

The 2012 Cloud Networking Report

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

Sponsored in part by:



Produced by:



Table of Contents

INTRODUCTION AND FORWARD TO THE 2012 EDITION	1
EXECUTIVE SUMMARY	3
BACKGROUND	3
THE EMERGENCE OF CLOUD COMPUTING AND CLOUD NETWORKING	3
THE EMERGING DATA CENTER LAN	4
SOFTWARE DEFINED NETWORKS (SDN)	7
THE WIDE AREA NETWORK	8
MANAGEMENT & SECURITY	10
THE EMERGENCE OF CLOUD COMPUTING AND CLOUD NETWORKING	14
THE GOAL OF CLOUD COMPUTING	14
CHARACTERISTIC OF CLOUD COMPUTING SOLUTIONS	16
CLASSES OF CLOUD COMPUTING SOLUTIONS	17
<i>Private Cloud Computing</i>	<i>17</i>
<i>Public Cloud Computing</i>	<i>18</i>
<i>The Role of Virtualized Network Services</i>	<i>24</i>
<i>Hybrid Cloud Computing</i>	<i>26</i>
EMERGING PUBLIC CLOUD COMPUTING SERVICES	29
<i>Data Center Services</i>	<i>29</i>
<i>Cloud Networking Services</i>	<i>29</i>
THE CULTURE OF CLOUD COMPUTING	31
THE EMERGING DATA CENTER LAN	33
FIRST AND SECOND GENERATION DATA CENTER LANS	33
DRIVERS OF CHANGE	34
THIRD GENERATION DATA CENTER LAN ARCHITECTURE AND TECHNOLOGY OPTIONS	37
<i>Two Tier Data Center LAN Design</i>	<i>37</i>
<i>Alternatives to the Spanning Tree Protocol</i>	<i>39</i>
<i>Scalability of Two Tier LAN Designs</i>	<i>44</i>
<i>Network Support for Dynamic Creation and Movement of VMs</i>	<i>50</i>
<i>Network Virtualization</i>	<i>52</i>
<i>Network Convergence and Fabric Unification</i>	<i>57</i>
<i>Security Services in Virtualized Data Centers</i>	<i>60</i>
<i>Summary of Third Generation Data Center LAN Technologies</i>	<i>62</i>
SOFTWARE DEFINED NETWORKING (SDN)	64
THE SDN NETWORK ARCHITECTURE	66
OPEN NETWORKING FOUNDATION	68
OPENFLOW	71
<i>Potential Benefits of OpenFlow</i>	<i>74</i>
THE MARKETPLACE REALITY	76
CROSSING THE CHASM	81
A PLAN FOR SDN	82

THE WIDE AREA NETWORK (WAN).....	84
INTRODUCTION	84
<i>Background</i>	<i>84</i>
<i>Contrasting the LAN and the WAN.....</i>	<i>85</i>
<i>WAN Budgets.....</i>	<i>86</i>
<i>Drivers of Change.....</i>	<i>86</i>
<i>WAN Requirements.....</i>	<i>89</i>
TRADITIONAL WAN SERVICES	92
<i>Background</i>	<i>92</i>
<i>WAN Design Criteria and Challenges</i>	<i>92</i>
<i>Local Access to the Internet.....</i>	<i>94</i>
<i>Cloud Networking Without the Internet</i>	<i>95</i>
<i>Service Level Agreements.....</i>	<i>96</i>
OPTIMIZING THE PERFORMANCE OF IT RESOURCES	97
<i>Background</i>	<i>97</i>
<i>WAN Optimization Controllers (WOCs).....</i>	<i>98</i>
<i>Modeling Application Response Time.....</i>	<i>99</i>
<i>Application Delivery Controllers (ADCs).....</i>	<i>99</i>
<i>Virtual Appliances.....</i>	<i>100</i>
<i>Optimizing Access to Public Cloud Computing Solutions.....</i>	<i>102</i>
ALTERNATIVE WAN SERVICES	104
<i>An Internet Overlay</i>	<i>104</i>
<i>An Integrated Private-Public Solution.....</i>	<i>105</i>
<i>Dual ISP Internet VPN with Policy Based Routing.....</i>	<i>106</i>
<i>Hybrid WANs with Policy Based Routing.....</i>	<i>107</i>
<i>Aggregated Virtual WANs</i>	<i>107</i>
<i>Network-as-a-Service</i>	<i>109</i>
<i>Cloud-Based Network and Application Optimization</i>	<i>110</i>
<i>VPLS.....</i>	<i>111</i>
<i>Software Defined Networking (SDN).....</i>	<i>113</i>
EMERGING CLOUD NETWORKING SPECIFIC SOLUTIONS	114
<i>Cloud Balancing.....</i>	<i>114</i>
<i>WAN Optimization and Application Delivery for Cloud Sites.....</i>	<i>115</i>
MANAGEMENT & SECURITY	119
MANAGEMENT	119
<i>A New Set of Management Challenges.....</i>	<i>119</i>
<i>The Traditional Management Environment.....</i>	<i>124</i>
<i>The Emerging Management Environment.....</i>	<i>128</i>
<i>Application Performance Management.....</i>	<i>136</i>
<i>Management as a Cloud Provided Service.....</i>	<i>142</i>
SECURITY	143
<i>The Current Environment for Security Breaches.....</i>	<i>143</i>
<i>The Current Environment for Implementing Security.....</i>	<i>144</i>
<i>Security as a Cloud Provided Service.....</i>	<i>148</i>
<i>Web Application Firewall Services.....</i>	<i>150</i>
CONCLUSIONS AND OBSERVATIONS.....	153

Introduction and Forward to the 2012 Edition

Numerous analyst reports have pointed out the broad interest that IT organizations have in deploying one or more classes of cloud computing. For example, Gartner¹ recently stated that they expected that cloud computing would grow 19% in 2012, and would become a \$109 billion industry. Gartner also stated that they expected that by 2016, that cloud computing would be a \$207 billion industry. The high growth rate in the cloud computing market is in sharp contrast to the annual growth rate of the overall IT market, which Gartner estimates to be 3%. The broad interest in cloud computing is understandable given that the goal of cloud computing is to enable IT organizations to become dramatically more agile and cost effective and that evidence exists that that goal is achievable.

The primary goal of this report is to describe the network related 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.

As will be discussed in this report, a traditional network will not be able to successfully support cloud computing.

In order to support cloud computing, a cloud network must be dramatically more agile and cost effective than a traditional network.

In order to describe the networking challenges that are associated with enabling cloud computing, the rest of this section of the report will identify what cloud computing is today and will also describe how cloud computing is likely to evolve in the near term. Subsequent sections focus on the key components of a cloud network: Data Center LANs, WANs, and Network Management. A subsequent section will also focus on an emerging component of a cloud network: software defined networks (SDNs). Given the breadth of fundamental technology changes that are impacting the data center LAN, the data center LAN section is very technical. The sections on WANs, SDNs and Network Management are moderately technical. This year's edition of the cloud networking report leverages last year's edition of the report². However, every section of The 2011 Cloud Networking Report has been significantly updated to reflect the changes that have occurred in the last year.

As noted, the primary goal of this report is to describe the network related challenges and solutions that are associated with cloud networking. A secondary goal of this report is to identify how IT organizations are currently approaching both cloud computing and cloud networking and where possible, indicate how that approach is changing. To accomplish that goal, this report includes the results of surveys that were recently given to the subscribers of Webtorials.com. Throughout this report, the IT professionals who responded to those surveys will be referred to as the **Survey Respondents**. In some cases, the results of the surveys given to the Survey Respondents will be compared to the results of surveys given in 2011. In addition, the SDN section of The Report will include the results of a survey that was conducted in conjunction with

¹ <http://www.networkworld.com/news/2012/071312-gartner-cloud-260882.html>

² <http://www.webtorials.com/content/2011/11/2011-cloud-networking-report.html>

Information Week. Throughout the SDN section of this report, the IT professionals who responded to the SDN survey will be respectively to as the ***Information Week Respondents***.

The results of surveys such as the ones described in the preceding paragraph that ask IT organizations about their plans are always helpful because they enable IT organizations to see how their own plans fit with broad industry trends. Such surveys are particularly beneficial in the current environment when so much change is occurring.

Executive Summary

Background

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

The first section of this report will identify what cloud computing is today and will also describe how cloud computing is likely to evolve in the near term. Subsequent sections focus on the key components of a cloud network: Data Center LANs, WANs, and Network Management. There is also a separate section on Software Defined Networking (SDN). This year's edition of the cloud networking report leverages last year's edition of the report³. However, every section of [The 2011 Cloud Networking Report](http://www.webtorials.com/content/2011/11/2011-cloud-networking-report.html) has been significantly updated to reflect the changes that have occurred in the last year.

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

The Emergence of Cloud Computing and Cloud Networking

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

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

³ <http://www.webtorials.com/content/2011/11/2011-cloud-networking-report.html>

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

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

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

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

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

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

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

The Emerging Data Center LAN

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

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

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

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

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

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

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

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

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

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

VXLAN⁴ virtualizes the network by creating a Layer 2 overlay on a Layer 3 network via MAC-in-UDP encapsulation. The VXLAN segment is a Layer 3 construct that replaces the VLAN as the mechanism that segments the data center LAN for VMs. The VXLAN segment has a 24 bit VXLAN Network identifier and VXLAN is transparent to the VM, which still communicates using MAC addresses. NVGRE⁵ uses the GRE tunneling protocol defined by RFC 2784 and RFC 2890. NVGRE is similar in most respects to VXLAN with two major exceptions. While GRE encapsulation is not new, most network devices do not parse GRE headers in hardware, which may lead to performance issues and issues with 5-tuple hashes for traffic distribution in multi-path data center LANs. The other exception is that the current IETF NVGRE draft does not address the control plane question, leaving that for a future draft or possibly as something to be addressed by (Software Defined Networking) SDN controllers. STT⁶ is a third overlay technology for creating Layer 2 virtual networks over a Layer 2/3 physical network within the data center. Conceptually, there are a number of similarities between VXLAN and STT. However, STT encapsulation differs from NVGRE and VXLAN in two ways. First, it uses a stateless TCP-like header inside the IP header, which allows tunnel endpoints within end systems to take advantage of TCP segmentation offload (TSO) capabilities of existing TOE server NICs. STT also allocates more header space to the per-packet metadata, which provides added flexibility for the virtual network control plane.

