T1/E1 for access, since a single xDSL often results in a lower level of availability and performance predictability.

In many cases, Internet-based VPNs that use DSL for access are 'good enough'.

Drivers of Change

The Webtorials Respondents were asked to indicate the anticipated change in their organization's use of the eleven WAN services. For each WAN service, Table 5.1 indicates the percentage of The Webtorials Respondents that expect to reduce their use of that service; to make no change in the use of that service; and to increase their use of that service.

Service	% Decrease	% Stay the Same	% Increase
Private lines between your sites	26.2%	62.7%	11.1%
Private lines to access network services	20.7%	62.8%	16.5%
Frame Relay	33.9%	64.4%	1.7%
ATM	22.3%	71.4%	6.3%
MPLS	4.8%	43.2%	52.0%
VPLS	8.0%	68.1%	23.9%
Internet-based VPNs with T1/T3/OC-3/OC-12 for access	12.4%	64.5%	23.1%
Internet-based VPNs with DSL or cable for access	10.4%	52.0%	37.6%
IP VPN	4.9%	61.5%	33.6%
An Internet overlay from a company like Akamai or CDNetworks	6.3%	83.0%	10.7%
Internet traffic to external sites	2.4%	57.3%	40.3%

Table 5.1: Expected Change in the Use of WAN Services

As previously noted, the primary WAN services used by IT organizations are MPLS and the Internet. As shown in Table 5.1, the majority of IT organizations are expecting to increase their use of MPLS. In addition, the majority of IT organizations are also increasing their use of one or more forms of Internet services.

One form of centralization of resources that is likely to drive a further increase in WAN traffic is desktop virtualization. To put the challenge of desktop virtualization into perspective, The Interop Respondents were asked about their organization's current and planned deployment of desktop virtualization. Their responses are shown in Table 5.2.

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

Table 5.2: The Percentage of Desktops that Already Have or Will be Virtualized

The data in Table 5.2 indicates that within the next year:

- The number of IT organizations that have virtualized the majority of their desktops will almost double.

- The number of IT organizations that have not implemented desktop virtualization will be cut in half.

As pointed out in ***Virtualization: Benefits, Challenges and Solutions***, desktop virtualization results in a number of new protocols, such as Teradici's PC-over-IP (PCoIP), transiting the WAN. As that report details, in many instances these new protocols consume considerable WAN bandwidth and are latency sensitive. The performance challenges associated with desktop virtualization are described in detail in a subsequent section of this report.

As previously noted, given that storage tends to follow Moore's Law, an IT organization may well be able to support a significant increase in storage requirements without a dramatic increase in cost. Unfortunately, WAN services do not follow Moore's Law.

> ***The price/performance of MPLS tends to improve by only a couple of percentage points per year.***

As such, IT organizations will not be able to support a significant increase in bandwidth requirements without a significant increase in cost. To put the challenge of supporting significant increases in WAN bandwidth into context, The Webtorials Respondents were asked how their budget this year for all WAN services compares to what it was last year. Their responses are contained in Figure 5.1

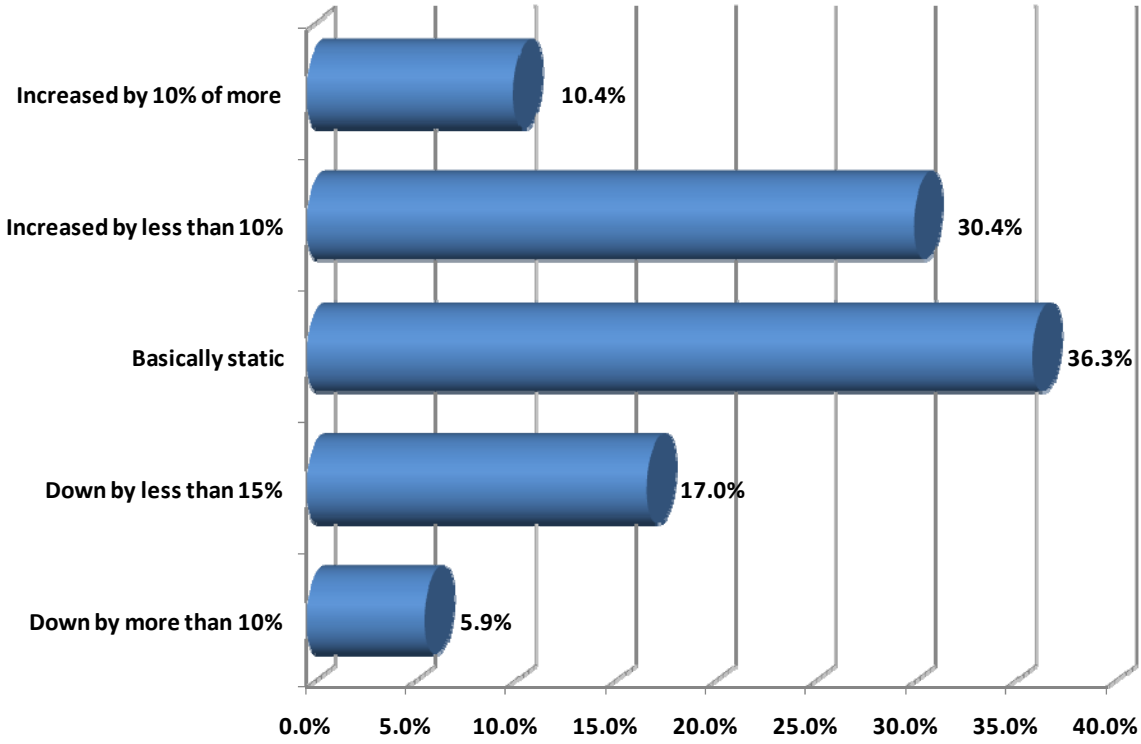

Figure 5.1: Anticipated Change in WAN Budgets

The data in Figure 5.1 shows that while over forty percent of IT organizations are experiencing an increase in their WAN budget, few are experiencing a significant increase. In addition, almost a quarter of IT organizations are experiencing a decrease in their WAN budget.

> ***IT organizations will not be able to support a significant increase in the use of WAN services with constrained budgets unless they make changes to how they use WAN services.***

Cloud Networking Without the Internet

There is a temptation to associate the WAN component of *cloud networking* exclusively with the Internet. For example, some definitions of cloud computing[10] assume that cloud services are always delivered over the Internet. Due to a variety of well-known issues (e.g., packet loss at peering points, BGP's inability to choose the path with the lowest delay, the TCP Slow start algorithm) the Internet often exhibits performance problems. As such, the Internet is not always the most appropriate WAN service to use to access cloud computing solutions. To put the use of the Internet into context, The Webtorials Respondents were asked to indicate which WAN service they would most likely use when accessing public and private

[10] http://en.wikipedia.org/wiki/Cloud_computing

cloud computing services over the next year. Their responses are shown in Table 5.3.

	The Internet	An Internet overlay from a company such as Akamai	A traditional WAN service such as MPLS	WAN Optimization combined with a traditional WAN service; e.g. MPLS
Public Cloud Computing Services	57.1%	9.0%	19.5%	14.3%
Private Cloud Computing Services	28.8%	4.0%	30.4%	36.8%

Table 5.3: WAN Services to Access Cloud Computing Services

The data in Table 5.3 indicates that IT organizations understand the limitations of the Internet relative to supporting cloud computing.

In many cases, the WAN service that IT organizations plan to use to support cloud computing is not the Internet.

WAN Design

IT organizations typically design their WAN based on the following criteria:

1. Minimize cost
2. Maximize availability
3. Ensure appropriate performance

One of the key WAN design challenges is that WAN traffic is typically comprised of widely varying traffic types. This includes traffic generated by enterprise applications that are business critical and delay-sensitive (e.g., SCM, ERP); highly-visible, delay-sensitive real-time applications (e.g., voice, video and telepresence); and the growing use of public cloud computing services (e.g., Salesforce.com, Amazon's EC2) and social networking sites; e.g., LinkedIn and Facebook.

As previously demonstrated, the majority of IT organizations utilize MPLS. One of the reasons for the popularity of MPLS is that the major suppliers of MPLS services offer a number of different classes of service (CoS) designed to meet the QoS requirements of the types of applications discussed in the preceding paragraph. Real-time applications are typically placed in what is often referred to as a Differentiated Services Code Point (DSCP) Expedited Forwarding class that offers minimal latency, jitter, and packet loss. Mission critical business applications are typically relegated to what is often referred to as a DSCP Assured Forwarding Class. Each class of service is typically associated with a service level agreement (SLA) that

specifies contracted ranges of availability, latency, packet loss and possibly jitter. Unfortunately, in many cases the SLAs are weak. In particular, it is customary to have the SLAs be reactive in focus; i.e., the computation of an outage begins when the customer opens a trouble ticket. In most cases, the carrier's SLA metrics are calculated as network-wide averages rather than for a specific customer site. As a result, it is possible for a company's data center to receive notably poor service in spite of the fact that the network-wide SLA metrics remain within agreed bounds. In addition, the typical level of compensation for violation of service level agreements is quite modest.

To gauge the effectiveness of SLAs that IT organizations receive from their network service providers (NSPs), The Webtorials Respondents were asked to indicate which of the following best describes the SLAs that they get from their NSPs for services such as MPLS.

- The SLAs go a long way towards ensuring that we get a quality service from the network service provider.
- The SLAs are better than nothing, but not by much.
- The SLAs are not worth the paper they are written on.

Their responses are shown in Figure 5.2.

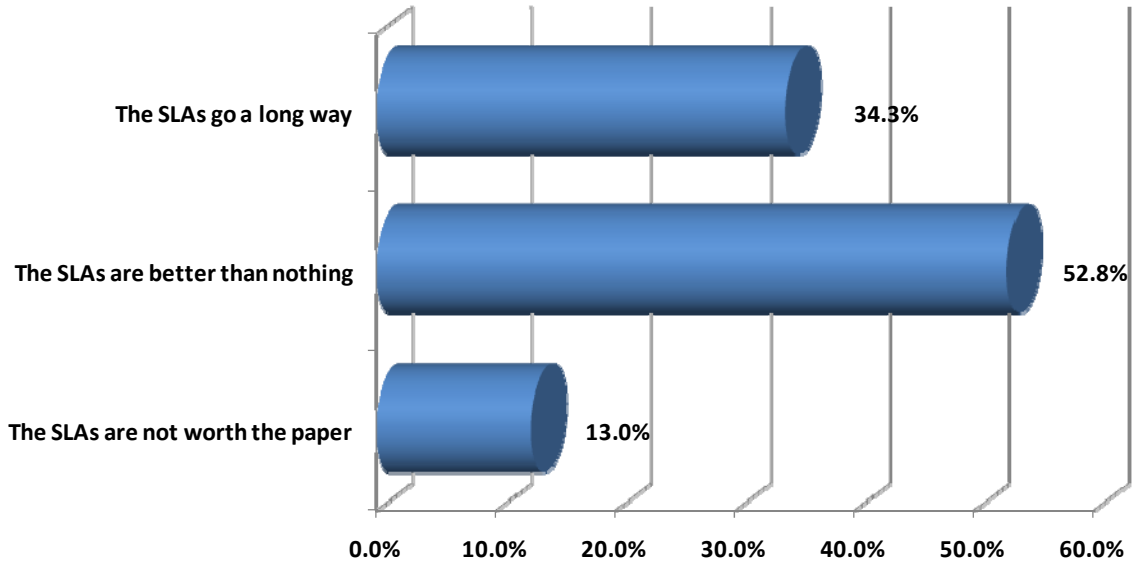

Figure 5.2: The Effectiveness of SLAs

The fact that two thirds of The Webtorials Respondents indicated that the SLAs that they receive from 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, demonstrates the weak nature of most SLAs.

The majority of IT organizations don't regard the SLAs that they receive from their NSPs as being effective.

While the pricing and other characteristics of MPLS and Internet services vary on an international basis, Table 5.4 depicts how these services typically compare.

	MPLS	Internet
Cost	High	Low
Availability	High	High
Performance	Predictable	Non-Predictable
QoS	Sophisticated	None

Table 5.4: Characteristics of MPLS and Internet Services

One obvious conclusion that can be drawn from Table 5.4 is that IT organizations won't be able to satisfy the primary WAN design criteria with a simple network design that supports all traffic types. For example a network design that uses MPLS to interconnect an organization's large sites and uses a traditional approach to Internet access (e.g., a single DSL circuit) to connect their small sites would have two key limitations. This design would overpay for WAN services at the organization's large sites while not supporting any QoS functionality at their small sites.

WAN Service Alternatives

As noted, there is no new generation of WAN technology currently under development. However, as is described below, there are a number of WAN service alternatives that are variations on existing WAN technologies and services that better enable IT organizations to meet their WAN design goals.

VPLS

As previously mentioned:

VPLS represents the combination of Ethernet and MPLS.

While VPLS is not widely implemented today, the data in Table 5.1 indicates that roughly one quarter of IT organizations will increase their use of VPLS over the next year.

VPLS is a class of VPN that supports the connection of customer edge (CE) Layer 2 switches at multiple sites in a single bridged (multipoint-to-multipoint) domain over a provider's IP/MPLS network, as shown in Figure 5.3. VPLS presents an Ethernet interface to customers, that simplifies the LAN/WAN boundary for Service Providers and customers, and enables rapid and flexible service provisioning. All sites in a

VPLS appear to be on the same LAN, regardless of location. A companion technology, Virtual Private Wire Services (VPWS), provides point-to point services.