⁴ <http://searchservervirtualization.techtarget.com/news/2240074318/VMware-Cisco-propose-VXLAN-for-VM-mobility>

⁵ <http://tools.ietf.org/html/draft-sridharan-virtualization-nvgre-00>

⁶ <http://tools.ietf.org/html/draft-davie-stt-01>

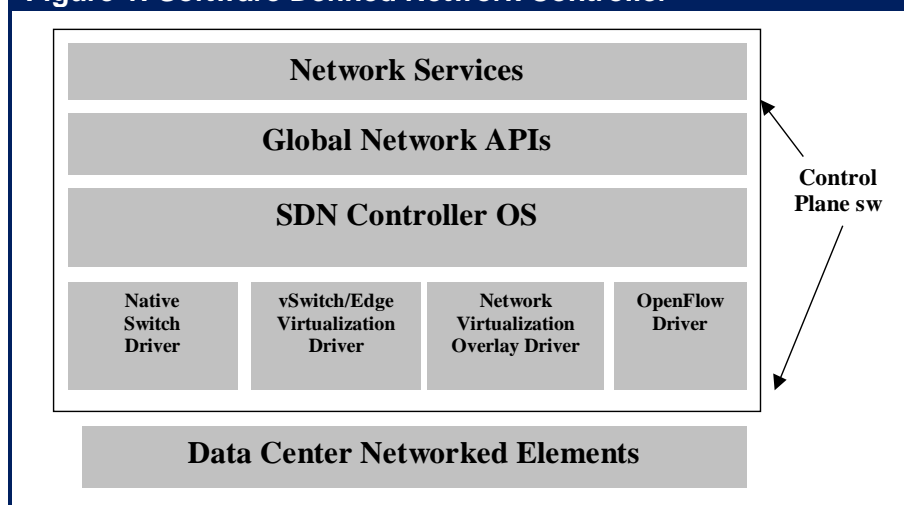
Security is generally considered by enterprise IT departments to be the primary concern in today's highly virtualized data centers and in the implementation of private or public cloud computing environments. In the traditional data center, internal security has generally been implemented by deploying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

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

Software Defined Networks (SDN)

As is typical of emerging technologies and new approaches to networking, there is currently somewhat of a broad definition relative to how the industry, particularly vendors, define SDN. The most common way that SDN is described is based on a layered architecture as shown in **Figure 16**. In **Figure 16**, the control plane function is centralized in SDN Controller software that is installed on a server or on a redundant cluster of servers for higher availability and performance. The SDN controller is used to control the actions of the subtending networked elements.

Figure 1: Software Defined Network Controller



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

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

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

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

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

The Wide Area Network

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

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

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

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

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

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

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

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

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

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

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

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

Management & Security

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

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

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

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

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

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

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

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

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

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

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

A key component of application performance management is the ability to perform root cause analysis. A prerequisite to being able to perform effective root cause analysis is the automatic discovery of all the elements in the IT infrastructure that support each service or application. For example, if IT organizations can effectively identify which components of the infrastructure

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

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

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

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

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

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

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

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

The Emergence of Cloud Computing and Cloud Networking

The Goal of Cloud Computing

Within the IT industry there still isn't a universally accepted definition of what is meant by cloud computing. The Report 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.

In order to demonstrate the concept behind the phrase *good enough*, consider just the availability of an IT service. In those cases in which the IT service is business critical, *good enough* could mean five or six 9's of availability. However, in many other cases *good enough* has the same meaning as *best effort* and in these cases *good enough* could mean two or three 9's of availability. The instances in which an approach that provides two or three 9's of availability is acceptable are those instances in which the IT service isn't business critical and that approach is notably less expensive than an alternative approach that offers higher availability.

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.

In most instances the SLAs that are associated with public cloud computing services such as Salesforce.com or Amazon's Simple Storage System are weak and as such, it is reasonable to say that these services are delivered on a best effort basis. For example, the SLA⁷ that Amazon offers for its Amazon Web Services (AWS) states that, "AWS will use commercially reasonable efforts to make Amazon EC2 available with an Annual Uptime Percentage of at least 99.95% during the Service Year." As part of the Amazon definition of Annual Uptime Percentage, Amazon excludes any outage of 5 minutes or less. The Amazon SLA also states that if their service doesn't meet the Annual Uptime Percentage commitment, the customer will receive 10% off its bill for the most recent month that the customer included in the SLA claim that it filed.

A key attribute of the vast majority of the SLAs that are associated with public cloud computing services is that they don't contain a goal for the end-to-end performance of the service. The reason for the lack of performance guarantees stems from the way that most public cloud computing services are delivered. As shown in **Figure 2**, one approach to providing public cloud computing services is based on the service being delivered to the customer directly from an independent software vendor's (ISV's) data center via the Internet. This is the distribution model currently used for Salesforce.com's CRM application. Another approach is for an ISV to

⁷ <http://aws.amazon.com/ec2-sla/>

leverage an IaaS provider such as Amazon to host their application on the Internet. Lawson Software's Enterprise Management Systems (ERP application) and Adobe's LiveCycle Enterprise Suite are two examples of applications hosted by Amazon EC2. Both of these approaches rely on the Internet and it is not possible to provide end-to-end quality of service (QoS) over the Internet. As a result, neither of these two approaches lends itself to providing an SLA that includes a meaningful commitment to critical network performance metrics such as delay, jitter and packet loss.

The fact that cloud computing service providers (CCSPs) don't provide an end-to-end performance SLA for applications delivered over the Internet will not change in the foreseeable future. However, as will be described in a subsequent section of this report, there are things that can be done to improve the performance of applications delivered over the Internet.

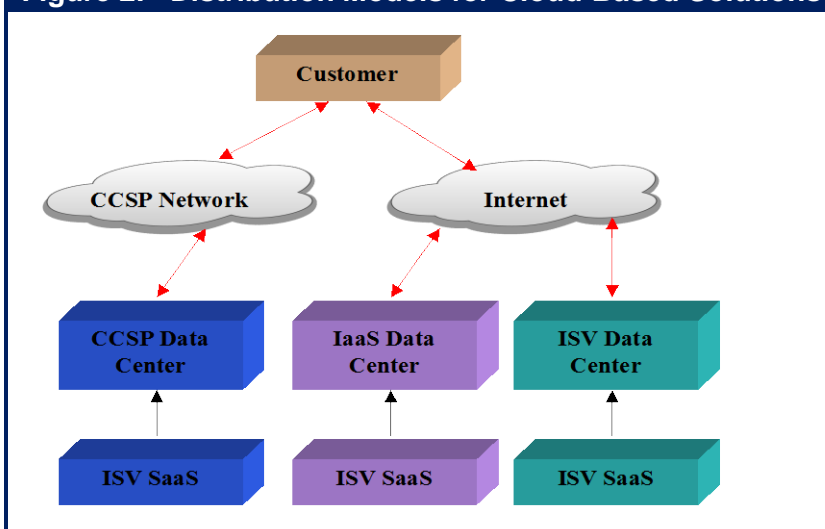
An approach to providing public cloud computing services that does lend itself to offering more meaningful SLAs is based on a CCSP providing these solutions to customers

from the CCSP's data center and over a network that is provided by the CCSP and based on a technology such as MPLS.

Organizations that utilize best effort cloud computing services do so with the implicit understanding that if the level of service they experience is not sufficient; their primary recourse is to change providers. It may seem counter-intuitive that a company would utilize public cloud computing services for which end-to-end performance SLAs are essentially non-existent. However, as described in a subsequent section of this report, two thirds of The Webtorials Respondents indicated that the SLAs that they receive from their network service providers for services such as MPLS 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.

Figure 2: Distribution Models for Cloud-Based Solutions



Characteristic 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

There are three classes of cloud computing solutions that will be described in this section of the report. Those classes are private, public and hybrid.

Private Cloud Computing

Many IT organizations have decided to implement some of the characteristics of cloud computing solutions described in the preceding subsection within their internal IT environment. This approach is usually referred to as a *Private Cloud*. One of the primary ways that IT organizations have adopted private cloud computing solutions is by implementing some or all of the previously mentioned characteristics of cloud computing solutions in order to be able to provide Infrastructure-as-a-Service (IaaS) solutions that are similar to the solutions offered by IaaS providers such as Rackspace.

The initial set of IaaS solutions that were brought to market by IaaS providers were the basic compute and storage services that are necessary to run applications. However, the IaaS market is highly dynamic and IaaS providers are deploying myriad new services including:

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

The Survey Respondents were given a set of 7 possible approaches to IaaS and were asked to indicate which approach best described their company's approach to using IaaS solutions, either provided internally by their own IT organization, or provided externally by an IaaS provider. The Survey Respondents were allowed to indicate as many approaches as were appropriate. Their responses are shown in **Table 1**.

Table 1: Approach to IaaS		N=171
Approach	Percentage of Respondents	
We are in the process of developing a strategy	48.0%	
We provide IaaS solutions internally for a wide range of applications	19.9%	
We provide IaaS solutions internally for a small range of applications	19.9%	
We have a well-defined and understood strategy	15.2%	
We only use IaaS solutions from a CSP for a small set of applications that are not business critical	14.6%	
We use IaaS solutions from a CCSP for a wide range of applications	12.3%	
Other	7.0%	
We only outsource either a trial of the initial deployment of an application to a CCSP	6.4%	
We have a policy against using any IaaS solutions provided by a CCSP	3.5%	

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

Roughly half of all IT organizations are currently in the process of developing a strategy for how they will use public and private IaaS solutions.

The Survey Respondents were asked to indicate the two primary factors that limit their company's interest in using internally provided IaaS solution. The five inhibitors to the adoption of private IaaS solutions that were indicated the most times by the Survey Respondents and the percentage of times that they were mentioned were:

- Concerns about the security and confidentiality of data (36.3%)
- Their lack of an internal strategy about IaaS (28.7%)
- Their lack of personnel to design and implement the solutions (25.7%)
- The relative immaturity of the technologies that would have to be installed and managed (19.9%)
- The lack of significant enough cost savings (19.3%)

While the conventional wisdom in our industry is that security and confidentiality of data is the major impediment to the adoption of public cloud based IaaS solutions, it is somewhat surprising that:

Concern about the security and confidentiality of data is the primary impediment to the broader adoption of private IaaS solutions.


Public Cloud Computing

This section of The Report will focus on the two most popular types of public cloud computing solutions: Software-as-a-Service and Infrastructure-as-a-Service.

Software-as-a-Service

According to Gartner⁸, the Software as a Service (SaaS) market will have worldwide revenues of \$22.1 billion by 2015. One of the key characteristics of the SaaS marketplace is that:

The SaaS marketplace is comprised of a small number of large players such as Salesforce.com, WebEx and Google Docs as well as thousands of smaller players.



One of the reasons why there are so many players in the SaaS market is that the barrier to entry is relatively low.

⁸ <http://www.slideshare.net/rajeshdgr8/global-saa-s-2012>

The Survey Respondents were asked to indicate if their company currently acquires applications from a SaaS provider or if they are likely to within the next twelve months. Their responses are shown in **Figure 3**.

The Survey Respondents were then given a set of 7 types of applications and were asked to indicate the types of applications that their company currently acquires from a SaaS provider and the types of applications that their organization would likely acquire from a SaaS provider over the next twelve months. Their responses are shown in **Table 2**.

Table 2: Interest in SaaS		<i>N=153</i>
	Currently Acquire	Will Acquire
Collaboration	55%	31%
Customer Relationship Management (CRM)	53%	22%
Human Resources	45%	18%
Office Productivity	40%	33%
Project and Portfolio Management	27%	54%
Enterprise Resource Planning (ERP)	24%	16%
Supply Chain Management (SCM)	15%	27%

The Survey Respondents were given a set of ten factors and were asked to indicate the two factors that were the primary drivers of their organization's interest in using SaaS solutions. The responses of the Survey Respondents are shown in **Table 3**. In **Table 3**, the column on the right is labeled *Percentage of Respondents*. That column contains the percentage of the Survey Respondents that indicated that the factor in the left hand column of **Table 3** was one of the two primary drivers of their organization's interest in using SaaS solutions.

Table 3: Factors Driving the Adoption of SaaS Solutions		<i>N=153</i>
Factor	Percentage of Respondents	
Lower cost	39%	
Reduce the amount of time it takes to implement an application	35%	
Free up resources in the IT organization	29%	
Deploy applications that are more robust; e.g., available and scalable	27%	
Easier to justify OPEX than CAPEX	26%	
Leverage the expertise of the SaaS provider	19%	
Reduce risk	11%	
Management mandate as our strategic direction	8%	
Meet temporary requirements	3%	
Other	2%	

One conclusion that can be drawn from the data in **Table 3** is that:

The primary factors that are driving the adoption of SaaS are the same factors that drive the adoption of any form of out-tasking.

Given the concerns that IT organizations have relative to the security and confidentiality of their data, it appears to be counter intuitive that 11% of the Survey Respondents indicated that reducing risk was a factor that would cause them to use a public cloud computing solution. In most cases the Survey Respondents' reasoning was that acquiring and implementing a large software application (e.g., ERP, CRM) presents considerable risk to an IT organization and one way to minimize this risk is to acquire the functionality from a SaaS provider.

Infrastructure as a Service (IaaS)

The barrier to enter the IaaS marketplace is notably higher than is the barrier to enter the SaaS marketplace. That is one of the primary reasons why there are fewer vendors in the IaaS market than there are in the SaaS market. Representative IaaS vendors include Amazon, AT&T, CSC, GoGrid, IBM, Joyent, NTT Communications, Orange Business Services, Rackspace, NaviSite (acquired by Time Warner), Savvis (acquired by Century Link), Terremark (acquired by Verizon) and Verizon. As the preceding sentence indicates, the IaaS market is going through a period that is characterized by mergers and acquisitions. The IaaS market is also expected to exhibit significant growth in the next few years. For example, Gartner⁹ estimates that the IaaS market will grow from \$3.7 billion in 2011 to \$10.5 billion in 2014.

The Survey Respondents were asked to indicate the IaaS services that their organization currently acquires from a CCSP and the services that their organization will likely acquire from a CCSP during the next year. Their responses are shown in **Table 4**.

Table 4: Current and Planned Adoption of IaaS Services			<i>N = 142</i>
	Currently Acquire	Will Likely Acquire	
Storage	26.8%	16.9%	
Computing	26.8%	9.2%	
Virtual Private Data Center	17.6%	14.1%	
Disaster Recovery	16.2%	21.8%	
High Performance Computing	10.6%	9.9%	

Because storage and computing were the initial set of IaaS services that were brought to market, it was not at all surprising to see that over a quarter of the Survey Respondents indicated that they currently used those services. In addition, given that high performance computing (HPC) is somewhat of a niche application, it was not surprising that there was relatively little interest in acquiring HPC from an IaaS supplier. However it was somewhat of a surprise to see that:

There is strong interest on the part of IT organizations in acquiring both virtual private data center and disaster recovery services from IaaS providers.

⁹ http://www.gas.com/company/data-quality-news/iaas_market_to_record_strong_growth_7178.htm

Drivers and Inhibitors

This section will discuss the factors that are driving and the factors that are inhibiting the deployment of IaaS solutions.

- **Drivers**

The Survey Respondents were given a set of eleven factors and were asked to indicate the two factors that were the primary drivers of their organization's interest in using Cloud-based IaaS solutions. The responses of the Survey Respondents are shown in **Table 5**. In **Table 5**, the column on the right is labeled *Percentage of Respondents*. That column contains the percentage of the Survey Respondents that indicated that the factor in the left hand column of **Table 5** was one of the two primary drivers of their organization's interest in using Cloud-based IaaS solutions.

Table 5: Factors Driving the Adoption of IaaS Solutions		<i>N = 171</i>
Factor	Percentage of Respondents	
Lower cost	30.4%	
The ability to dynamically add capacity	30.4%	
Reduce time to deploy new functionality	26.3%	
Obtain functionality we are not able to provide ourselves	22.2%	
Deploy more highly available solutions	19.3%	
Free up resources	17.0%	
Easier to justify OPEX than CAPEX	15.8%	
Prefer to only pay for services that we use	14.0%	
Satisfy temporary requirements	11.7%	
Other	4.7%	
Our strategy is to use IaaS providers wherever possible	4.1%	
Leverage the security expertise of the provider	4.1%	

The conventional wisdom in the IT industry is that lower cost is the primary factor driving the adoption of Cloud-based IaaS solutions and that factors such as the ability to dynamically add new capacity, while important, are nowhere near as important. As the data in **Table 5** highlights, the reality is that the ability to dynamically add new capacity is as important a driver of the adoption of Cloud-based IaaS solutions as is lowering cost. In addition, another very important driver of the adoption of Cloud-based IaaS solutions is the ability to reduce the time it takes to deploy new functionality. It is reasonable to look at the ability to dynamically add capacity and the ability to reduce the time it takes to deploy new functionality as two components of a single factor – agility. Looked at this way,

By a wide margin, agility is the most important factor driving the adoption of Cloud-based IaaS solutions.

- **Inhibitors**

The Survey Respondents were asked to indicate the two primary factors that limit their company's interest in using a Cloud-based IaaS solution. Those factors and the percentage of times that they were indicated by the Survey Respondents are shown in **Table 6**.

Table 6: Inhibitors to the adoption of Cloud-based IaaS Solutions <i>N</i> = 171	
Factor	Percentage of Respondents
We are concerned about the security and confidentiality of our data	57.9%
We don't see significant enough cost savings	24.0%
The lack of time and resources to sufficiently analyze the offerings and the providers	19.9%
Uncertainty about the provider living up to their promises	19.9%
We have concerns about the availability of the solutions	16.4%
Our lack of confidence in a shared infrastructure	15.2%
The lack of a meaningful SLA	14.6%
We don't believe that the gains in the agility of these solutions justifies the cost and/or the risk	11.7%
Our policy is to either limit or totally avoid using IaaS providers	8.8%
The provider is not capable of adding capacity in a dynamic enough fashion	4.7%

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

Concern about the security and confidentiality of data is by a wide margin the number one factor inhibiting the adoption of Cloud-based IaaS solutions

A component of the concerns that IT organization have about security and confidentiality stems from the overall increase in the sophistication of hackers. For example, until relatively recently the majority of security attacks were caused by individual hackers, such as Kevin Mitnick, who served five years in prison in the late 1990s for computer- and communications-related hacking crimes. The goal of this class of hacker is usually to gain notoriety for themselves and they often relied on low-technology techniques such as dumpster diving.

However, over the last few years a new class of hacker has emerged and this new class of hacker has the ability in the current environment to rent a botnet or to develop their own R&D lab. This new class includes crime families and hactivists such as Anonymous. In addition, some national governments now look to arm themselves with Cyber Warfare units and achieve their political aims by virtual rather than by physical means.

The sophistication of the current generation of hackers was highlighted in the Blue Coat Systems 2012 Web Security Report¹⁰, which focused on a number of topics including malnets and social networking. A malware network, or malnet, gathers users, most frequently when they are visiting trusted sites and routes them to malware. According to the Blue Coat Report, “In 2011, malnets emerged as the next evolution in the threat landscape. These infrastructures last beyond any one attack, allowing cybercriminals to quickly adapt to new vulnerabilities and repeatedly launch malware attacks. By exploiting popular places on the Internet, such as search engines, social networking and email, malnets have become very adept at infecting many users with little added investment.”

The report noted the increasing importance of social networking and stated that, “Since 2009, social networking has increasingly eclipsed web-based email as a method of communications.” The report added that, “Now, social networking is moving into a new phase in which an individual site is a self-contained web environment for many users – effectively an Internet within an Internet.” For example, according to the Blue Coat report 95% content types that are found on the Internet are also found within social networking sites. The five most requested subcategories of content that were requested from social networking sites, and the percentage of times that they were requested are shown in **Table 7**.

Table 7: Most Requested Content from Social Media Sites	
Subcategory of Content	Percentage of Times it was Requested
Games	37.9%
Society/Daily Living	23.8%
Personal Pages/Blogs	6.4%
Pornography	4.9%
Entertainment	4.2%

Part of the challenge that is associated with social network sites being so complex is that IT organizations cannot just look at a social media site as one category and either allow or deny access to it. Because these sites contain a variety of classes of content, IT organizations need the granular visibility and control to respond differently to requests at the same social media site for different types of content.

Another component of the concern that IT organizations have about security and confidentiality of their data stems from the fact that in most cases IT organization perceive that there is a higher security risk if their data is being stored on a device that is shared with other users which is typically the case when an IT organization is using an IaaS solution. The security risk that is associated with all forms of cloud computing was discussed in IBM's X-Force 2011 Trend and Risk Report¹¹ that was published in March 2012. According to the IBM report, in 2011, there were many high profile cloud breaches affecting well-known organizations and large populations of their customers.

¹⁰ http://www.bluecoat.com/sites/default/files/documents/files/BC_2012_Security_Report-v1i-optimized.pdf

¹¹ [X-Force 2011 Trend and Risk Report](#)

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

The Role of Virtualized Network Services

As previously noted, one of the primary goals of The Report is to identify what functionality is needed in the network to support cloud computing. With that goal in mind, the Survey Respondents were given a number of questions that related to the role that virtualized network services play in their evaluation and selection of Cloud-based IaaS services.

One of the questions contained a set of network services and the Survey Respondents were asked to indicate if they thought the network service should be part of a Cloud-based IaaS service and if they did, whether they preferred to manage the network service themselves or have the CSP manage it. The vast majority of the Survey Respondents (87+%) thought that each one of the network services listed in **Table 8** should be part of a Cloud-based IaaS service. Columns two and three of **Table 8** respectively contain the percentage of the Survey Respondents who prefer to manage the service themselves as well as the percentage of the Survey Respondents who prefer to have a CSP manage the service.