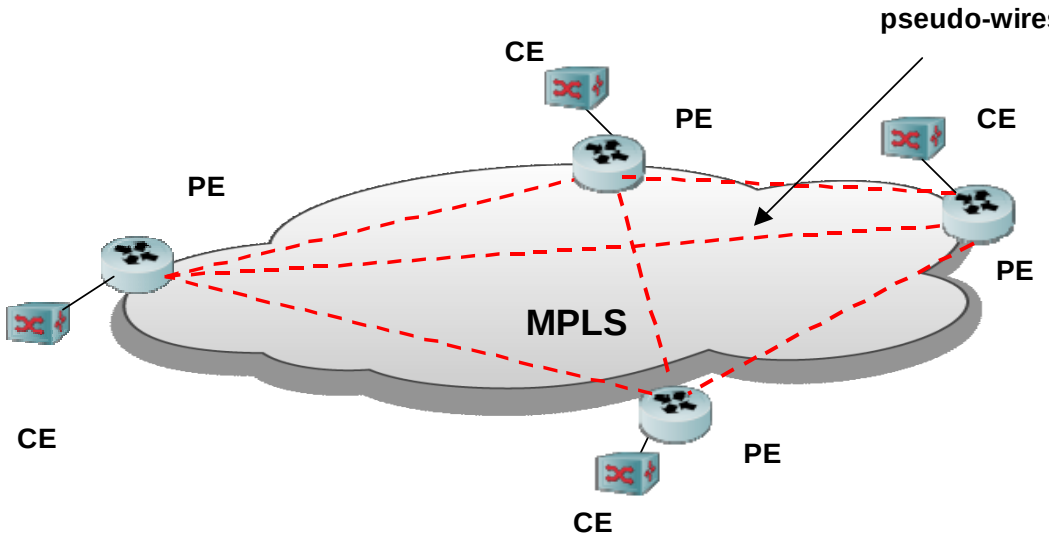

Figure 5.3: A VPLS Service Linking Four Customer Sites

With VPLS, either the Border Gateway Protocol (BGP) or the Label Distribution Protocol (LDP) is used to create the required pseudo-wires to fully mesh the provider edge (PE) devices serving the customer sites. Meshed pseudo-wires support the multipoint-to-multipoint nature of the virtual LAN and improve reliability. Reliability is enhanced because in case of failure in the MPLS network, traffic will automatically be routed along available backup paths, providing very short failover times.

Within VPLS, MPLS packets have a two-label stack. The outer label is used for normal MPLS forwarding in the service provider's network. If BGP is used to establish the VPLS, the inner label is allocated by a PE as part of a label block. If LDP is used, the inner label is a virtual circuit ID assigned by LDP when it first establishes a mesh between the participating PEs. Every PE keeps track of the assigned inner labels, and associates these labels with the VPLS instance.

Table 5.5 provides a high level comparison of the different types of Ethernet WAN provider services available for LAN extension between data centers. It should be noted that there are other options for LAN extension, such as Ethernet over leased dark fiber and Ethernet over GRE tunneling through a private IP network. Cisco's Overlay Transport Virtualization (OTV) falls into the latter category, as a proprietary option.

Service Topology	Access Link	Provider Core	Service Type	Tunneling
Ethernet end-end	Ethernet	Ethernet	Pt-Pt or Mpt-Mpt	802.1Q or Q in Q
Ethernet/IP	Ethernet	IP	Pt-Pt or Mpt-Mpt	L2TPv3
VPLS/VPWS	Ethernet	MPLS	Pt-Pt or Mpt-Mpt	EoMPLS

Table 5.5: Ethernet WAN Service Types

Hybrid WANs/WAN Virtualization

As previously noted, IT organizations won't be able to cost-effectively satisfy the primary WAN design criteria with a simplistic network design. As was also noted, in the current environment it is somewhat common to create a new WAN service by making minor variations to existing WAN services. One example of that approach is the utilization of Policy Based Routing (PBR) to implement a hybrid WAN that leverages multiple WAN services such as MPLS and the Internet. When a router receives a packet it normally decides where to forward it based on the destination address in the packet, which is then used to look up an entry in a routing table. Instead of routing by the destination address, PBR allows network administrators to create routing policies to allow or deny paths based on factors such as the identity of a particular end system, the protocol or the application.

An advantage of the PBR approach to creating a hybrid WAN is that IT organizations already have the functionality to do so. There are, however, some disadvantages of this approach. For example, configuring PBR is complex, time consuming and error prone. There is also not a direct linkage between PBR and other critical WAN functionality, such as the ability to optimize network and application traffic and the ability to have visibility into the traffic that transits the WAN.

Perhaps the biggest limitation of this simple PBR approach it that it creates a static allocation of traffic to multiple links and it doesn't have the ability to reallocate the traffic when the quality of one of the links degrades. The static allocation of traffic causes all one type of traffic (e.g., SCM) to always transit an MPLS service while causing all traffic of another type (e.g., FTP) to always transit the Internet. There will be times, however, when sending some SCM traffic over the Internet will result in acceptable application performance and conserve expensive resources. Since PBR cannot respond in real time to changing network conditions, it will consume more costly bandwidth than would a dynamic approach to creating a hybrid WAN.

> ***In order to be effective, a hybrid WAN has to be able to perform adaptive path selection across two or more WAN links in a dynamic, intelligent fashion.***

In order to perform adaptive path selection a dynamic hybrid WAN must be able to select a WAN link in real time based on:

- The instantaneous end-to-end performance of each available network: This allows the solution to choose the optimal network path for differing traffic types.

- The instantaneous load for each end-to-end path: The load is weighted based on the business criticality of the application flows. This enables the solution to maximize the business value of the information that is transmitted.

- The characteristics of each application: This includes the type of traffic (e.g., real time, file transfer); the performance objectives for delay, jitter and packet loss; as well as the business criticality and information sensitivity.

As noted, one option is for a hybrid WAN to balance traffic over MPLS and the Internet. While this reduces cost vs. an approach that puts all traffic over MPLS, additional cost savings are possible. For example, another option is to balance traffic over multiple low cost Internet access services such as DSL and cable. As previously noted, many IT organizations have avoided utilizing these access options for branch offices because the traditional approach to using DSL or cable for Internet access results in an unacceptably low level of availability and performance predictability.

However, because of adaptive path selection, the availability of a hybrid WAN that consists of multiple parallel paths is very high even if the availability of each component path is only moderately high. For example, Figure 5.4 depicts a system that is composed of two components that are connected in parallel.

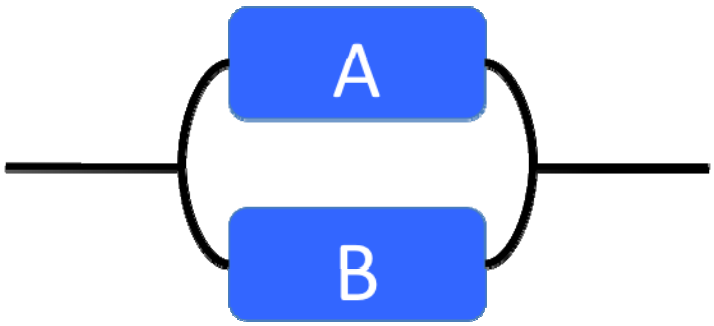

Figure 5.4: Two Components Connected in Parallel

The system depicted in Figure 5.4 is available unless both of the two components are unavailable. Assuming that each component is a diversely routed DSL or cable access line and that one of the access lines has an availability of 99% and the other has an availability of 98%, then the system has an availability of 99.98%. Alternatively, if both access lines have an availability of 99%, then the system is available 99.99% of the time. This level of availability is equal to or exceeds the availability of most MPLS networks.

As described above, one of the principal advantages of a dynamic hybrid WAN is that is allows IT organizations to add significant amounts of additional bandwidth to an existing MPLS-based WAN at a relatively low incremental cost. WAN Virtualization (Figure 5.5) can be thought of as a variation of a hybrid WAN. In addition to enabling the augmentation of an MPLS WAN with inexpensive Internet connectivity, WAN Virtualization also gives IT organizations the option to reduce its monthly ongoing expense by either eliminating or reducing its MPLS connections while simultaneously providing more bandwidth than did the original design.

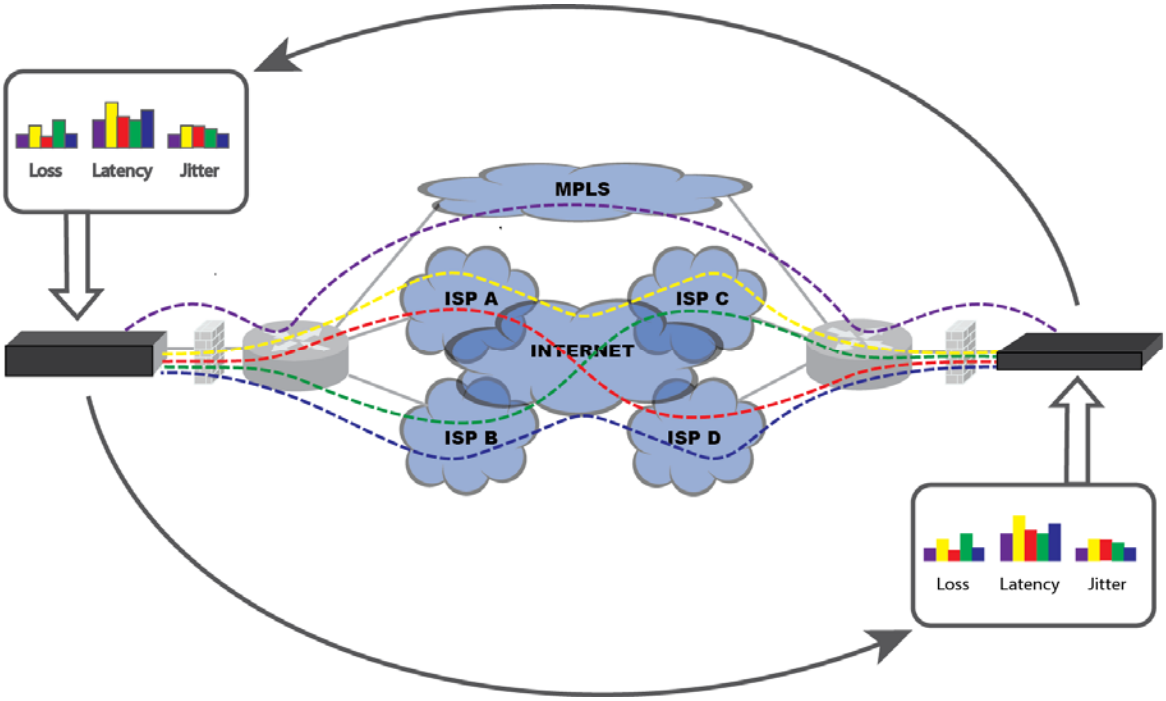

Figure 5.5: WAN Virtualization

As shown in Figure 5.5, because it continuously measures loss, latency, jitter and bandwidth utilization across all of the various paths between any 2 locations, in less than a second WAN Virtualization can switch traffic away from a path that is exhibiting an unacceptable level of performance. This capability, combined with the availability of parallel systems as depicted in Figure 5.4, means that all of the bandwidth in each of the paths can be used most of the time, and that most of the bandwidth can be used all of the time. This combination of capabilities also underscores the ability of WAN Virtualization to deliver performance predictability that equals, and in most cases exceeds, that of a single MPLS network.

Because of the high availability and performance predictability of WAN virtualization, IT organizations can now leverage a number of WAN services that are dramatically lower in cost than traditional MPLS services. This includes DSL and cable Internet access from branch offices and fiber access to the Internet from data centers. It also positions IT organizations to take advantage of the huge volumes of very inexpensive Internet access bandwidth that are typically available at co-location facilities.

While the preceding discussion focused on DSL and cable access to the Internet it is important to realize that over the next year or two, there will be a broad scale deployment of 4G services on the part of most wireless service providers. There will be some variability in the effective bandwidth of 4G services based in part on the fact that the wireless service providers will not all implement the same technologies. It should generally be possible, however, for users of these services to realize throughput in the range of three to four megabits per second, which is roughly equivalent to two T1 or E1 access lines. This will make 4G services a viable access service for some branch offices. For example, a 4G service could be combined with Internet access via DSL as part of a virtual WAN. In addition to providing cost savings, due to the inherent diverse routing associated with 4G and DSL, this design would provide a very high level of reliability.

Cloud Bridging

A hybrid cloud computing solution relies on a WAN to provide the connection between the enterprise locations, including the enterprise data center(s), and the public cloud data center(s) providing the IaaS or other cloud service. Ideally, the resulting hybrid cloud would appear to both users and administrators as a single cloud data center, with the physical location of application resources as transparent as possible.

As is the case for private clouds:

> ***Hybrid clouds depend heavily on VM migration among geographically dispersed servers connected by a WAN in order to ensure high availability and dynamic response to changes in user demand for services.***

In a hybrid cloud that leverages a WAN for access to the public portion of the cloud, transparency of application location has a number of implications:

- **VLAN Extension**
  The VLANs within which VMs are migrated must be extended over the WAN between the private and public data centers. This involves the creation of an overlay network that allows the Layer 2 VLAN traffic to be bridged or tunneled through the WAN.

- **Secure Tunnels**
  These tunnels must provide an adequate level of security for all the required data flows over the Internet. For the highest level of security, this would typically involve both authentication and encryption, such as that provided by IPsec tunnels.