Table 8: The Applicability and Management of Network Services N = 171		
Network Service	Manage Ourselves	CSP Manage
Load Balancer	61.9%	38.1%
SSL Load Balancer	62.2%	37.8%
Firewall	81.4%	18.6%
WEB application firewall	68.5%	31.5%
IDS/IPS	64.1%	35.9%
VPN	70.2%	29.8%
WAN optimization	50.8%	49.2%

One obvious conclusion that can be drawn from the data in **Table 8** is:

There is a strong desire on the part of IT organizations to manage the security related network services that are part of an IaaS service.

Because IT organizations expect that Cloud-based IaaS services are supported by a wide range of network services, this raises the question, “When evaluating IaaS services, how carefully do IT organizations evaluate the associated network services?” To answer that question, the Survey Respondents were asked, “When your organization evaluates cloud services such as computing, storage and virtual private data centers, how carefully does your organization evaluate the enabling network services such as Load Balancer, SSL Load Balancer, Firewall?” Their answers are contained in **Table 9**.

Table 9: Importance of Network Services		N = 171
How Carefully	Percentage of Respondents	
We don't evaluate them at all	8.6%	
We look at them as a check-off item, but don't evaluate	10.0%	
We pay some attention to them, but they are not a major component of the evaluation process	21.4%	
They are a major component of the overall evaluation process	33.6%	
They are a critical component of the overall evaluation process	26.4%	

One obvious conclusion that can be drawn from the data in **Table 9** is:

The evaluation of the supporting network services is a key component of the overall process of evaluating IaaS solutions.

Given the critical role that network services play in the evaluation of Cloud-based IaaS services, the Survey Respondents were asked to indicate the two most important criteria they look for when evaluating network services such as a Load Balancer, an SSL Load Balancer, or a Firewall, that enable cloud services. The criteria and the percentage of times that they were indicated by a survey respondent are shown in **Table 10**.

Table 10: Criteria to Evaluate Networking Services		N = 171
Criteria	Percentage of Respondents	
A robust feature set similar to traditional networking equipment	25.9%	
The ability to grow/shrink the capacity of the service on demand	23.8%	
The ability to rapidly provision the network service; e.g., 5 minutes or less	21.1%	
The ability to only pay for what we use	17.8%	
A brand name vendor	6.3%	
The ability to charge back to business units based on usage	5.1%	

The conventional wisdom is that when IT organizations evaluate network services, that a name brand vendor is an important criterion. The data in **Table 10** refutes that belief as the data in the table highlights the fact that a robust feature set is the single most important criterion that IT organizations examine with evaluating networks services. However, another way to evaluate the data in **Table 10** is based on the previous definition of agility¹². Looked at this way, the data in **Table 10** clearly indicates that the agility of network services is the most important criterion that IT organizations examine with evaluating networks services.

¹² In this context, agility is the ability to dynamically add capacity and the ability to reduce the time it takes to deploy new functionality.

In order to understand the organizational dynamic that underlies the decision to use an IaaS solution from a CSP, the Survey Respondents were asked about the roles of the organizations that are involved in making that decision. Their responses, shown in **Table 11**, indicate how the decision is made.

Table 11: The Decision Making Process <i>N=160</i>	
Role	Percentage of Respondents
Largely by the IT organization with some input from the business or functional unit	40.0%
The IT unit and the business or functional unit participate equally	26.3%
Largely by the business or functional unit with some input from the IT organization	15.6%
Entirely by the IT organization	11.3%
Entirely by the business or functional unit	6.9%

One obvious conclusion that can be drawn from the data in **Table 11** is:

Roughly 20% of the times that a company is evaluating public IaaS solutions, the company's IT organization is either not involved at all or plays a minor role.

Hybrid Cloud Computing


Like so much of the terminology of cloud computing, there is not a uniformly agreed to definition of the phrase **hybrid cloud computing**. According to Wikipedia¹³, "Hybrid cloud is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together, offering the benefits of multiple deployment models. Briefly it can also be defined as a multiple cloud systems which are connected in a way that allows programs and data to be moved easily from one deployment system to another."

Based on this definition, one form of a hybrid cloud is an n-tier application in which the web tier is implemented within one or more public clouds while the application and database tiers are implemented within a private cloud. Another form of hybrid cloud that receives a lot of attention is cloud balancing. The phrase **cloud balancing** refers to routing service requests across multiple data centers based on myriad criteria. As shown in **Figure 4**, cloud balancing involves one or more corporate data centers and one or more public cloud data centers.

Cloud balancing can be thought of as the logical extension of global server load balancing (GSLB).

¹³ http://en.wikipedia.org/wiki/Cloud_computing#Hybrid_cloud

Figure 4: Cloud Balancing



The goal of a GSLB solution is to support high availability and maximum performance. In order to do this, a GSLB solution typically makes routing decisions based on criteria such as the application response time or the total capacity of the data center. A cloud balancing solution may well have as a goal supporting high availability and maximum performance and may well make routing decisions in part based on the same criteria as used by a GSLB solution. However, a cloud balancing solution extends the focus of a GSLB solution to a solution with more of a business focus. Given that extended focus, a cloud balancing solution includes in the criteria that it uses to make a routing decision the:

- Performance currently being provided by each cloud
- Value of the business transaction
- Cost to execute a transaction at a particular cloud
- Relevant regulatory requirements

Some of the benefits of cloud balancing include the ability to:

- **Maximize Performance**
Routing a service request to a data center that is close to the user and/or to one that is exhibiting the best performance results in improved application performance.
- **Minimize Cost**
Routing a service request to a data center with the lowest cost helps to reduce the overall cost of servicing the request.
- **Minimize Cost and Maximize Service**
Cloud balancing enables a service request to be routed to a data center that provides a low, although not necessarily the lowest cost while providing a level of availability and performance that is appropriate for each transaction.

- **Regulatory Compliance**

For compliance with regulations such as PCI, it may be possible to partition a web services application such that the PCI-related portions remain in the PCI-compliant enterprise data center, while other portions are cloud balanced. In this example, application requests are directed to the public cloud instance unless the queries require the PCI-compliant portion, in which case they are directed to the enterprise instance.

- **Manage Risk**

Hosting applications and/or data in multiple clouds increases the availability of both. Balancing can be performed across a number of different providers or it can be performed across multiple independent locations of a single cloud service provider.

Emerging Public Cloud Computing Services

Data Center Services

Most of the IaaS providers do not want to compete entirely based on providing commodity services such as basic compute and storage. As such, many IaaS providers are implementing higher value-added data center services such as the ones described below.

Private Cloud Data Center Services

These services are based on outsourcing the enterprise's multi-tier private data center to a service provider. The data center could be located at either a site controlled by the enterprise or at a service provider's site. In most cases service providers will structure these services so that the customers receive the highest levels of support, as well as assurances written into the corresponding SLA for high levels of availability, performance and security. A private WAN service would typically be used to provide access to these services.

Virtual Private Data Center (VPDC)

These services provide an instance of an entire data center hosted on a service provider's infrastructure that is optimized to provide a high level of security and availability for multiple tenants. From the service provider's perspective, the data center architecture for the VPDC would be similar to the architecture used for a private cloud data center except that the resources would be shared among a number of customers rather than being dedicated to a single customer or tenant. The service provider's architecture needs to effectively leverage virtualization in order to maximize the efficient usage of a shared pool of resources. The architecture also needs to allow for a high degree of flexibility in providing a broad range of required network capabilities. This includes WAN optimization, load balancing and firewall services. Service management software should be in place to enable the co-management of the VPDC by customers and providers.

The hybrid cloud computing model works best in those instances in which the VPDC and the private cloud data center are based on the same hypervisors, hypervisor management systems and cloud controllers. This maximizes the enterprise's control over the hybrid cloud and allows application and server management to remain the responsibility of the enterprise. Access to a VPDC could be provided either over the Internet or a private WAN service.

Cloud Networking Services

With the exception of collaboration, the applications that organizations have historically acquired from CCSPs have typically been enterprise applications such as CRM. Recently, a new class of solutions has begun to be offered by CCSPs. These are solutions that have historically been provided by the IT infrastructure group itself and include network and application optimization, VoIP, Unified Communications (UC), security, network management and virtualized desktops. Within The Report, this new class of solutions will be referred to as [Cloud Networking Services \(CNSs\)](#).

The Survey Respondents were given a set of 7 CNSs and were asked to indicate the CNSs that their organization currently acquires from a CCSP and the services that their organization will likely acquire from a CCSP during the next year. Their responses are shown in **Table 12**.

Table 12: Current and Planned Adoption of CNSs		<i>N = 142</i>
	Currently Acquire	Will Likely Acquire
VoIP	20.4%	17.6%
Network Management	19.7%	8.5%
Security	18.3%	9.9%
Unified Communications	15.5%	23.2%
Application Performance Management	10.6%	10.6%
Network and Application Optimization	8.5%	9.2%
Virtual Desktops	7.0%	19.0%

The data in **Table 12** shows that the interest in CNS is quite broad, as over twenty-five percent of the survey respondents indicated that over the next year that five of the seven services listed in the table would either likely be acquired, or would be acquired.

Cloud Networking Services represents the beginning of what could be a fundamental shift in terms of how IT services are provided.

Since CNS solutions are just one more form of public cloud computing, when evaluating these solutions IT organizations need to understand the degree to which these solutions overcome the factors that impede the use of any public cloud computing solution. Since concerns about security are typically one of the primary impediments to the adoption of public cloud computing solutions, evaluating the security of the CNS provider's facilities is a critical component of evaluating a CNS solution.

However, just as important as whether or not the CNS solution provides adequate security is whether or not the solution actually provides the benefits that drive IT organizations to use public cloud computing solutions. As previously discussed, the primary benefit of using a public cloud computing solution is typically lower cost. While it can be tricky to compare the usage sensitive pricing of the typical CNS solution with the fully loaded cost of a premise based solution, the cost information provided by the CCSP should give the IT organization all the information it needs to do that analysis. Another key benefit of using a public cloud computing solution is being able to reduce the time it takes to deploy new functionality. Evaluating the agility of a CCSP is notably more difficult than evaluating their cost structure.

One way for an IT organization to evaluate the agility of a CCSP is to identify the degree to which the CCSP has virtualized their infrastructure.

This follows because a virtual infrastructure is notably easier to initialize, scale and migrate than a physical infrastructure is. Since the vast majority of CCSPs implement virtualized servers, server virtualization is unlikely to distinguish one CCSP from another. What can distinguish one CCSP from another is the degree to which they have virtualized other components of their infrastructure, most notably their network. That is one of the reasons why a subsequent section of The Report will discuss network virtualization.

The Culture of Cloud Computing

The rest of The Report will discuss the networking technologies that enable cloud computing. However, as much as cloud computing is about technologies it is also about changing the culture of the IT organization. One such cultural shift was described in the preceding subsection entitled “The Goal of Cloud Computing”.

To put this cultural shift into perspective, it is important to realize that it is implicit in the traditional IT culture to implement ongoing enhancements to make the network and the IT services that are delivered over the network, increasingly resilient. The adoption of cloud computing changes that model and as previously described, in some instances it is becoming acceptable for IT services to be delivered on a best effort basis. A clear indication of that change is the success of Salesforce.com. Salesforce.com has three million customers who use their solutions to support critical sales processes. Yet in spite of the importance of the application, in virtually all cases Salesforce.com will not give a customer an availability guarantee and since the application is typically accessed over the Internet, it doesn't come with an end-to-end performance guarantee.

One of the other cultural shifts that is associated with the adoption of cloud computing is that IT organizations become less of a provider of IT services and more of a broker of IT services. In the traditional IT environment, the IT organization is the primary provider of IT services. Part of the challenge that is associated with the IT organization being the primary provider of IT services is that sometimes the IT organization can't meet the needs of the business units in a timely fashion. In the past the way that business unit managers have dealt with this lack of support is by having their own shadow IT organization whereby the business unit managers have some people on their staff whose role is to provide the IT services that the business unit manager can't get from the IT organization¹⁴. In the current environment, public cloud providers often play the role of a shadow IT organization by providing a company's business unit managers services or functionality that they either can't get from their IT organization or they can't get in a timely manner. In some instances the IT function is in a position to stop the non-sanctioned use of public cloud computing once they find out about it. However, in many other instances they aren't.

Instead of trying to prevent business unit managers from acquiring public cloud services, a better role for an IT organization is to modify their traditional role of being the primary provider of IT services and to adopt a role in which they provide some IT services themselves and act as a broker between the company's business unit managers and cloud computing service providers for other services. In addition to contract negotiations, the IT organization can ensure that the acquired application or service doesn't create any compliance issues, can be integrated with other applications as needed, can scale, is cost effective and can be managed.

IT organizations provide considerable value by being the broker between the company's business unit managers and cloud computing service providers.

Another cultural change that is associated with the adoption of cloud computing is the implementation of more usage sensitive chargeback. Usage sensitive chargeback is not new. Many IT organizations, for example, allocate the cost of the organization's network to the company's business unit managers based on the consumption of that network by the business

¹⁴ The data in Table 11 provides some insight into how often this occurs.

units. Since there has traditionally been a lot of overhead associated with usage sensitive chargeback, usage sensitive chargeback has only made sense in those situations in which the IT organization is in a position both to explain to the business unit managers in easily understood language, what they are paying for and to provide suggestions as to how the business unit managers can reduce their cost. In the current environment, roughly fifty percent of all IT organizations implement usage sensitive chargeback for at least some components of IT. However, relatively few implement it broadly. Input from the Survey Respondents indicates that over the next two years IT organizations will make increased use of usage sensitive chargeback. Most of this increased use will come from having the business unit managers pay the relevant cloud computing service providers for the services that their organization consumes. The movement to implement more usage sensitive chargeback over the next two years will not be dramatic because:

The culture of an IT organization changes very slowly.

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, which was the only viable, high-speed First Generation LAN technology, 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 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.
- Support for applications that are neither bandwidth intensive nor sensitive to latency.
- Switches with relatively low port densities.
- High over-subscription rate on uplinks.
- The separation of the data network from the storage network.
- VLANs to control broadcast domains and to implement policy.
- The need to primarily support client server traffic; a.k.a., north-south traffic.
- Redundant links to increase availability.
- Access Control Lists (ACLs) for rudimentary security.
- The application of policy (QoS settings, ACLs) based on physical ports.

Drivers of Change

One of the key factors driving IT organizations to redesign their data center LANs is the requirement to support the growing deployment of virtual servers. With that in mind, The Survey 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. Their responses are shown in **Table 13**.

Table 13: Deployment of Virtualized Servers					N = 112
	None	1% to 25%	26% to 50%	51% to 75%	76% to 100%
Have already been virtualized	18%	30%	25%	16%	11%
Expect to be virtualized within a year	11%	28%	24%	25%	12%

The way to read the data in **Table 13** is that in the current environment only 18% of IT organizations have not virtualized any data center servers and that within a year, that only 11% of IT organizations will not have virtualized any of their data center servers.

As pointed out in [Virtualization: Benefits, Challenges and Solutions](#)¹⁵, server virtualization creates a number of challenges for the data center LAN. 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. In order to quantify the extent to which IT organizations move VMs between physical servers, The Survey Respondents were asked to indicate whether they agreed or disagreed with the statements in the left hand column of **Table 14**.

Table 14: Movement of VMs		N = 265
	Agree	Disagree
We currently manually migrate VMs between servers in the same data center	66%	34%
We currently automatically migrate VMs between servers in the same data center	55%	45%
We currently manually migrate VMs between servers in disparate data centers	48%	52%
We currently automatically migrate VMs between servers in disparate data centers	26%	74%

The data in **Table 14** indicates the great interest that IT organizations have in moving VMs between physical servers. However, as will be described throughout this section of the report, moving VMs between physical servers can be very complex.

¹⁵ <http://www.webtutorials.com/content/2010/06/virtualization.html>

Manually configuring parameters such as QoS settings and ACLs in order to support the dynamic movement of VMs is not the only challenge that is associated with server virtualization. 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.
- Limited VM-to-VM Traffic Visibility**
 Traditional vSwitches 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 server domains in both private, public and hybrid clouds.
- 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 that is causing IT organizations to redesign their data center LANs. The left hand column in [Table 15](#) contains a list of the factors that are driving data center redesign. The center column shows the percentage of The Survey Respondents who in 2011 indicated that the corresponding factor was the primary factor that is driving their organization to redesign their data center LAN. The right hand column shows the percentage of The Survey Respondents who recently indicated that the corresponding factor was the primary factor that is driving their organization to redesign their data center LAN.

Table 15: Factors Driving Data Center LAN Redesign		<i>N = 265</i>
Factor	% of The Survey Respondents In 2011	% of The Survey Respondents in 2012
To reduce the overall cost	24.6%	20.8%
To support more scalability	20.8%	9.1%
To create a more dynamic data center	12.6%	10.2%
To support server virtualization	12.1%	14.0%
To reduce complexity	5.3%	12.5%
To make it easier to manage and orchestrate the data center	13.0%	14.3%
To support our storage strategy	3.4%	3.4%

Table 15: Factors Driving Data Center LAN Redesign		N = 265
Factor	% of The Survey Respondents In 2011	% of The Survey Respondents in 2012
To reduce the energy requirements	1.0%	0.8%
Other (please specify)	3.4%	9.1%
To make the data center more secure	3.9%	6.0%

The data in **Table 15** indicates that a broad range of factors are driving IT organizations to re-design their data center LANs. There is, however, significant overlap between some of the factors in **Table 15**. For example, there is significant overlap between creating a more dynamic data center and supporting server virtualization. There is also significant overlap between reducing complexity and making it easier to manage and orchestrate the data center. Combining the factors that overlap indicates that:


The primary factors driving IT organizations to re-design their data center LAN is the desire to reduce cost, support server virtualization and reduce complexity.

The conventional wisdom in the IT industry is that the cost of the power consumed by data center LAN switches is not significant because it is a small percentage of the total amount of power that is consumed in the typical data center. There is the potential for that situation to change going forward as 10 Gbps, 40 Gbps and 100 Gbps LAN interfaces will potentially consume considerably more power than 1 Gbps LAN interfaces currently do. As such, a requirement of third generation data center LAN switches is that the amount of power that they consume is only marginally more than what is consumed by second generation data center LAN switches and that these switches provide functionality to intelligently manage the power consumption during off peak hours.

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. Alternative technologies included Ethernet, Token Ring, FDDI/CDDI, 100VG-AnyLAN 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 leveraging Ethernet to be 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.

Many of the technologies that are discussed in this chapter and in the chapter on Software Defined Networks are still under development and will not be standardized for another year or two. In order to understand whether or not IT organizations account for emerging technologies in their planning, The Survey Respondents were asked to indicate their company's planning horizon for the evolution of their data center LANs. To avoid ambiguity, the survey question stated "A planning horizon of three years means that you are making decisions today based on the technology and business changes that you foresee happening over the next three years." Their answers are shown in **Figure 5**.



The data in **Figure 5** indicates that almost 75% of IT organizations have a planning horizon of three years or longer. Since most of the technologies discussed in this chapter will be standardized and ready for production use in three years, that means that the vast majority of IT organizations can incorporate most of the technologies discussed in this chapter into their plans for data center LAN design and architecture.

Below is a discussion 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.

Two Tier Data Center LAN Design

There are many on-going IT initiatives that are aimed at improving the cost-efficiency of the enterprise data center. This includes server virtualization, Services Oriented Architecture (SOA), Web 2.0, access to shared network storage as well as the implementation of HPC and cluster computing. In many cases these initiatives are placing a premium on IT organizations being able to provide highly reliable, low latency, high bandwidth communications among both physical and virtual servers. Whereas the hub and spoke topology of the traditional three-tier

Second Generation LAN was optimized for client-to-server communications that is 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.

As discussed below, 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. As is also discussed below, the requirement for increased reliability and availability creates a requirement for redundant switch configurations in both tiers of the network.

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 10-20 VMs per server that is typical today. With more VMs per server, I/O requirements increase proportionally. Thankfully, the traditional economics of Ethernet performance improvement¹⁶ is falling into place for 10 Gigabit Ethernet (10 GbE). As a result, Third Generation data center 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, plus 100 GbE ports when these are available and become cost-effective for data center applications. These high-speed ports will be used for multiple purposes, including 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 768 non-blocking 10 GbE ports or 192 40 GbE ports. The typical maximum port density for TOR switches which are generally based on merchant silicon, is 64 10 GbE ports (or alternatively 48 10 GbE ports and 4 40 GbE ports). Today, high-speed uplinks are often comprised of multiple 10 GbE links that leverage Link Aggregation (LAG)¹⁷. 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 whereby a majority of traffic is concentrated in a small number of flows. Most high performance modular switches

¹⁶ 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.

¹⁷ www.ieee802.org/3/hssg/public/apr07/frazier_01_0407.pdf

already have a switch fabric that provide 100 Gbps of bandwidth to each line card, which means that as 40 GbE and 100 GbE line cards become available, these can be installed on existing modular switches, preserving the investment in these devices. Most vendors of modular switches are currently shipping 40 GbE line cards, while 100 GbE line cards will not be widely deployed until 2013 or later due primarily to economic considerations. Currently, most 100 GbE deployments have restricted to service providers, such as Internet exchanges.