- **Universal Access to Central Services**
  All application services, such as load balancing, DNS, and LDAP, should be available and function transparently throughout the hybrid cloud. This allows these application services to be provisioned from the private enterprise data center and

eliminates the need for manual intervention to modify server configurations as the application and its VM are transferred from the private cloud to the public cloud.

- **Application Performance Optimization**
  Application performance must meet user expectations regardless of user location within the enterprise network and the server location within the hybrid cloud. This means that public cloud data center extensions need to offer the same WAN optimization and application acceleration capabilities as are described below. In addition, WAN optimization controllers serving the bridged connection between the enterprise private cloud data center and the public cloud data center can accelerate VM migration, system backups, and other bulk data transfers between these data centers.

The Evolving Branch Office

As recently as 2005 the typical branch office was IT heavy. By that is meant that the branch office circa 2005 contained a lot of network equipment (e.g., switches and branch office routers) as well as a multitude of servers, applications, and storage. The devices in the typical branch office, however, were not aware of application level traffic detail. In addition, the network optimization of application traffic, if it existed at all, was very primitive (e.g., QoS, TOS) and it was focused on a couple of obvious applications. It was also unusual in this time frame for IT organizations to have visibility into the application traffic that was either originated in, or was destined to branch office users.

The current branch office infrastructure is less IT centric than it was in 2005. For example, the majority of IT organizations have removed at least some servers, applications and storage out of the branch office and placed these resources in centralized data centers. This centralization of IT resources is driven both by the desire to reduce computing and IT management costs as well as the desire to get more control over those resources and to implement better security.

In order to improve the performance of applications delivered to branch office employees, many IT organizations have deployed into their branch offices the WAN optimization controllers (WOCs) that are described below. In addition, in order to improve security, many IT organizations have deployed firewalls, intrusion detection systems (IDS) and intrusion protection systems (IPS) into their branch offices.

> ***One of the design challenges facing IT organizations is how to strike a balance between having branch offices be IT heavy and consolidating all IT functionality out of branch offices and into centralized data centers.***

To help IT organizations strike that balance, many vendors have developed a branch office box (BOB). The advantage of the BOB is that it can potentially consolidate several functions (i.e., WOC, firewall, IDS, IPS, print server, file server, domain controller, etc.) into a single physical device, while also allowing all of these capabilities to be managed via a single system management interface.

The Webtorials Respondents were asked to indicate the type of device that they would most likely build their branch office infrastructure around. Their responses are shown in Table 5.6.

Type of Device	Percentage of Respondents
A router that supports VMs	30.7%
An appliance, such as a WOC, that supports VMs	25.7%
A virtualized server	27.7%
Other	15.8%

Table 5.6: BOB Form Factor

In the *other* category, the most common response was that based on the situation, all three types of devices are provided.

Local Access to the Internet

The traditional approach to providing Internet access to branch office employees was to carry the Internet traffic on the organization's enterprise network (e.g., their MPLS network) to a central site where the traffic was handed off to the Internet. The advantage of this approach is that it enables IT organizations to exert more control over the Internet traffic and simplifies management in part because it centralizes the complexity of implementing and managing security policy. One disadvantage of this approach is that it results in extra traffic transiting the enterprise WAN, which adds to the cost of the WAN. Another disadvantage of this approach is that it adds additional delay to the Internet traffic.

The Webtorials Respondents were asked to indicate how they currently route their Internet traffic and how that is likely to change over the next year. Their responses are contained in Table 5.7.

Percentage of Internet Traffic	Currently Routed to a Central Site	Will be Routed to a Central Site within a Year
100%	39.7%	30.6%
76% to 99%	24.1%	25.4%
51% to 75%	8.5%	13.4%
26% to 50%	14.2%	14.2%
1% to 25%	7.1%	6.7%
0%	6.4%	9.7%

Table 5.7: Routing of Internet Traffic

Driven in part to save money and in part to improve application performance:

> ***IT organizations will make an increased use of distributed access to the Internet from their branch offices.***

Network and Application Optimization

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

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

Table 5.8 shows the ten optimization tasks that are the most important for IT organizations to improve on in the next year. Included in Table 5.8 are the tasks and the percentage of The Webtorials Survey Respondents who indicated that the task was either very or extremely important for their organization to get better at over the next year.

Optimization Tasks	Importance: Very or Extremely
Relating the performance of applications to their impact on the business	70%
Ensuring acceptable performance for VoIP traffic	68%
Improving the performance of applications used by mobile workers	60%
Ensuring acceptable performance for video or telepresence traffic	57%
Ensuring acceptable performance for the applications that you acquire from a Software-as-a-Service (SaaS) provider	56%
Optimizing the performance of TCP	54%
Controlling the cost of the WAN by reducing the amount of traffic that transits the WAN	50%
Optimizing the Web tier of a multi-tiered application for peak utilization	50%
Optimizing the performance of specific applications such as SharePoint	49%
Optimizing the performance of protocols other than TCP; e.g., HTTP and MAPI	49%

Table 5.8: The Importance of Improving Optimization Tasks

The data in Table 5.8 shows that even though acquiring applications from a SaaS provider is a relatively recent phenomenon, more than half of The Webtorials Respondents stated that ensuring acceptable performance for the applications that they acquire from a SaaS provider is either very or extremely important. More detail on the optimization challenges facing IT organizations can be found in the report ***Application Delivery: A Reality Check***.

WAN Optimization Controllers (WOCs)

Goals of a WOC

The goal of a WOC is to improve the performance of applications delivered across the WAN from the data center either to the branch office or directly to the end user, typically over a network such as MPLS. A model is helpful to illustrate how the performance of a WAN can impact the performance of an application and also serves to illustrate how a WOC can improve application performance. The following model (Figure 5.6) is a variation of the application response time model created by Sevcik and Wetzel[11]. Like all mathematical models, the following is only an approximation, and as a result it is not intended to provide results that are accurate to the millisecond level.

As shown below, the application response time (R) is impacted by amount of data being transmitted (Payload), the WAN bandwidth, the network round trip time (RTT), the number of application turns (AppTurns), the number of simultaneous TCP sessions (concurrent requests), the server side delay (Cs) and the client side delay (Cc).

$$R \approx \frac{Payload}{Goodput} + \frac{(\#\ of\ AppsTurns\ *\ RTT)}{\text{Concurrent Requests}} + Cs + Cc$$

Figure 5.6: Application Response Time Model

In order to improve the performance of applications that are delivered over the WAN, WOCs implement a variety of techniques. These techniques are summarized in Table 5.9 and are explained in detail in ***The 2010 Application Delivery Handbook***.

[11] Why SAP Performance Needs Help, NetForecast Report 5084, http://www.netforecast.com/ReportsFrameset.htm

WAN Characteristics	WAN Optimization Techniques
Insufficient Bandwidth	Data Reduction: • Data Compression • Differencing (a.k.a., de-duplication) • Caching
High Latency	Protocol Acceleration: • TCP • HTTP • CIFS • ICA • RDP Application Acceleration • SharePoint • SAP • Oracle Mitigate Round-trip Time • Request Prediction • Response Spoofing
Packet Loss	Congestion Control Forward Error Correction (FEC) Packet Reordering
Network Contention	Quality of Service (QoS)

Table 5.9: Techniques to Improve Application Performance

Virtually all WAN Optimization Controllers (WOCs) on the market support the functions listed above. However, as described below, there are some significant differences in terms of how the functionality is implemented and how well it performs. In addition, some WOC vendors provide functionality not included in the above list. A recent report[12] provides insight into the primary WOC vendors and their products.

Enabling Virtual Desktops

As was previously mentioned, one of the factors that will drive more traffic over the WAN is the implementation of virtual desktops. As explained in the report entitled ***Virtualization: Benefits, Challenges and Solutions***, the two fundamental forms of desktop virtualization are client side (a.k.a., streamed desktops) and server side; a.k.a., hosted desktops.

> ***The ICA and RDP protocols employed by many hosted desktop virtualization solutions are examples of protocols that can be difficult to optimize.***

One of the reasons that these protocols can be difficult to optimize is that they only send small request-reply packets. Byte-level caching best optimizes this form of

[12] http://searchenterprisewan.techtarget.com/generic/0,295582,sid200_gci1381156,00.html

communications. Unfortunately, not all WOCs support byte-level caching. Implementers of desktop virtualization need to understand the functionality provided by the various WOCs and to evaluate that functionality in the context of the types of desktop virtualization that they want to deploy.

As shown in Table 5.10, techniques such as byte level compression, caching, protocol (e.g., ICA, RDP) optimization, and QoS can provide benefits for hosted desktops. Before implementing them, however, an IT organization must determine which acceleration techniques are compatible with the relevant display protocols. For example, in order to be able to compress ICA traffic, a WOC must be able to decrypt the ICA workload, apply the optimization technique, and then re-encrypt the data stream.

	Streamed Desktops	Hosted Desktops
Block Level Compression	X	
Byte Level Compression	X	X
Caching	X	X
Staging	X	
Protocol Optimization (e.g., TCP, IP, UDP)	X	X
Protocol Optimization (e.g., ICA, RDP)		X
Protocol Optimization (e.g., CIFS, HTTP, MAPI)	X	
QoS	X	X

Table 5.10: Applicability of Common WAN Optimization Techniques

The support of streamed desktops also creates some significant WAN performance problems that may require the deployment of a WOC. For example, the code for streamed applications is typically transferred via a distributed file system protocol, such as CIFS, which is well known to be a chatty protocol. Hence, in order to effectively support application streaming, IT organizations need to be able to optimize the performance of protocols such as CIFS, MAPI, HTTP, and TCP. In addition, IT organizations need to implement other techniques that reduce the bandwidth requirements of application streaming. For example, by using a WOC it is possible to cache the virtual application code at the client's site. Caching greatly reduces the volume of traffic for client-side virtualized applications and it also allows applications to be run locally in the event of network outages. Staging is a technique that is similar to caching but is based on pre-positioning and storing streamed applications at the branch office on the WOC or on a branch server. With staging, the application is already locally available at the branch when users arrive for work and begin to access their virtualized applications.

Whether it is done by the WOC itself, or in conjunction with the WOC, supporting desktop virtualization will require that IT organizations are able to apply the right mix of optimization technologies for each situation. For example, protocols such as ICA and RDP already incorporate a number of compression techniques. As a result, any compression performed by a WAN optimization appliance must adaptively orchestrate with the hosted virtualization infrastructure to prevent compressing the traffic twice - a condition that can actually increase the size of the compressed payload.

Mobile Workers

In the last few years there has been an explosive growth in the number of mobile workers.

In order to enable this growing population of mobile workers to access enterprise applications as easily as do workers in branch offices, the communications between the mobile worker and the data center (whether it is owned by the enterprise or a third party provider such as a cloud computing service provider) has to be optimized. The importance of optimizing this traffic is reflected in the data in Table 5.8. As shown in that table, sixty percent of The Webtorials Respondents stated that improving the performance of applications used by mobile workers is either very or extremely important.

One way to optimize this communications is to deploy client software that provides WOC functionality onto the user's mobile device. In many cases the mobile worker will use some form of wireless access. Since wireless access tends to exhibit more packet loss than does wired access:

> ***The WOC software that gets deployed to support mobile workers needs functionality such as forward error correction that can overcome the impact of packet loss.***

In addition, as workers move in and out of a branch office, it will be necessary for a seamless handoff between the mobile client and the branch office WOC.

Until recently, the typical device that mobile workers used to access enterprise applications was a laptop. While that is still the most common scenario, today many mobile workers use their smartphones to access enterprise applications. Therefore, over the next few years it is reasonable to expect that many IT organizations will support the use of smartphones as an access device by implementing server-side application virtualization for those devices. This means that in a manner somewhat similar to remote workers, mobile workers will access corporate applications by running protocols such as ICA, RDP and PCoIP over a WAN.

Many IT organizations, however, resist putting any more software on the user's device. In addition, many users resent having multiple clients (e.g., WOC, SSL VPN, IPSec VPN, wireless/cellular access) that are not integrated on their access device. On a going forward basis, IT organizations should look to implement WOC software that is integrated with the other clients used by mobile workers.

Application Delivery Controllers (ADCs)

The current generation of ADCs evolved from the earlier generations of Server Load Balancers (SLBs) that were deployed in front of server farms. While an ADC still functions as an SLB, the ADC has assumed, and will most likely continue to assume, a wider range of sophisticated roles that enhance server efficiency and provide asymmetrical functionality to accelerate the delivery of applications from the data center to individual remote users.

An ADC provides more sophisticated functionality than a SLB does.

Referring back to Figure 5.6, one of the factors that increase the application response time is server side delay. An ADC can reduce server side delay and hence can reduce the application response time. In particular, the ADC can allow a number of compute-intensive functions, such as SSL processing and TCP session processing, to be offloaded from the server. Server offload can increase the transaction capacity of each server, reducing the number of servers required for a given level of business activity.

The deployment of an SLB enables an IT organization to get a *linear benefit* out of its servers. That means that if an IT organization that has implemented an SLB doubles the number of servers supported by that SLB that it should be able to roughly double the number of transactions that it supports. The traffic at most Web sites, however, is not growing at a linear rate, but at an exponential rate. To exemplify the type of problem this creates, assume that the traffic at a hypothetical company's (Acme) Web site doubles every year. If Acme's IT organization has deployed a linear solution, such as an SLB, after three years it will have to deploy eight times as many servers as it originally had in order to support the increased traffic. However, if Acme's IT organization were to deploy an effective ADC then after three years it would still have to increase the number of servers it supports, but only by a factor of two or three – not a factor of eight. The phrase *effective ADC* refers to the ability of an ADC to have all features turned on and still support the peak traffic load.

Key ADC Functionality

Below is a listing of the type of functionality that is provided by an ADC. This functionality is explained in detail in ***The 2010 Application Delivery Handbook***.

- **TCP Offload**
  TCP offload functionality is designed to deal with the complexity associated with the fact that each object on a Web page requires its own short-lived TCP connection. Processing all of these connections can consume an inordinate about of the server's CPU resources. Acting as a proxy, the ADC terminates the client-side TCP sessions and multiplexes numerous short-lived network sessions initiated as client-side object requests into a single longer-lived session between the ADC and the Web servers.

- **SLB and Global SLB**
  As noted, an ADC sits in front of a server farm and receives service requests from clients and delivers the requests for service to the most appropriate servers. As such, an ADC functions as a traditional SLB. In addition, an ADC can function as a global server load balancer (GSLB). In this role the ADC balances the load across geographically dispersed data centers by sending a service request to the data center that is exhibiting the best performance metrics.

- **SSL Offload**
  The ADC terminates the SSL session by assuming the role of an SSL Proxy for the servers. SSL offload can provide a significant increase in the performance of secure intranet or Internet Web sites. SSL offload frees up server resources, allowing existing servers to process more requests for content and handle more transactions.

- **XML Offload**
  XML is a verbose protocol that is CPU-intensive. Hence, another function that can be provided by the ADC is to offload XML processing from the servers by having an ADC serve as an XML gateway.

- **Scripting**
  One of the characteristics of most IT environments is that the environment is comprised of a large and growing number of servers and applications. Another characteristic is that most IT organizations have very limited control as to which users access which applications and which servers. An ADC gives control to the IT organization through functionality sometimes referred to as scripting, and sometimes referred to as a rules engine. This functionality allows the IT organization to directly classify and modify the traffic of any IP-based application.

- **Application Firewalls**
  ADCs may also provide an additional layer of security for Web applications by incorporating application firewall functionality. Application Firewalls are focused on blocking increasingly prevalent application-level attacks. Application firewalls are typically based on Deep Packet Inspection (DPI), coupled with session awareness and behavioral models of normal application interchange.

Virtual Appliances

Section 4 of this report used the phrase *virtual switch* in two fundamentally different ways. One way referred to making two or more physical switches appear to be a single logical switch. The other way referred to the switching functionality that resides inside of a virtualized server.

In similar fashion, it is possible to look at a *virtual appliance* in two fundamentally different ways. For example, two or more appliances, such as ADCs, can be combined to appear as a single logical ADC. However, the more common use of the phrase *Virtual Appliance* refers to what is typically appliance-based software, together with its

operating system, running in a VM. Virtual appliances can include WOCs, ADCs, firewalls, routers, IDS, IPS and performance monitoring solutions among others. As previously discussed, in the branch office a suitably placed virtualized server could potentially host a virtual WOC appliance as well as other virtual appliances. Alternatively, a router or a WOC that supports VMs could also serve as the infrastructure foundation of the branch office.

One of the compelling advantages of a virtualized appliance is that the acquisition cost of a software-based appliance can be notably less than the cost of a hardware-based appliance with same functionality.

In many cases the cost of a software-based appliance can be a third less than the cost of a hardware-based appliance[13]. In addition, a software-based solution can potentially leverage the functionality provided by the hypervisor management system to provide a highly available system without having to pay for a second appliance[14].

In addition to cost savings, another advantage of a virtual appliance is that it offers the potential to alleviate some of the management burdens in branch offices because most of the provisioning, software updates, configuration, and other management tasks can be automated and centralized at the data center.

In many instances the benefits of the dynamic movement of a VM from one server to another are maximized if the supporting infrastructure can also be dynamically moved.

If virtualized appliances have been deployed, then it is notably easier than it is in a more traditional environment for various networking functions (WOC, ADC, firewall, router, etc.) to be migrated along with VMs in order to replicate the VMs's networking environment in its new location.

A virtualized ADC makes it easy for an IT organization to package and deploy a complete application. One example of this packaging is the situation in which an entire application resides on VMs inside a physical server. The virtualized ADC that supports the application resides in the same physical server and it has been tuned for the particular application. This makes it easy to replicate or migrate that application as needed. In this case, a virtualized ADC also provides some organizational flexibility. For example, the ADC might be under the control of a central IT group or it might be under the control of the group that supports that particular application. The later is a possibility because any actions taken by the application group relative to the ADC will only impact their application.

[13] The actual price difference between a hardware-based appliance and a software-based appliance will differ by vendor.
[14] This statement makes a number of assumptions, including the assumption that the vendor does not charge for the backup software-based appliance.

One approach to monitoring and troubleshooting inter-VM traffic is to deploy a virtual performance management appliance or probe (vProbe).

One of the characteristics of a virtualized server is that each virtual machine only has at its disposal a fraction of the resources (i.e., CPU, memory, storage) of the physical server on which it resides. As a result, in order to be effective, a vProbe must not consume significant resources. The way that a vProbe works is similar to how many IT organizations monitor a physical switch. In particular, the vSwitch has one of its ports provisioned to be in promiscuous mode and hence forwards all inter-VM traffic to the vProbe. As a result, the use of a vProbe gives the IT organization the necessary visibility into the inter-VM traffic.

A virtual firewall appliance can help IT organizations meet some of the challenges associated with server virtualization. That follows because virtual firewall appliances can be leveraged to provide isolation between VMs on separate physical servers as well as between VMs running on the same physical server. The security appliance can potentially provide highly integrated functionality to help secure virtual machines, applications, and traffic. This includes firewall, VPN, anti-malware, IDS/IPS, integrity monitoring (e.g., registry changes), and log inspection functionality.

Virtualized security management makes it is possible to meet critical regulatory compliance requirements for full application segregation and protection within the confines of virtualized physical servers. Through tight integration with the virtual server management system, firewall appliances can also be dynamically migrated in conjunction with VM migration where this is necessary to extend a trust zone to a new physical location. In addition, hypervisor APIs, such as VMware's Vsafe, can allow physical/virtual firewall consoles to monitor servers for abnormal CPU, memory, or disk activity without the installation of special agent software.

The recently published report entitled ***Virtualization: Benefits, Challenges and Solutions***, contains more detail on virtual appliances. Included in that report is a discussion of the challenges associated with virtual appliances, as well as suggested evaluation criteria.

Software as a Service (SaaS) Based Solutions

As previously noted, one of the primary forms of public cloud computing is SaaS. As was also noted, the phrase ***cloud networking*** refers to the LAN, WAN and management functionality that must be in place to support cloud computing. As described below, there is a distinct synergy in which one of the approaches that typifies public cloud computing, SaaS, can be used as part of a cloud network to better support cloud computing. In particular, IT organizations can acquire the optimization and security functionality they need to support cloud computing from a SaaS provider. Currently, the most developed form of a SaaS solution that provides optimization and security functionality is an Internet overlay.

An Internet Overlay

As previously described, IT organizations often implement WAN Optimization Controllers (WOCs) and Application Delivery Controllers (ADCs) to improve network and application performance. However, these solutions make the assumption that performance characteristics within the WAN itself can't be optimized because they are determined by the relatively static service parameters controlled by the WAN service provider. This assumption is reasonable in the case of WAN services such as MPLS. However, this assumption does not apply to enterprise application traffic that transits the Internet because there are significant opportunities to optimize performance within the Internet itself based on implementing an Internet overlay. An Internet overlay leverages service provider resources that are distributed throughout the Internet in order to optimize the performance, security, reliability, and visibility of the enterprise's Internet traffic. As shown in Figure 5.7, all client requests to the application's origin server in the data center are redirected via DNS to a server in a nearby point of presence (PoP) close to application users, typically within a single network hop. This edge server then optimizes the traffic flow to the server closest to the data center's origin server.

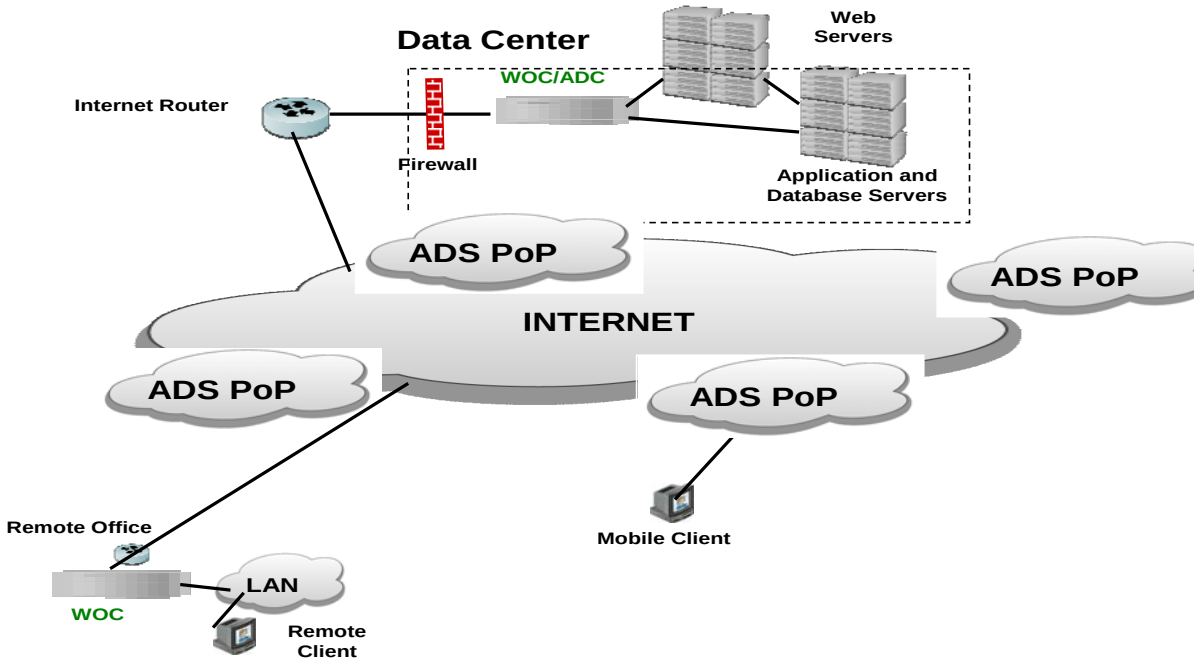

Figure 5.7: An Internet Overlay

An Internet overlay provides a variety of optimization functions that generally complement solutions such as an ADC rather than overlap or compete with them. One such function is content offload. This calls for taking static content out of a data-center and placing it in caches in servers and in replicated in-cloud storage facilities. IT organizations that offload content and storage improve response time and simultaneously reduce both their server utilization as well as the bandwidth utilization of their data center access links.

Some of the other common functionality associated with an Internet overlay include:

- Route optimization
- Transport optimization
- HTTP protocol optimization
- Visibility

In addition to the functionality listed above, some Internet overlays incorporate Web application firewall functionality.

6. Management

Background

As pointed out in ***Virtualization: Benefits, Challenges and Solutions***, server virtualization creates a number of management challenges. For example, the need to manually reconfigure the network to support VM migration that was previously mentioned can be regarded as either a LAN challenge or a management challenge. Additional management challenges include:

- **Breakdown of Network Design and Management Tools**
  The workload for the operational staff can spiral out of control due to the constant stream of configuration changes that must be made to the static date center network devices in order to support the dynamic provisioning and movement of VMs.

- **Limited VM-to-VM Traffic Visibility**
  The first generation of vSwitches doesn't have the same traffic monitoring features as does physical access switches. This limits the IT organization's ability to do security filtering, performance monitoring and troubleshooting within virtualized server domains.

- **Poor Management Scalability**
  Many IT organizations have experienced VM proliferation sometimes called VM sprawl. In addition, the normal best practices for virtual server configuration call for creating separate VLANs for the different types of traffic to and from the VMs. The combined proliferation of VMs, and VLANs places a significant strain on the manual processes traditionally used to manage servers and the supporting infrastructure.

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

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

The Webtorials Respondents were asked to indicate how important it is to their organization to get better over the next year at managing some key tasks related to server virtualization. They were given the following five-point scale:

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

Included in Table 6.1 are the tasks and the percentage of The Webtorials Survey Respondents who indicated that the task was either very or extremely important for their organization to get better at over the next year.

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

Table 6.1: Importance of Managing Server Virtualization

As shown in Table 6.1, The Webtorials Survey Respondents indicated that getting better at each of the individual challenges associated with server virtualization is important to their organization. In addition, it is reasonable to look at the five challenges contained in Table 6.1 as being a single challenge - managing server virtualization. When looked at that way, getting better at server virtualization is extremely important to The Webtorials Survey Respondents.

The Evolving Management Environment

One of the primary management challenges associated with cloud networking is the movement to focus on managing services as defined in section two of this report. To put that challenge into perspective, The Webtorials Survey Respondents were asked to indicate the approach their organization takes to management. They were given the following choices and allowed to choose all that applied to their environment.

- We have a focus primarily on individual technology domains such as LAN, WAN and servers

- We have a focus on managing the performance of applications as seen by the end user
- We have a focus on managing the performance of services as seen by the end user, where service refers to multiple, inter-related applications
- Other

Their responses are shown in Figure 6.1.