In the case of stackable Top of Rack (ToR) switches, adding 40 or 100 GbE uplinks often requires new switch silicon, which means that at least some of the previous generation of ToR switches will need to be swapped out in order to support 40 GbE and, at some future date, 100 GbE uplink speeds.

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 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. Multi-chassis Link Aggregation Group is described below. Implementing this technology also tends to increase availability because it enables IT organizations to dual home servers to separate physical switches.

Alternatives to the Spanning Tree Protocol

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. Another approach is to implement one of the two emerging shortest path first bridging protocols, TRILL and SPB, that eliminate loops and support equal cost multi-path bridging. TRILL and SPB are also described below.

Switch Virtualization and Multi-Chassis Link Aggregation Group

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) technology¹⁸, as shown in **Figure 6**. MC LAG allows the links of the LAG to span the multiple physical switches that comprise a virtual switch. The re-convergence time associated with MC LAG is typically under 50 ms., which means that real time applications such as voice are not impacted by the re-convergence of the LAN. 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 6**, 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 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 that is not shown in **Figure 6**. 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.

¹⁸ http://en.wikipedia.org/wiki/Link_aggregation

Figure 6: 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 by vendors 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. From the server perspective, 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 for some differentiation between LAN switches by improving load balancing across the MC LAG links. In addition, there are some differences in the number of ports or links 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.

Currently MC Lags are based on proprietary implementations that have a variety of different names. As a result, MC LAG interoperability between switches from different vendors cannot be expected. Most vendors recommend MC LAG 2 tier topologies similar to the one shown on **Figure 6**. MC LAG are generally not recommended in configurations with more than two aggregation switches, such as large 2 tier fat tree topologies.

SPB and TRILL

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 flows 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) standard for a Layer 2 shortest-path first (SPF) routing protocol for Ethernet. The TRILL RFC (RFC 6325) is currently supported by some vendors as part of their proprietary Layer 2 fabric implementations. However, most of the current implementations of TRILL are based on pre-standard drafts in combination with added proprietary features and are not interoperable. In the future, vendors that provided early support for TRILL are likely to offer two versions: openTRILL which is strictly standards compliant and interoperable and a proprietary fabric solution based partly on TRILL.

Shortest Path Bridging (SPB) as defined in IEEE 802.1aq is a competing standard for equal cost multi-path bridging Ethernet fabrics. There are two variants of SPB: SPBM where packets are encapsulated at the edge using 802.1ah MAC-in-MAC frame formats and SPBV where packets are tagged with 802.1D/802.1ad tags. Three switch vendors (Avaya, Alcatel Lucent, and Huawei) have demonstrated interoperability with SPBM.

With either TRILL or 802.1aq SPB, it would be possible to achieve load-balanced, active-active link redundancy without having to resort entirely to switch virtualization, MC LAG, and VSL/VSI interconnects. For example, dual homing of servers can be based on MC LAG to a virtual access switch comprised of two physical access switches, while the rest of the data center LAN is based on TRILL or SPB.

There is currently considerable debate in the industry about which is the best technology – TRILL or SPB. While that is an important debate:

In many cases, the best technology doesn't end up being the dominant technology in the marketplace.

TRILL and SPB have some points of similarity but they also have some significant differences that preclude interoperability. Both approaches use IS-to-IS as the Layer 2 routing protocol and both support equal cost multi-path bridging, which eliminates the blocked links that are a characteristic of STP. Both approaches also support edge compatibility with STP LANs. Some of the major differences include:

- TRILL involves a new header for encapsulation of Ethernet packets, while SPB uses MAC-in-MAC Ethernet encapsulation. Therefore, TRILL requires new data plane hardware, while SPB doesn't for Ethernet switches that support 802.1ah (MAC-in-MAC), 802.1ad (Q-in-Q) and 802.1ag (OAM).
- SPB's use of MAC-in-MAC Ethernet encapsulation eliminates the potential for a significant increase in the size of MAC address tables that are required in network switches.
- SPB forwards unicast and multicast/broadcast packets symmetrically over the same shortest path, while TRILL may not forward multicast/broadcast packets over the shortest path.
- SPB eliminates loops using Reverse Path Forwarding (RPF) checking for both unicast and multicast traffic, while TRILL uses Time to Live (TTL) for unicast and RPF for multicast.
- TRILL can support multi-pathing for an arbitrary number of links, while SPB is currently limited to 16 links.

- TRILL is supported by vendors with large market share in LAN switching. SPB is currently supported by vendors with a relatively small market share.
- With TRILL, Layer 2 network virtualization is limited to 4K VLANs, while SPBM supports a 16 million virtual network service instances via its 24 bit I-SID field in the encapsulating header.
- SPBM can also support Layer 3 network virtualization as described in an IETF draft (IP/SPBM)
- SPB is compatible with IEEE 802.1ag and ITU Y.1731 OAM which means that existing management tools will work for SPB, while TRILL has yet to address OAM capability.
- SPB is compatible with Provider Backbone Bridging (PBB), the protocol used by many service providers to provide MPLS WAN services. This means that SPB traffic can be directly mapped to PBB. Also, virtual data centers defined with SPB can be mapped to separate traffic streams in PBB and given different QoS and security treatment.

In the future TRILL and SPB 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 either TRILL or SPB and network designs based on these technologies. 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 and SPB, the difference between access switches and core switches may shrink significantly.

As a result of TRILL or SPB, the switch topology may shift from a two-tier hub and spoke, such as the one in **Figure 6**, to a highly meshed or even fully meshed array of switches that appears to the attached devices as a single switch. TRILL and SPF bridging can support a variety of other topologies, including the fat tree switch topologies¹⁹ that are popular in cluster computing approaches to HPC. Fat trees have also gotten a lot of attention as a topology for highly scalable data center LANs, such as Cisco's FabricPath and Juniper's QFabric. Fat tree topologies are also used by Ethernet switch vendors to build high density, non-blocking 10 GbE switches using merchant silicon switch chips. This trend may eventually lead to the commoditization of the data plane aspect of Ethernet switch design. **Figure 7** 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 constructed²⁰. A number of high density 10 GbE switches currently on the market use this design approach.

¹⁹ www.mellanox.com/pdf/./IB_vs_Ethernet_Clustering_WP_100.pdf

²⁰ The maximum density switch that can be built with a two-tier fat tree architecture based on 24 port switch chips has 288 ports.

Figure 7: TOR Switch Fat Tree Internal Architecture



A discussion of the alternatives to STP amongst six of the primary data center LAN switch vendors can be found at Webtorials²¹.

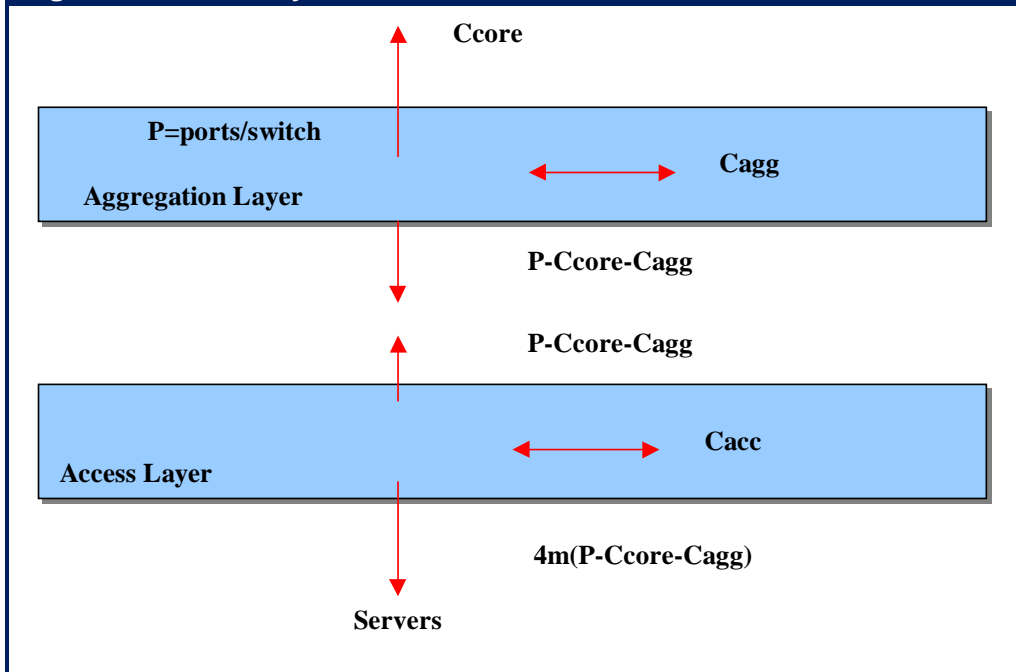
Scalability of Two Tier LAN Designs

The scalability of a LAN architecture is determined by the number of server ports that can be supported with a given level of redundancy and over-subscription at different points within the LAN topology. Many data center LANs being deployed today are based on a two tier design that provides high levels of redundancy and low over-subscription levels for server-to-server traffic. Two tier LAN designs are frequently implemented with Top of Rack (TOR) access switches in conjunction with chassis-based aggregation switches. The aggregation switches are connected to the LAN core and to the Internet, but all the server-to-server traffic within the data center flows only through the two tiers of access and aggregation switches.

Figure 8 shows a general model for two tier switched LANs that takes into account both connections for redundancy and connections to the LAN core. It is assumed that all servers are attached to the access/TOR switches via 10 GbE ports. Any inter-switch links at the access layer are assumed to be 10 GbE, and all other inter-switch links (i.e., inter-aggregation, access-to-aggregation and aggregation-to-core) are assumed to be 40 GbE. If a given model of switch does not yet support 40 GbE, a LAG with four 10 GbE member links could be substituted. It should be noted that as previously mentioned a 40 GbE link is preferable to a LAG of four 10 GbE links because having a single 40 GbE link avoids the issues that can occur when attempting to load balance traffic that consists of a small number of high volume flows.

²¹ <http://www.webtorials.com/content/tls.html>

Figure 8: Scalability Model for Two Tier Data Center LANs



Definition of Symbols

P: The number of 40 GbE ports per aggregation switch

m: The effective over-subscription ratio

S: The number of aggregation switches

Ccore: The number of 40 GbE ports per aggregation switch that are used to connect to the LAN core

Cagg: The number of 40 GbE ports per aggregation switch used to connect to other aggregation switches (e. g., for ISL/VSL). There may also be 10 GbE inter-switch links within the access/TOR tier to support virtual switch/router functions such as multi-chassis LAG (MLAG) or VRRP.

Cacc: The number of connections between TOR switches

P – Ccore – Cagg: The number of 40 GbE ports per aggregation switch available for connections to the access layer

4 x m x (P-Ccore-Cagg): The number of 10 GbE access layer ports that are available for server connection per aggregation

4 x S x m x (P-Ccore-Cagg): For two tier LAN design with multiple aggregation switches, the number of available server ports

This model can be applied equally well to two tier LANs based on MC LAGs and two tier fat trees. The model focuses on P, the number of 40 GbE ports per aggregation switch and the number of ports required to make connections both within and among network tiers.

In the model, *Ccore* is the number of 40 GbE ports per aggregation switch that are used to connect to the LAN core, *Cagg* is the number of 40 GbE ports per aggregation switch that are used to connect to other aggregation switches (e. g., for ISL/VSL). There may also be 10 GbE inter-switch links within the access/TOR tier to support virtual switch/router functions such as multi-chassis LAG (MLAG) or VRRP.

The access/TOR switches may be oversubscribed with more switch bandwidth allocated to server connections vs. the amount of bandwidth that is provided from the access tier to the aggregation tier. The over-subscription ratio is given by the following ratio:

The amount of bandwidth allocated to server access / The amount of bandwidth allocated to access-to-aggregation connectivity.

A typical high density TOR switch has 48 10 GbE ports for server connectivity and four 40 GbE ports for inter-switch connectivity. Where servers are single-attached to these TOR switches, m is equal to $(48 \times 10) / (4 \times 40) = 3$. Where the servers are dual-attached to a pair of TOR switches with active-passive redundancy, $m = 3$, but the effective over-subscription ratio is 1.5:1 because only one of the pair of server ports is active at any given time. Where the servers are dual-attached to a pair of TOR switches with active-active MC LAG redundancy, the requirement for inter-switch connections (C_{acc}) between the TOR switches means there are two fewer 10 GbE ports per TOR switch available for server connectivity and the over-subscription ratio is equal to $m = (46 \times 10) / (4 \times 40) = 2.88$

As shown in **Figure 8**, the number of 40 GbE ports per aggregation switch that is available for connections to the access layer is equal to $P \cdot C_{core} \cdot C_{agg}$ and the number of 10 GbE access layer ports that are available for server connection per aggregation is equal to $4 \times m \times (P \cdot C_{core} \cdot C_{agg})$. For a two tier LAN design with multiple aggregation switches, the number of available server ports is $4 \times S \times m \times (P \cdot C_{core} \cdot C_{agg})$, where S is the number of aggregation switches.

It should be noted that the model presented in **Figure 8** is based on having a single aggregation switch, and the factor S needs to be included to account for an aggregation tier with multiple aggregation switches. For an MC LAG 2 tier network S is generally limited to 2. For fat trees, the number of aggregation switches, or spine switches, is limited by the equal cost forwarding capabilities (16 paths is a typical limit), as well as the port density P . The port configuration of the access/TOR switch also imposes some limitations on the number of aggregation/spine switches that can be configured. For example, for a TOR switch with 48 10 GbE ports and four 40 GbE ports the number of 40 GbE aggregation switches is limited to four. Scaling beyond $S=4$, requires both a denser access switch with more 40 GbE ports and more 10 GbE port as well to maintain a desired maximum over-subscription ratio. The ultimate fat tree scalability is attained where the 10 GbE/40 GbE access switch has same switching capacity as the aggregation/spine switches.

With these caveats, the model takes into account redundancy and scalability for various Layer 2 and Layer 3 two-tier network designs as summarized in **Table 16**.

Table 16: Scalability of Two Tier 10/40 GbE Data Center LANs				
Parameter	2 Tier L2	2 Tier Layer 2 MC LAG	2 Tier Layer 2 Fat Tree	2 Tier Layer 3 Fat Tree
Redundancy	none	Full	full	Full
Ccore	variable	Variable	variable	variable
Cagg	0	ISL/VSL 2 per agg switch	0	0
Cacc	0	active/passive server access: 0 active/active: 2 per TOR	active/passive server access: 0 active/active: 2 per TOR	active/passive: 2 per TOR active/active: 2 per TOR
Max 10 GbE server ports	$4Sm(P \cdot C_{core} \cdot C_{agg})$ $S=1$	$4Sm(P \cdot C_{core} \cdot C_{agg})$ $S=2$	$4Sm(P \cdot C_{core} \cdot C_{agg})$; $S = \#$ of aggregation switches	$4Sm(P \cdot C_{core} \cdot C_{agg})$; $S = \#$ of aggregation switches
Scaling	Larger P, m	Larger P, m	Larger P, m, S	Larger P, m, S

As highlighted in **Table 16**, the only way that the scalability of the data center LAN can be increased is by increasing the:

- Number of aggregation switches
- Number of 40 GbE ports per aggregation switch
- Level of over-subscription

As stated earlier, a typical initial design process might start from identifying the required number of server ports, the required redundancy, and an upper limit on the over-subscription ratio. As shown in **Figure 9**, calculating the required number of 40 GbE ports per aggregation switch to meet these requirements is accomplished by inverting the scaling formula. An IT organization could use the following process to utilize the formula:

1. Determine required number of server ports
2. Select the desired network type from Table 4. This will determine Cagg
3. Select an access/TOR switch model. This together with the network type will determine Cacc and m.
4. Select the desired Ccore. This will determine over-subscription ratio for client/server traffic via the core
5. Calculate the required port density of the aggregation switch using the following formula:

Figure 9: Required Aggregation Switch Port Density

$$P=((\text{\# of server ports})/4Sm)+Ccore+Cagg$$

To exemplify the formula shown in **Figure 9**, consider the following network parameters:

The number of servers ports = 4512

Network type; MC LAG

m = 3

S = 2

Ccore = 2

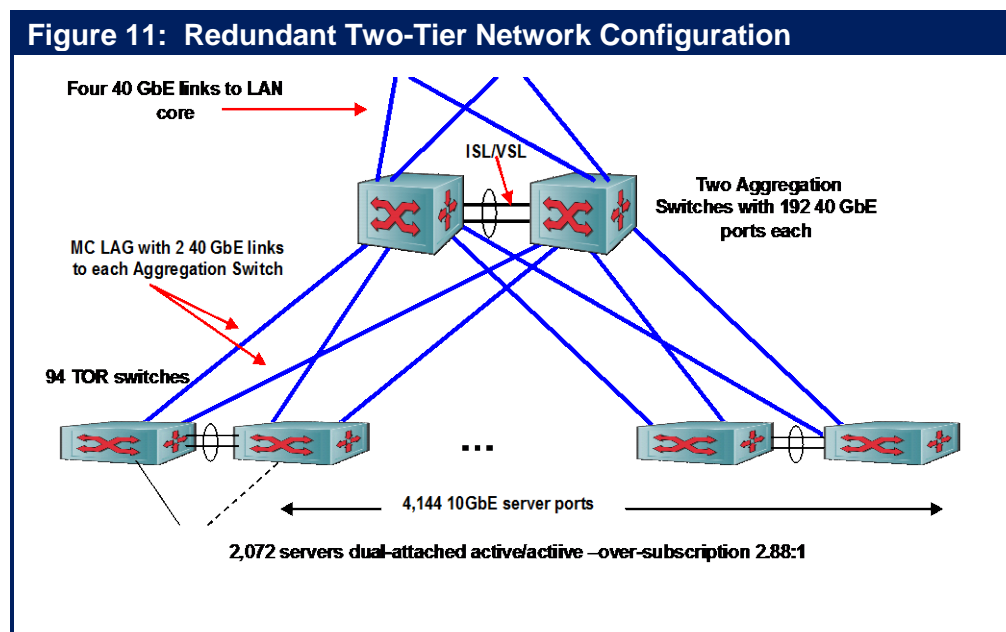
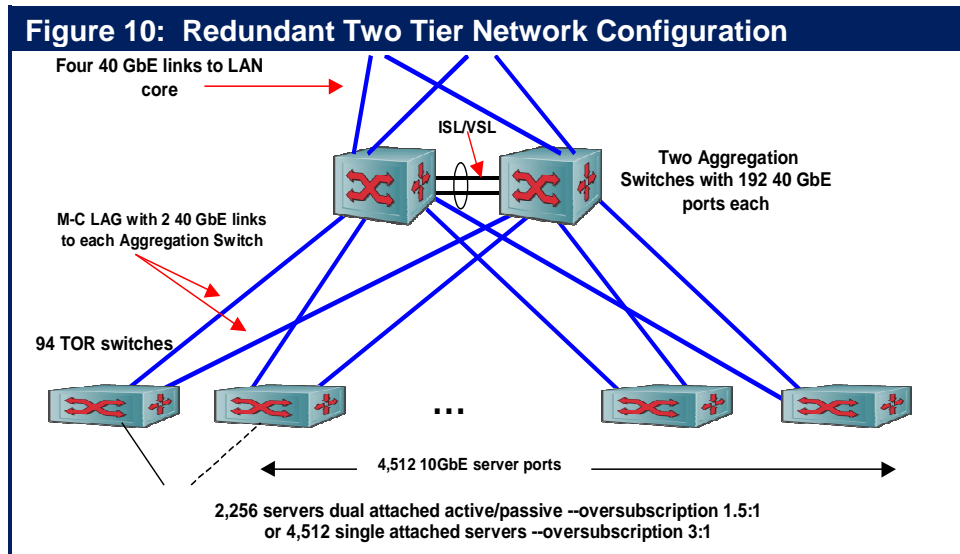
Cagg = 2

The formula in **Figure 9** indicates that in order to support the indicated network parameters, an aggregation switch with 192 40 GbE ports is required.

Figure 10 shows an example of a data center network that provides fully redundant Layer 2 server-to-server connectivity based on 94 TOR switches, each having 48 10 GbE ports and 4 40 GbE ports plus a pair of high density aggregation switches with 192 40 GbE ports each. The topology is an MC LAG Layer 2 network with oversubscribed TOR switches. Each of the 2,256 servers is connected to two TOR switches in an active/passive mode. The same configuration could also support 4,512 single-attached servers. With active/passive redundancy, the over-subscription of access switches for server-to-server traffic is 1.5:1.

For active-active server connectivity, each pair of TOR switches would need to be configured as a virtual switch with a pair of inter-TOR 10 GbE links for the ISL/VSL connectivity required for the virtual switch, as shown in **Figure 11**. This would reduce the number

of servers per TOR switch from 24 to 23 and the number of dual-attached servers to 2,072. With active/active redundant MLAG server connectivity, the over-subscription ratio for server-to-server traffic is 2.88:1.