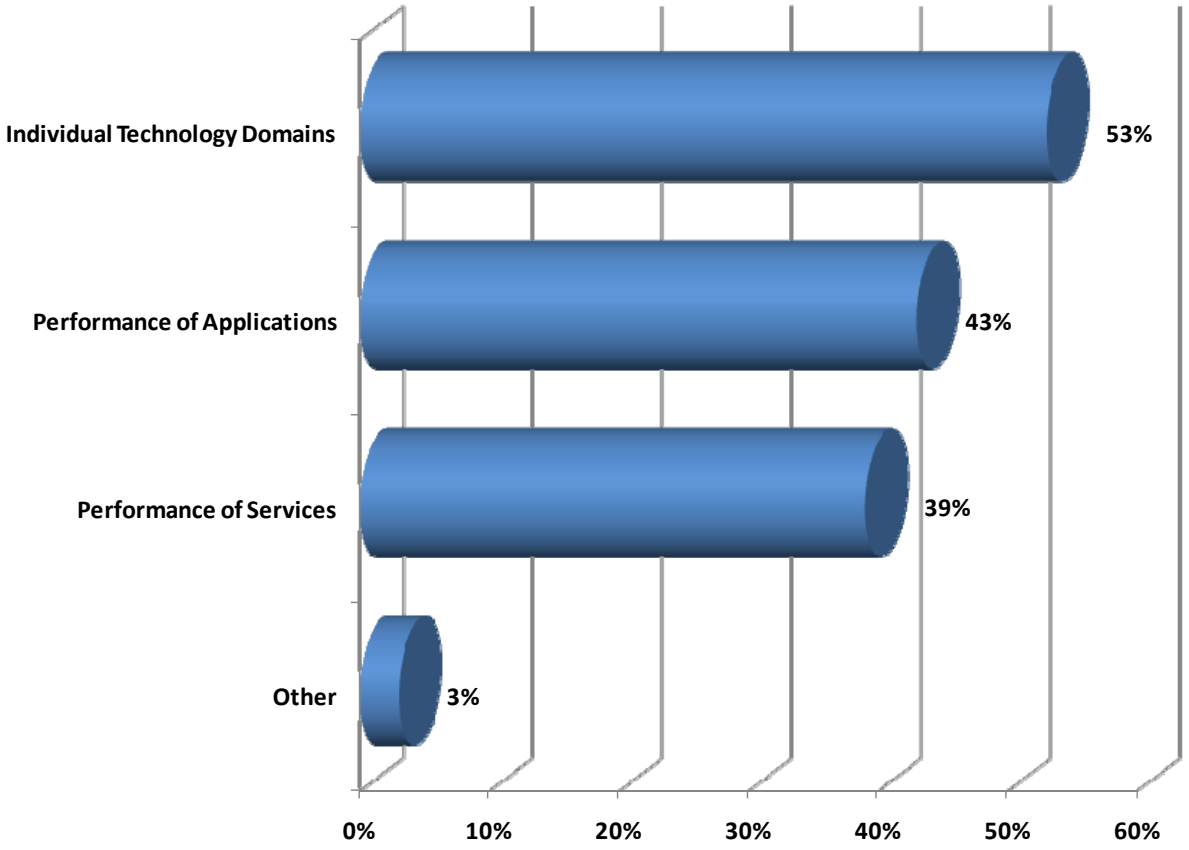

Figure 6.1: Focus of Management

The data in Figure 6.1 indicates that the most frequent approach that IT organizations take to management is to focus on individual technology domains. However:

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

The Webtorials Survey Respondents were also asked to indicate how important it is to their organization to get better at twenty different management tasks over the next year. They were given the same five-point scale as was discussed relative to Table 6.1. Included in Table 6.2 are the tasks and the percentage of The Webtorials Survey Respondents who indicated that the task was either very or extremely important for their organization to get better at over the next year.

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

Table 6.2: The Importance of Improving Management Tasks

More detail on the management challenges facing IT organizations can be found in the report ***Application Delivery: A Reality Check***.

The fact that sixty one percent of The Webtorials Respondents indicated that effectively managing services, where services are comprised of multiple, inter-related applications was either very or extremely important to them underscores the growing importance of IT organizations having a focus on services.

Communications Based Applications

The fact that two thirds of The Webtorials Respondents indicated that effectively managing QoS was either very or extremely important to them is consistent with some of the results contained in the previous section of this report. In particular, the data in Table 5.8 indicates that two thirds of The Webtorials Respondents indicated that ensuring acceptable performance for VoIP traffic was either very or extremely important to them. In addition, over half of The Webtorials Respondents indicated that ensuring acceptable performance for video or telepresence traffic was either very or extremely important to them. QoS is an important technique to allow IT organizations to ensure the performance of communications based applications such as VoIP, video and telepresence.

The survey data discussed in the preceding paragraph clearly indicates the importance of managing communications based applications such as VoIP, traditional video conferencing, as well as high definition video conferencing and telepresence. This importance will only grow as the use of those services grows. To illustrate the growth in communications based applications, Table 6.3 shows the percentage of The Webtorials Respondents that currently make at least some use the application and the percentage that expect to increase their use of the application over the next year.

Application	Percent Currently Using	Percent that Expect to Increase Use
Traditional voice (POTS)	93.0%	18.2%
VoIP	91.6%	76.9%
Traditional videoconferencing	86.2%	36.5%
HD Videoconferencing or telepresence	60.1%	58.3%

Table 6.3: Use of Communications Based Applications

Internal SLAs

As recently as two years ago, few IT organizations offered an SLA to the company's business and functional managers; a.k.a., an internal SLA. However, that situation has changed and now it is common for IT organizations to offer internal SLAs. To understand the prevalence and effectiveness of internal SLAs, The Webtorials Respondents were asked to indicate their agreement or disagreement with three statements. The three statements and the percentage of The Webtorials Respondents that agreed with the statement are shown in Table 6.4.

Statement	Percent that Agree
We provide an SLA internally for every application that we support	30.0%
We provide an SLA internally for at least some applications	69.9%
We do a good job of managing our internal SLAs	55.8%

Table 6.4: Status of Internal SLAs

The data in Table 6.4 highlights the growing interest that IT organizations have in providing internal SLAs for at least some applications. However, as previously noted, the SLAs that are associated with public cloud computing services such as Salesforce.com or Amazon's Simple Storage System are generally weak or non-existent.

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

Root Cause Analysis

It is not surprising that rapidly identifying the root cause of degraded application performance is so important to IT organizations in part because on an ever increasing basis a company's key business processes rely on a handful of business critical applications. That means that if those applications are not running well, neither are those key business processes.
Even in the traditional IT environment[15] when the performance of an application is degrading the degradation is typically noticed first by the end user and not by the IT organization. In addition, when IT is made aware of the fact that application performance has degraded, the process to identify the source of the degradation can be lengthy.

Unfortunately:

The adoption of cloud computing makes troubleshooting application performance an order of magnitude more difficult.

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

One of the challenges associated with managing a private cloud environment was referred to in section 2 of this report. In particular, 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 tasks such as troubleshooting. 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.

Another one of the challenges associated with managing a private cloud environment is that in order to manage BusApp in the type of virtualized environment described above, an IT organization needs detailed information on each of the three VMs that support the application and the communications amongst them. For the sake of example, assume

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

that the IT organization has deployed the tools and processes to gather this information and has been able to determine that the reason that BusApp sporadically exhibits poor performance is that the application server occasionally exhibits poor performance. However, just determining that it is the application server that is causing the application to perform badly is not enough. The IT organization also needs to understand why the application server is experiencing sporadic performance problems. The answer to that question might be that other VMs on the same physical server as the application server are sporadically consuming resources needed by the application server and that as a result, the application server occasionally performs poorly. Part of the challenge associated with troubleshooting this scenario is that as previously noted, in most cases once an IT organization has virtualized its servers it loses insight into the inter-VM traffic that occurs within a physical server.

Staying with this example, now assume that SmartCompany has decided to evaluate the viability of deploying BusApp using either a public or hybrid cloud computing solution.

One of the fundamental issues relative to managing either a public or hybrid cloud computing service is that the network topology becomes even more complex and hence understanding the end-to-end path becomes more difficult.

For example, as described below, in some hybrid cloud environments the client request goes to a web server hosted by a cloud computing service provider that then queries the internal database over an MPLS network. Part to the complexity here is a result of the fact that:

- For added reliability, there are likely to be multiple BGP-based Internet peering points
- MPLS is difficult to manage using most traditional management tools
- All of the tiers of the application (i.e., Web, application, database) are running on VMs that are being dynamically moved between servers, both within and between data centers.

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

For the sake of this example, consider two alternative approaches that SmartCompany might evaluate. Those approaches are:

1. Public Cloud Computing

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

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

2. Hybrid Cloud Computing

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

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

Management Solutions

Route Analytics

In a traditional IT environment it is sometimes difficult to know the end-to-end path that packets take across a network. This management complexity comes in part from the distributed nature of IP. In particular, routers exchange reachability information with each other via a routing protocol such as OSPF (Open Shortest Path First). Based on this information, each router makes its own decision about how to forward a packet. There is, however, no single repository of routing information in the network. As described in the preceding section, the difficulty of understanding the end-to-end path is magnified in a cloud network.

As shown in Figure 6.2, route analytics provides IT organizations and service providers with insight into the routing layer.

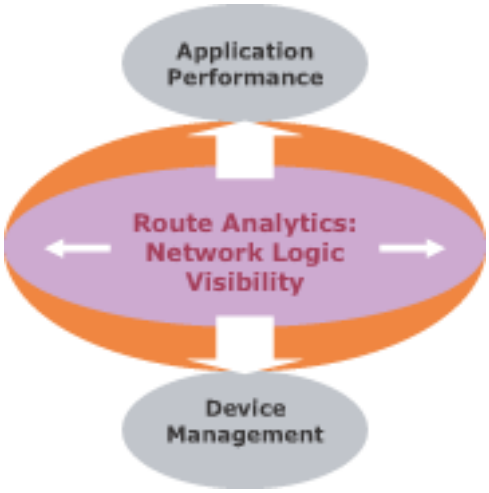

Figure 6.2: The Positioning of Route Analytics

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

Dynamic Infrastructure Management

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

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

Virtualized Performance and Fault Management

Another example of a management capability in the traditional physical environment that is important to implement in a virtual environment is adaptive performance thresholding. This capability identifies systemic deviations from normal as well as time over threshold violations, and can automatically update thresholds based on changes to historic levels of utilization. That same capability is needed in a virtualized environment so that IT organizations can monitor the performance of individual VMs.

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

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

Orchestration and Provisioning

> ***Service orchestration is an operational technique that helps IT organizations automate many of the manual tasks that are involved in provisioning and controlling the capacity of dynamic virtualized services.***

By automatically coordinating provisioning and resource reuse across servers, storage, and networks, service orchestration can help IT organizations streamline operational workloads and overcome technology and organizational silos and boundaries. Orchestration engines use business policies to define a virtual service and to translate that service into the required physical and virtual resources that are needed for deployment. The orchestration engine then disseminates the needed configuration commands to the appropriate devices across the network in order to initiate the requested service. The orchestration engine can automatically initiate the creation of the required virtual machines while simultaneously deploying the network access and security models across all of the required infrastructure components. This includes routers, switches, security devices, and core infrastructure services. The entire process can allow setup and deployment of network routes, VPNs, VLANs, ACLs, security certificates, firewall rules and DNS entries without any time consuming manual entries via device-specific management systems or CLIs.

Orchestration engines are generally limited in the range of devices with which they can interface due to differences in device and/or vendor management interfaces. Therefore, orchestration solutions mirror to some extent the constraints of virtual data center solutions that result from vendor partnerships among manufacturers of virtual server software, networks, and networked storage. The initial focus of such partnerships has been on promulgating validated network designs and architectures rather than on fully integrated or automated management. The next logical step for such partnerships is to include orchestration capabilities.

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

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

7. Summary & Call to Action

For the foreseeable future IT organizations will increasingly adopt cloud computing. Cloud networking is the LAN, WAN and management functionality that enables IT organizations to support cloud computing. The key characteristics of a cloud network are that a cloud network:

- Has the same goal as cloud computing
- Supports the characteristics of a cloud computing solution
- Does no harm to cloud computing solutions
- Provides solutions that are good enough

Data Center LANs

The majority of IT organizations have the goal of evolving their data center infrastructure to be one that can dynamically provide each application and network service with the required resources. Because of the complexity and risk associated with achieving that goal, many IT organizations will choose to implement a Greenfield data center and then gracefully cut over to a new data center.

IT organizations that are going to implement a Greenfield data center should:

- Give primary consideration to two-tier designs that avoid the spanning tree protocol and which are based on switch virtualization and MC LAG.

- Consider new servers with multi-core processors that are capable of supporting a large number of VMs, and which incorporate dual 10 GbE LAN connections. For example, with four-processor, 48 core servers, a realistic goal for the number of VMs per server would fall in the range of ten to fifty, depending on the characteristics of the applications.

- Base the data center LAN design on both high-density modular access and core switches that can provide non-blocking support for 40 and 100 GbE when available. If TOR switches are used due to cost or cabling considerations, IT organizations should make sure that the switches can provide 10 GbE server connections and can support the desired over-subscription ratios.

- Acquire switches from a vendor whose product roadmap includes TRILL/SPB as this allows for possible modification of the switch topology.

- Acquire switches from a vendor whose product roadmap includes some form of early EVB/VEPA support during 2011 or early 2012.

In addition, if storage access (i.e., NAS, SAN) over Ethernet is an important consideration, IT organizations should make sure that all the switches have a solid roadmap for supporting DCB. DCB will also benefit other applications including real-time video and voice.

IT organizations need to adopt a storage networking strategy. Three possibilities are:

- Continue to deploy Fibre Channel.
- Cap the use of Fibre Channel and implement FCoE. This allows IT organizations to both preserve the investment in Fibre Channel and to migrate to a unified fabric.
- Cap the use of Fibre Channel and deploy 10/40/100 GbE iSCSI in large part because this is a simpler approach to fabric unification.

If IT organizations choose to implement FCoE, they need to determine whether they want to connect their Fibre Channel SANs to the access tier or to the core tier of the network. The trade-off between these approaches is primarily between the benefits of sharing storage across the data center vs. the complexity and the impact of loading the uplinks and the core switches with a mixture of storage and traffic.

Wide Area Networking

Unlike the way things were during the entire twenty-year period that began in the mid to late 1980s, today there is not a fundamentally new generation of WAN technology in development. As such, IT organizations need to maximize their use of the existing WAN technologies and leverage whatever new products and services are developed. With this in mind, IT organizations should:

- Consider the use of VPLS. As is typically the case with WAN services, the viability of using this service vs. traditional services will hinge largely on the relative cost of the services. This will vary by service provider and by geography.
- Implement a dynamic virtual WAN/network virtualization. One of the many positive aspects of this approach to wide area networking is that it can be based on a number of different WAN services; e.g., MPLS as well as cable, DSL, T1/E1 and 4G access to the Internet. In addition, this approach can be deployed at a limited number of sites and expanded incrementally if desired.
- Explore cloud bridging solutions if hybrid cloud solutions are of interest.
- Consider local access to the Internet. The trade-off here is that centralized access to the Internet reduces the complexity of providing security, but this approach increases the amount of traffic on the WAN, and hence the cost of the WAN. This approach also adds to overall network delay.

- Implement WOC functionality both to save on the cost of WAN services and to improve application performance. Alternatives include both hardware and software-based WOCs, whether they are provided by the IT organization itself or by a third party as part of a managed service. Another alternative is to acquire WOC functionality from a SaaS provider.

- Implement ADC functionality both to improve the performance of the data center servers as well as to improve overall application performance. Alternatives include both hardware and software-based ADCs, whether they are provided by the IT organization itself or by a third party as part of a managed service.

Management

Almost every aspect of cloud computing (e.g., server virtualization, public and private cloud computing solutions) create significant management challenges. To respond to these challenges, IT organizations should:

- Analyze solutions for data center management automation and integration initiatives from the perspective of the organizational domains (e.g., servers, storage, and network) as well as the required expertise and staff development that are required to fully exploit vendor-supported APIs and the associated scripting languages.

- Evaluate the viability of implementing a route analytics solution to obtain visibility, analysis, and diagnosis capabilities of the issues that occur at the routing layer in complex, meshed networks, such as those found in public and hybrid cloud computing solutions.

- Increase their focus on managing services vs. focusing on managing individual technology domains.

- Work with the team responsible for data center LANs to determine the best way to get visibility into the traffic that goes between VMs on a given server.

- Implement the ability to perform standard management functions (e.g., troubleshooting, baselining) on a per-VM basis.

- Analyze the offerings of cloud computing service providers to determine if they provide APIs that can be leveraged to better manage public and hybrid cloud computing solutions.

- Implement a dynamic, highly scalable and integrated DNS/DHCP/IPAM solution, which is also well integrated with the virtual server management system.

- Evaluate service orchestration solutions relative to their ability to automate many of the manual tasks that are involved in provisioning and controlling the capacity of dynamic virtualized services.

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.

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

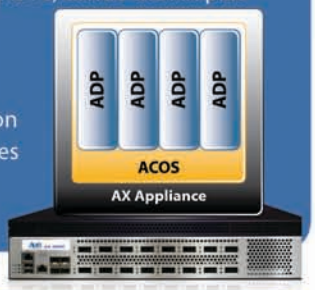

AX Virtual Chassis System (aVCS)

- 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.

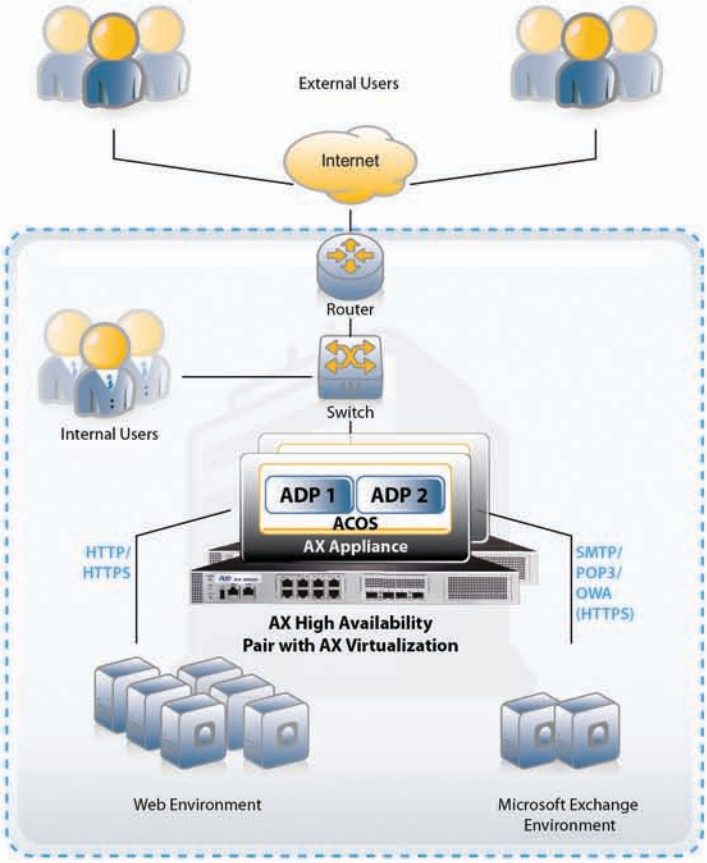

Summary

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

Software WAN Optimization

Accelerate Your Business

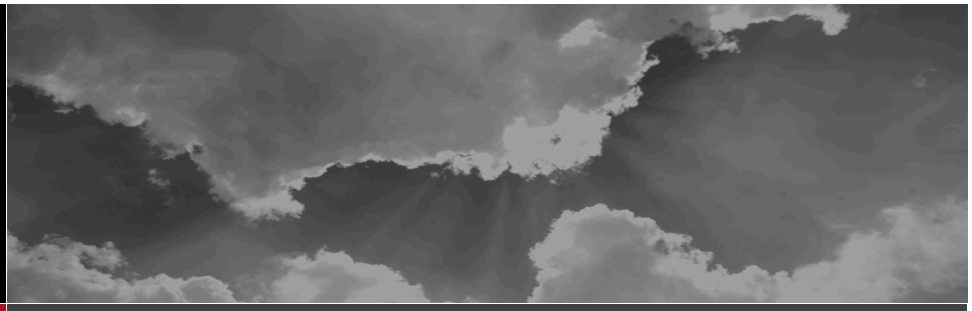

certeon®

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"aaS**ervice 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

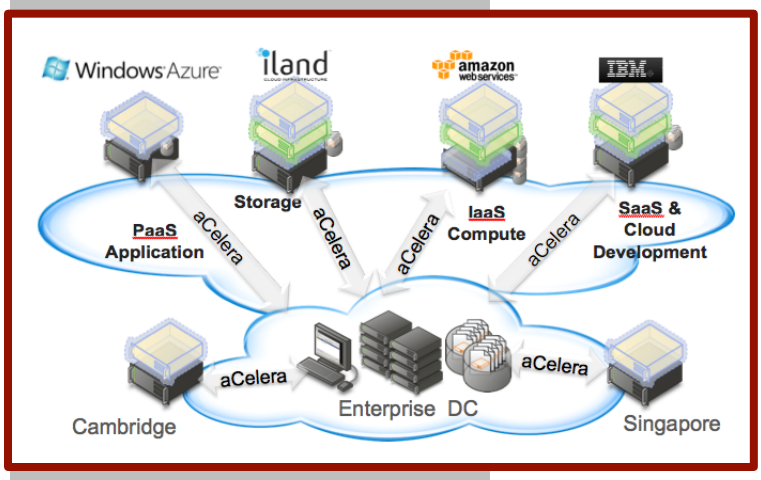

Enterprises and cloud service providers can deploy aCelera in any form factor, using any number of instances, delivering any throughput capacity, aligned with any application SLA requirement while meeting cost savings objectives and footprint limitations. This can be done in seconds on the enterprise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

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

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

certeon
THE Application Performance Company

1, Enterprise Strategy Group

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

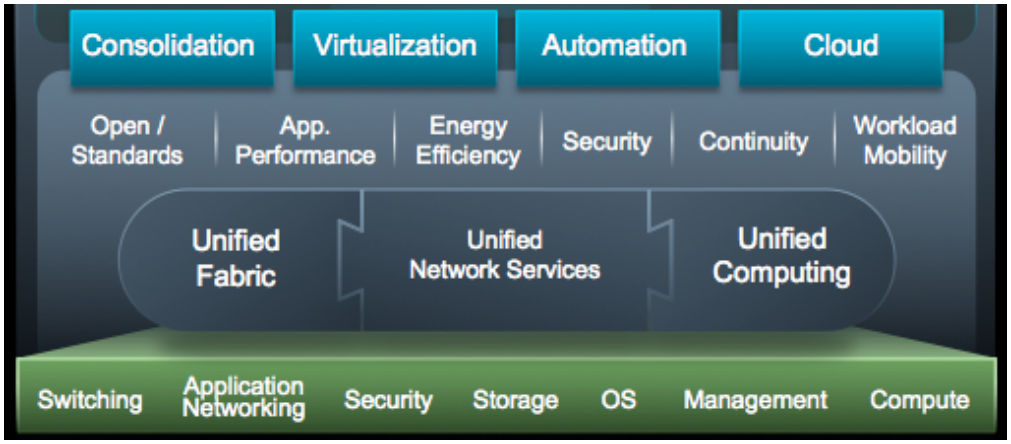

Unified Computing

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

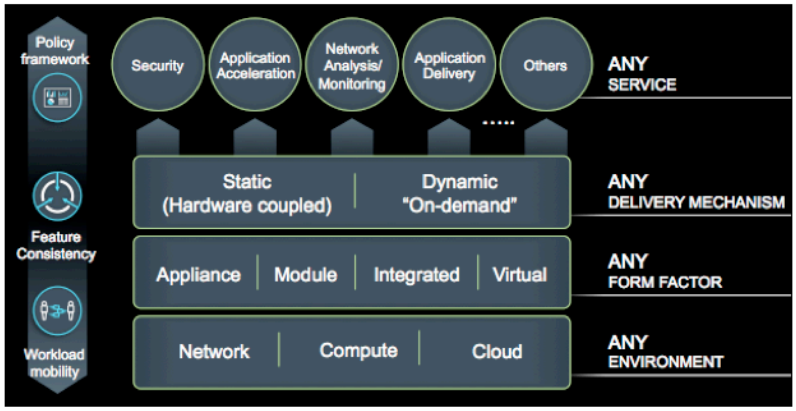

Policy-based centralized management. vWAAS is the industry's first cloud-ready WAN optimization solution. vWAAS works with the Cisco Nexus 1000v virtual switch's vPath capability to:

- 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.**

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.

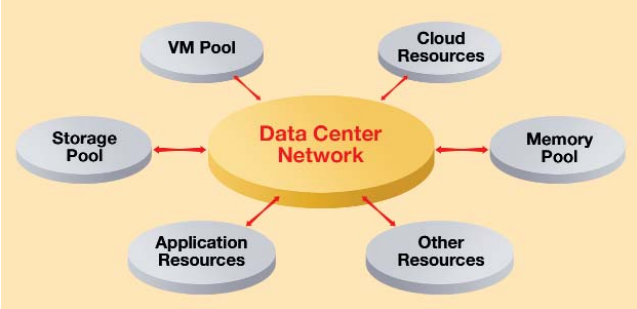

Figure 2. The virtualized data center is connected by the network

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.

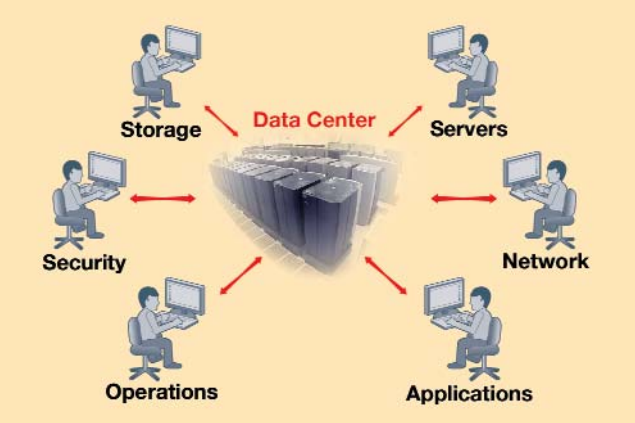

Figure 3. The virtual data center is everyone's responsibility

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

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

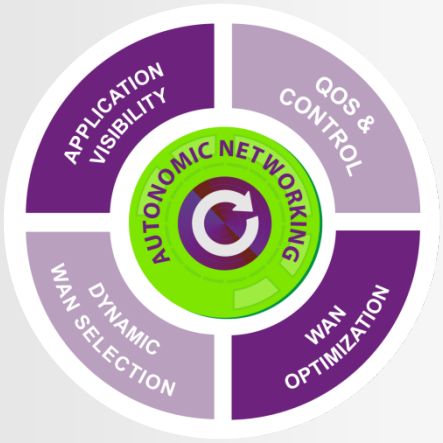

Powered by ANS™, WAN Governance brings tangible results

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

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.

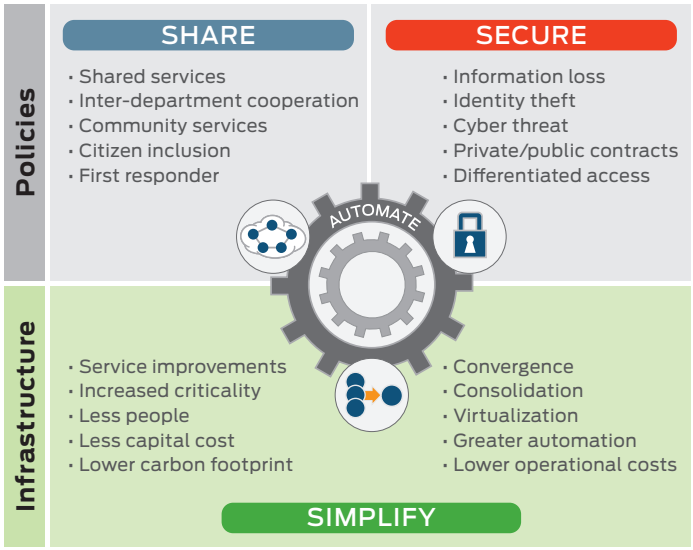

SIMPLIFY

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

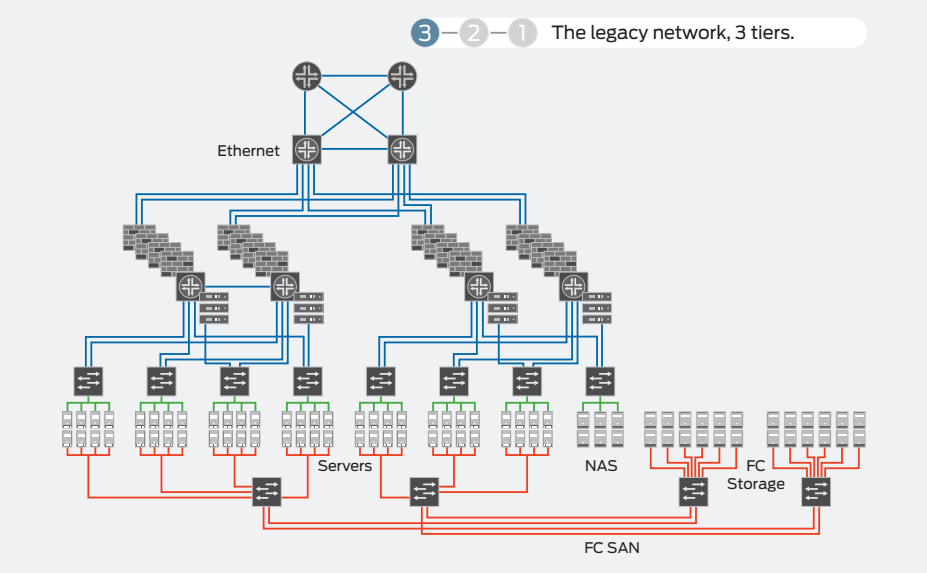

Figure 1: The legacy network

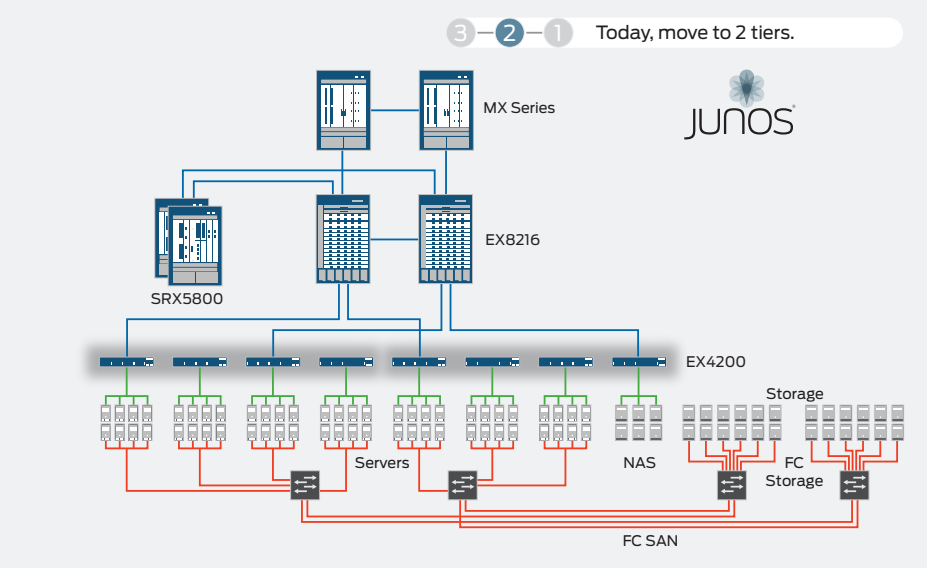

Figure 2: Juniper delivers a simplified two-tier network today with Virtual Chassis fabric technology.

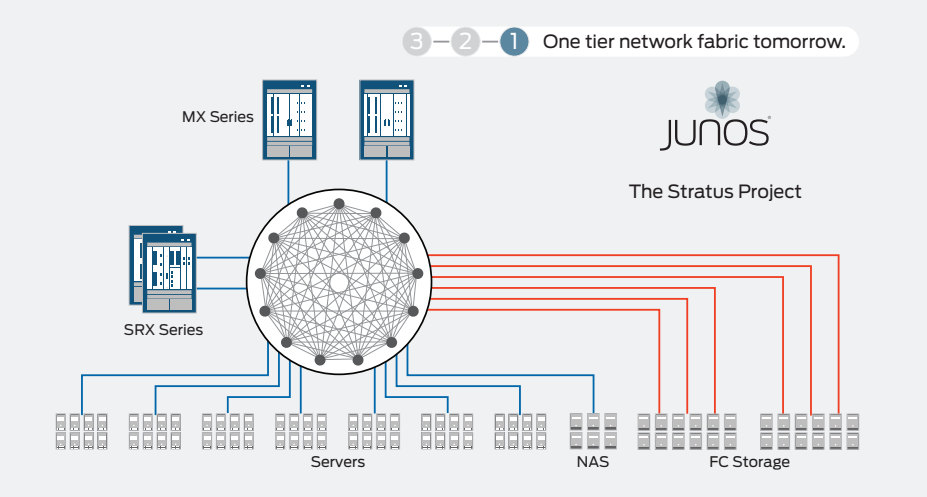

Figure 3: Juniper's vision for the ultimate simplification of the data center is Project Stratus, delivering a single fabric that unites Ethernet, Fibre Channel, and Infiniband networks.

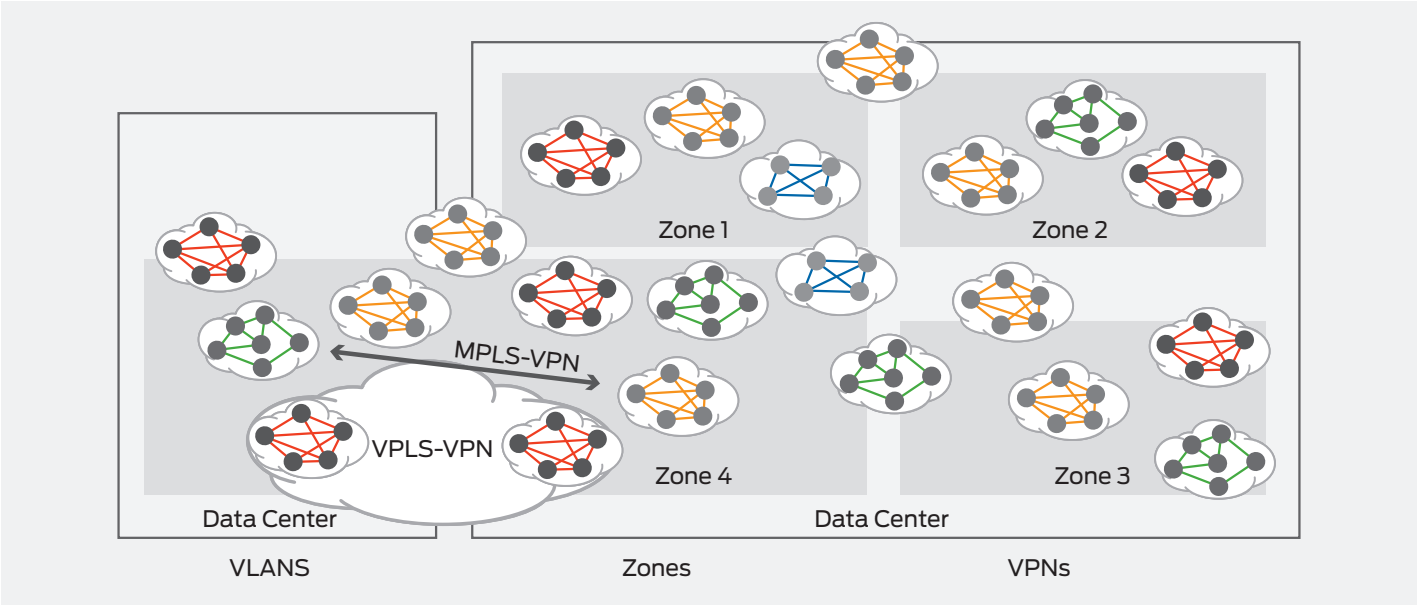

Figure 4: Scalable network virtualization technologies

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

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.

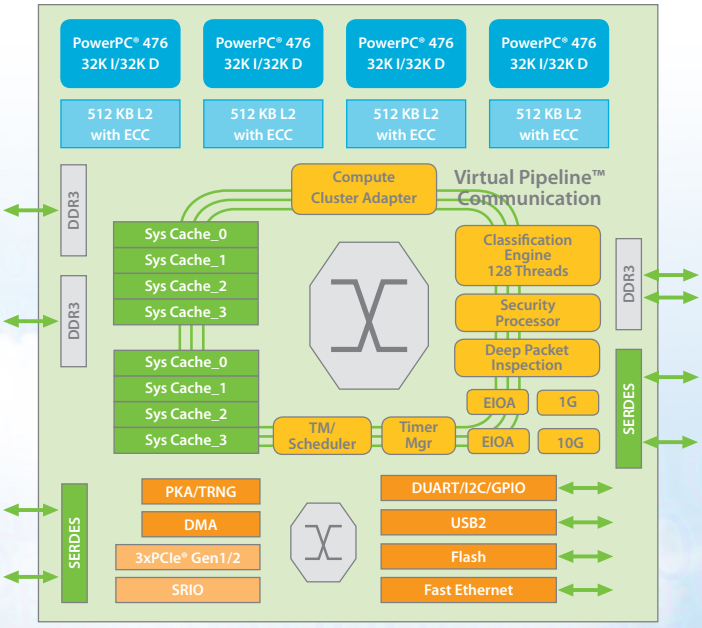

Figure 1 - ACP Block Diagram

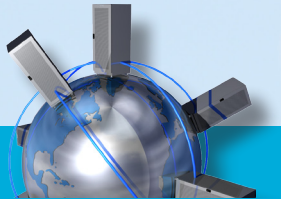

Axxia Intelligent Network Interface Card

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

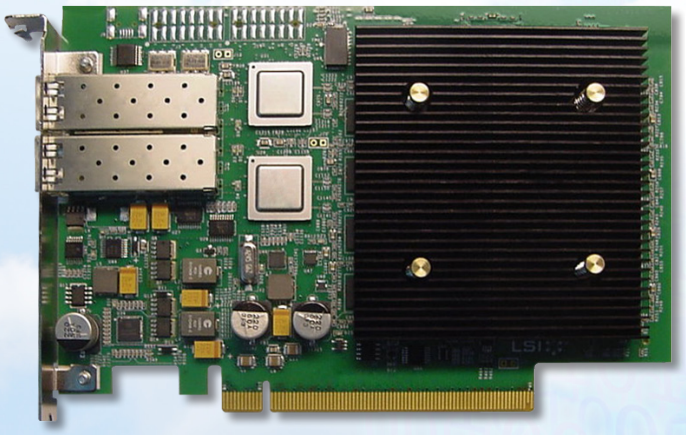

Axxia Intelligent NIC Software Architecture

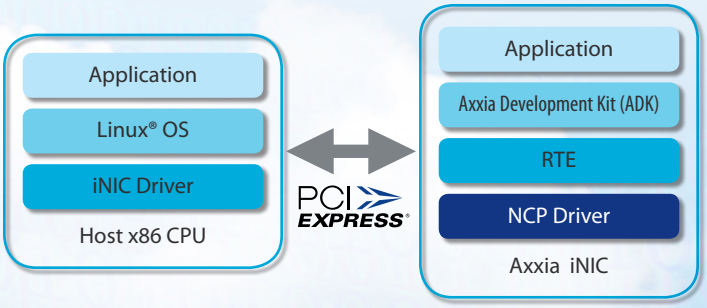

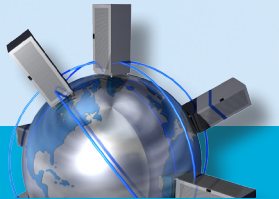

Packet Design Solutions:

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

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.

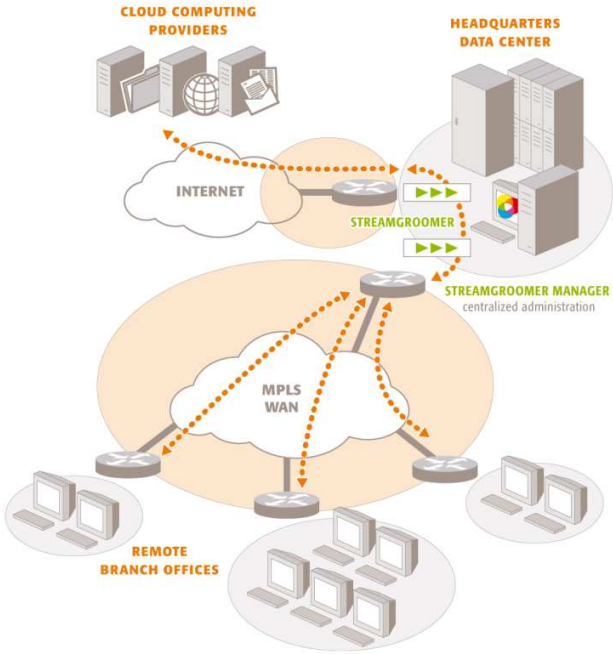

Fig. 1:
Centralized Internet access with branches over a private WAN

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.

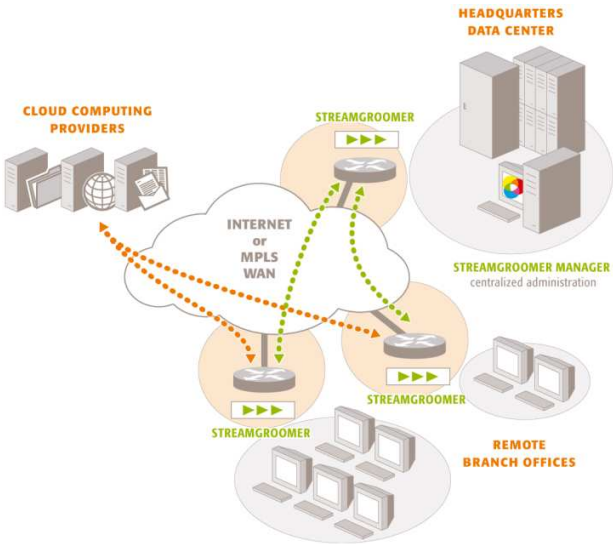

Fig. 2:
Branches with direct connections to the Internet

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.

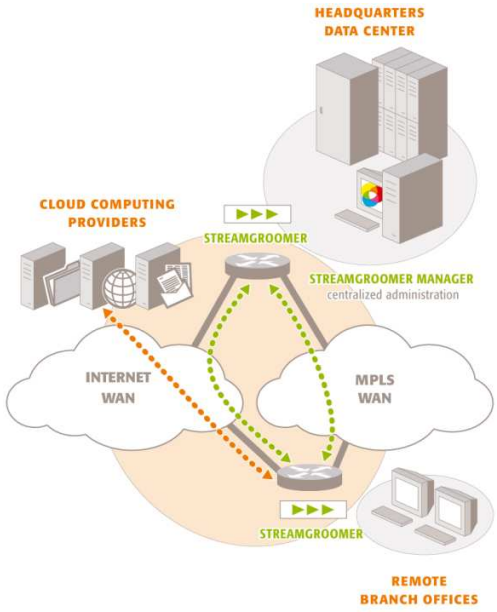

Fig. 3:
Hybrid network with load balancing per application performed by StreamGroomers

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.

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.

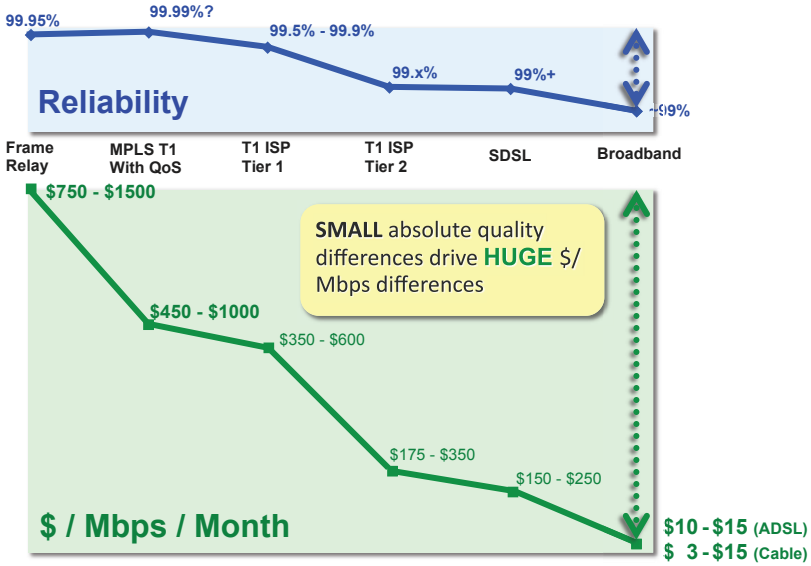

Figure 1: Private / Public WAN Pricing Disparity

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

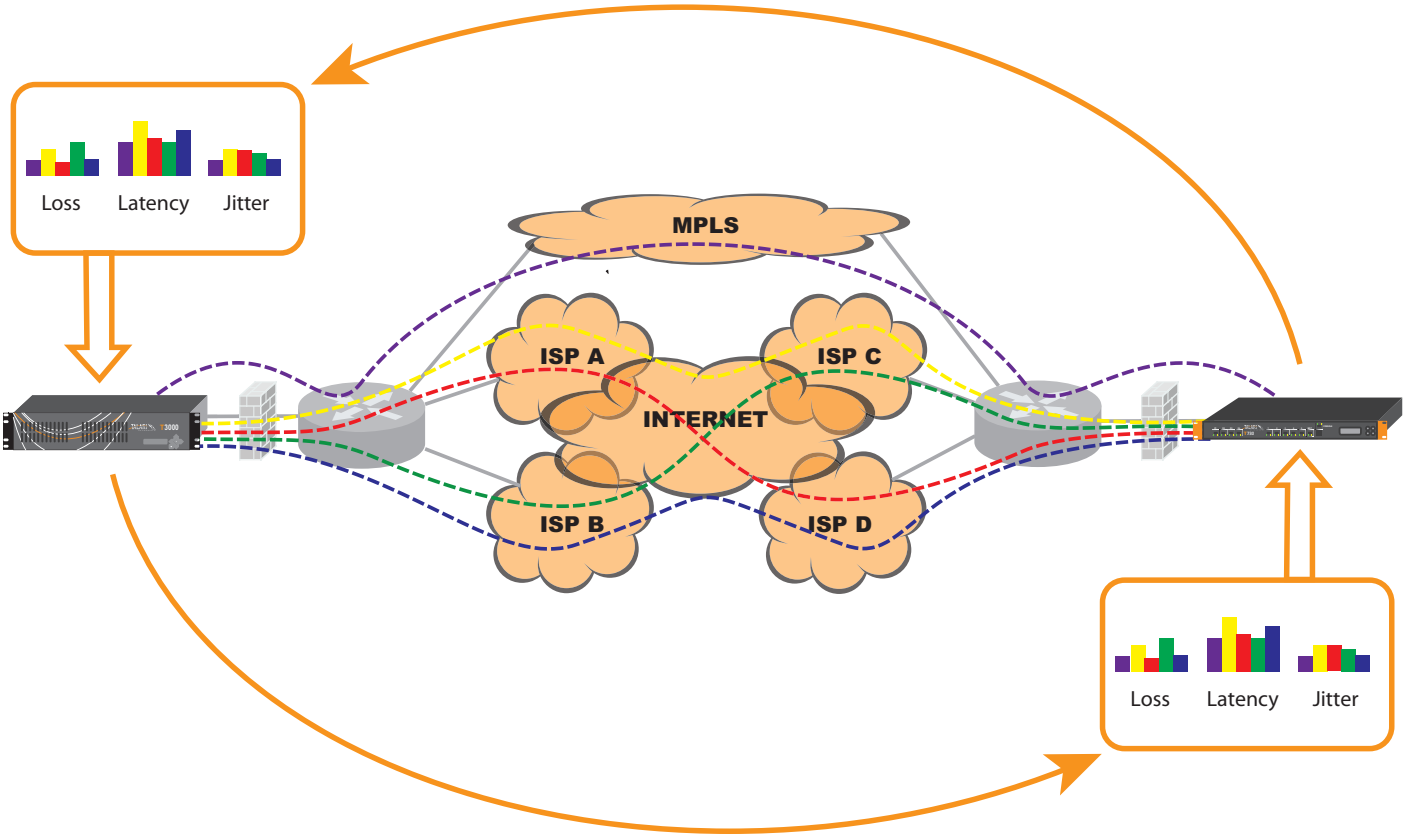

Figure 2: Continuous Measurment and Adaptation to Network Conditions

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**.

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.

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.

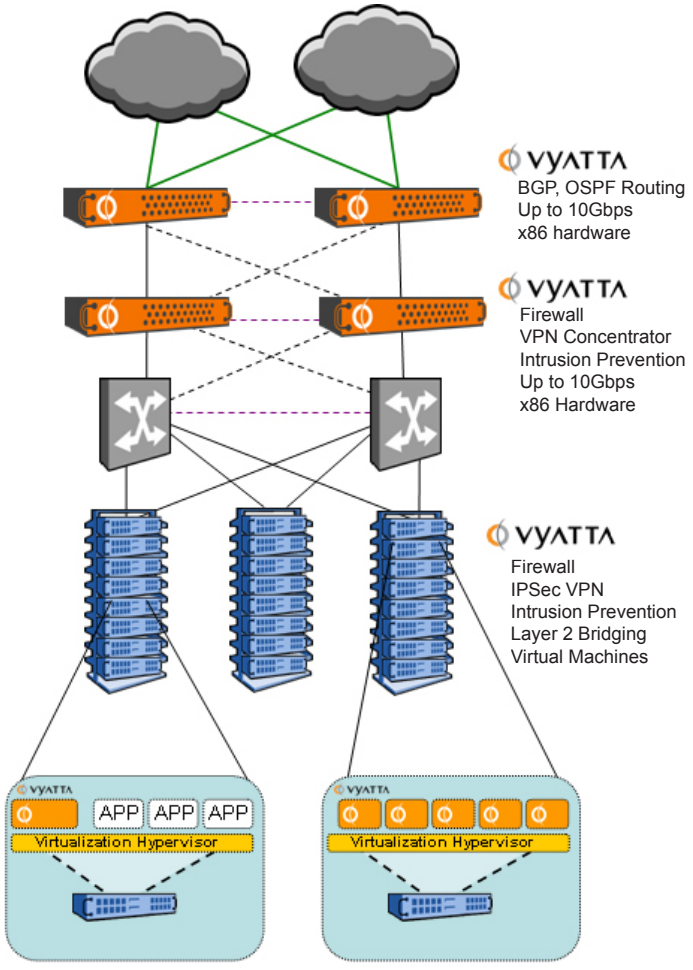

The Vyatta Network OS

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.

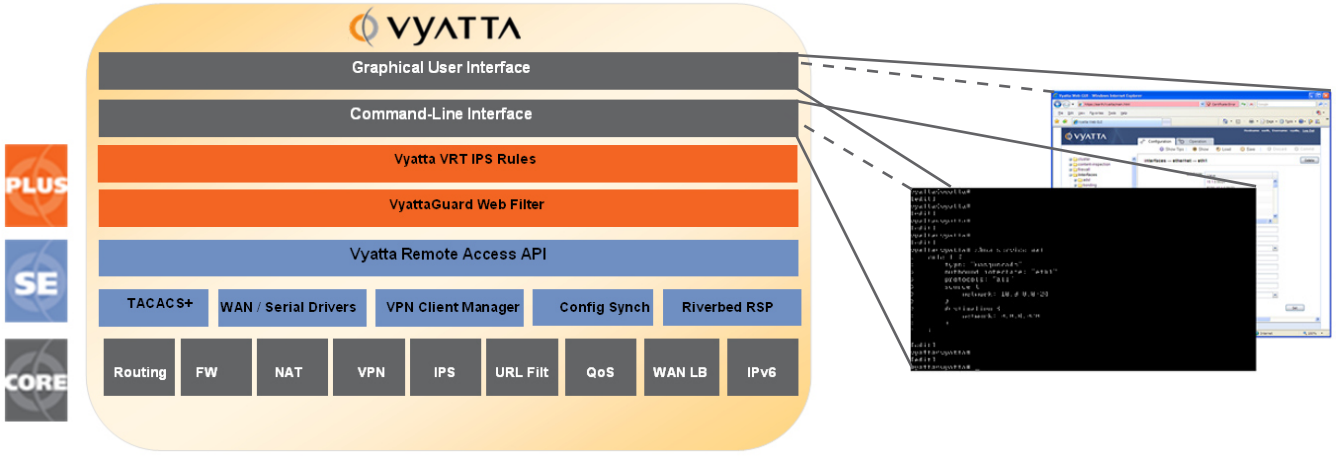

About Vyatta

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.