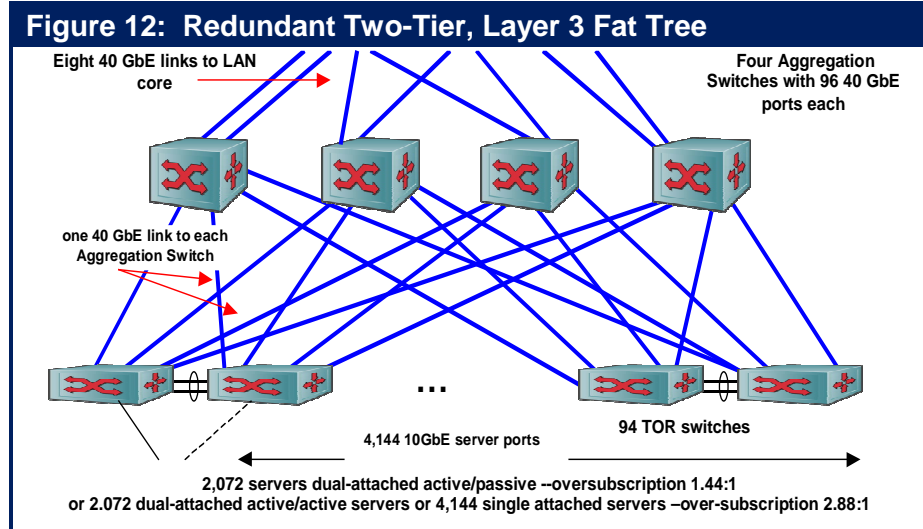


Building a comparable network with essentially the same number of 10 GbE server ports and similar over-subscription ratios using similar TOR switches and an aggregation switch with half the density (i.e., 96 40 GbE ports) requires some design

changes. Comparing the two designs provides an illustration of the effect that the density of the aggregation switch can have on the network design and the resulting TCO.

One possibility would be to build a Layer 2 fat tree network using four aggregation switches in the spine/aggregation layer and the same number of TOR switches (94) as the leaves/access switches. However, most TOR switches do not yet support Layer 2 equal cost multi-path forwarding alternatives other than with some form of MC LAG. One workaround is to move the Layer 3 boundary from the aggregation switch to the TOR switch and build a Layer 3 fat tree with OSPF ECMP providing the multi-path functionality. **Figure 12** shows what this could look like. Here the ISL links are only at the TOR level rather than the aggregation level and the

server connection can be made active/active without affecting the topology. With active/passive redundancy, the over-subscription of aggregation switches for server-to-server traffic is 1.44:1, while with active/active redundant server connectivity, the over-subscription ratio is 2.88:1. Note that Layer 2 and Layer 3 fat trees based on switches with the same port densities at the aggregation and access levels have the same physical topology.



If a TCO comparison is made of the two networks shown in **Figure 11** and **Figure 12**, some of the differences to consider are:

- Capex and Opex differences with four switches vs. two at the aggregation level, including switch cost, power capacity requirements, rack space requirements, annual power, annual routine administration, and annual service contract costs
- Difference in the number of server ports per TOR
- Differences in over-subscription ratios to the core
- Eight links vs. four links to the LAN core needed for redundancy
- Administrative cost and complexity differences with 98 Layer 3 devices if the fat tree is implemented at Layer 3 vs. two Layer 3 devices with MC LAG.

In addition, in a Layer 3 fat tree, there is a requirement for a Layer 2 over Layer 3 network virtualization to enable VM migration across Layer 3 boundaries

This example shows some of the complexities that can be encountered in comparing the TCOs of competing data center switching solutions that are based on switches of different port densities, as well as somewhat different functionality.

Network Support for Dynamic Creation and Movement of VMs

When VMs are migrated between servers, the network has to accommodate the constraints imposed by the VM migration utility; e.g., VMotion. Typically the VM needs to be on the same VLAN when migrated from source to destination server. This allows the VM to retain its IP address which helps to preserve user connectivity after the migration. When migrating VMs between disparate data centers, these constraints generally require that the data center Layer 2 LAN be extended across the physical locations or data centers without compromising the availability, resilience and security of the VM in its new location. VM migration also requires the LAN extension service have considerable bandwidth and low latency. VMware's VMotion, for example, requires at least 622 Mbps of bandwidth and less than 5 ms of round trip latency between source and destination servers over the extended LAN²².

The data storage location, including the boot device used by the virtual machine, must be accessible by both the source and destination physical servers at all times. If the servers are at two distinct locations and the data is replicated at the second site, the two data sets must be identical. One approach is to extend the SAN to the two sites and maintain a single data source. Another option is to migrate the data space associated with a virtual machine to the secondary storage location. In either case, there is a significant impact on the WAN.

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 was also noted earlier, the requirements for VM migration within VLAN boundaries have provided a major impetus for flattening the LAN with two-tier designs featuring Layer 2 connectivity end-to-end. Extending VLANs across the data center requires configuration of 802.1Q trunks between the intermediate switches, which can be a labor intensive task. With other forms of network virtualization (discussed in a later section of the report) virtual networks can be created without reconfiguration of intermediate switches.

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's 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 vendors²³. This type of solution is commonly referred to as Edge Virtualization.

²² <http://www.vce.com/pdf/solutions/vce-application-mobility-whitepaper.pdf>

²³ While this approach is the most common, some vendors have alternative approaches.

When a Virtual Machine is created or when the movement of a VM 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 can 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 proprietary form of VM network profile software, including linking their switches to at least one brand of hypervisor. Some differences exist between the range of hypervisors supported and the APIs that are 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.

Another approach to edge virtualization is the Distributed Virtual Switch (DVS). With DVS, the control and data planes of the embedded hypervisor vSwitch are decoupled. This allows the data planes of multiple vSwitches to be controlled by an external centralized management system that implements the control plane functionality. Decoupling the data plane from the control plane also 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, in addition to providing edge virtualization in support of VM creation and mobility.

The DVS is a significant improvement over earlier hypervisor vSwitches, but retains a number of characteristics of vSwitches that may be of concern to network designers, including:

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.

IEEE 802.1Qbg is a standard that addresses both edge virtualization and some of the potential issues with vSwitches. The standard includes Edge Virtual Bridging (EVB) in which all the traffic from VMs is sent to the physical network access switch. If the traffic is destined for a VM on the same physical server, the access switch returns the packets to the server over the same port on which it was received. The shipping of traffic from a VM inside of a physical server to an external access switch and then back to a VM inside the same physical server is often referred to as a hair pin turn or reflective relay. With Edge Virtual Bridging, the hypervisor can be relieved from all switching functions, which are now performed by the physical access network. With EVB, the vSwitch can perform the simpler function of a Virtual Ethernet Port Aggregator (VEPA) 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 802.1Qbg standard includes some additional protocols that standardize the switch side of edge virtualization. The additional protocols Edge TLV Protocol and VSI Discovery and Configuration Protocol (VDP) support edge virtualization where the Layer 2 configuration of the network to support VM creation and migration is automated. Using VDP, the target switch can be informed of the imminent VM deployment, allowing the target switch to be properly configured in advance of VM creation or movement. Therefore, Qbg provides a standards-based alternative to proprietary approaches to edge virtualization via integration between switch management systems and hypervisor management systems. A companion effort, the IEEE's 802.1BR Bridge Port Extension is defining a technique for a single physical port to support a number of logical ports and a tagged approach to deal with frame replication issues. Port Extension is used in fabric extenders for blade servers and rack mounted servers as an alternative to blade server switches and full function ToR switches.

Vendors of data center switches are expected to provide some level of support for 802.1Qbg. Some vendors may focus on either EVB or edge virtualization, while others will support the full range of Qbg capabilities. Some vendors may also offer DVS implementations that support Qbg-based edge virtualization.

Network Virtualization

Within the IT industry, the phrase *network virtualization* is used in a wide variety of ways. In order to eliminate confusion and ambiguity, The Survey Respondents were told that "Network virtualization is the creation of multiple logical networks that share a common physical network in a manner that is somewhat analogous to how multiple virtual machines share a common physical server. While techniques such as VLANs have been available for a long time, emerging technologies such as VXLAN, NVGRE and Software Defined Networks are enabling new forms of network virtualization."

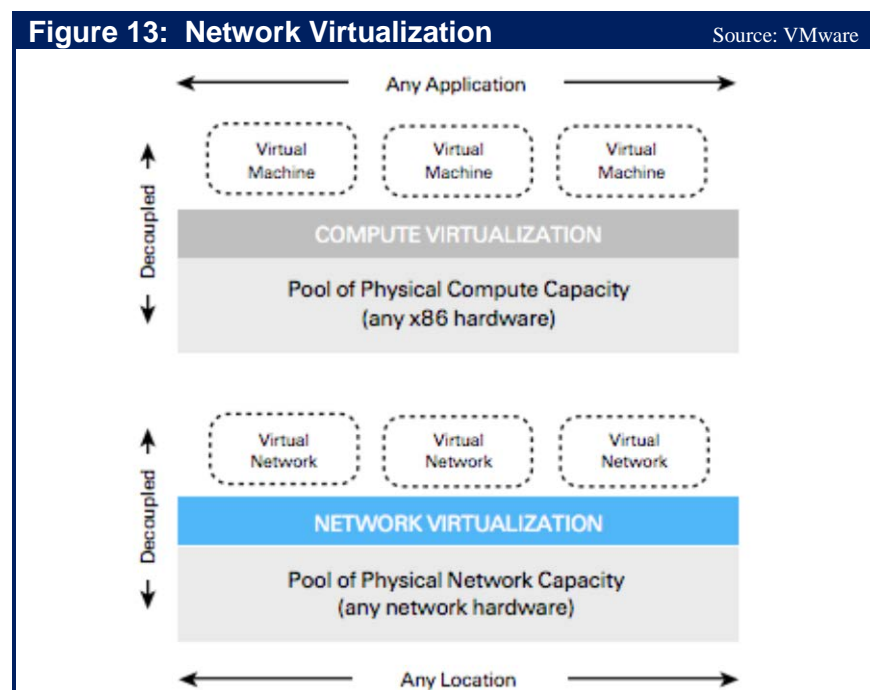
The Survey Respondents were then given a set of possible actions and were asked to indicate which of the actions best describes their organizations approach to these new forms of network virtualization. Their responses are shown in **Table 17**.

Table 17: Status of Network Virtualization		N = 307
Action	Percentage of The Survey Respondents	
We have already implemented network virtualization	23.8%	
We are interested in network virtualization, but we will not likely take any steps towards implementing it for at least a year	17.6%	
We are in the process of evaluating network virtualization	16.3%	
We are not currently taking any steps towards implementing network virtualization, but are likely to in the next twelve months	12.7%	
Don't know	10.1%	
We currently have no interest in network virtualization	9.1%	
We are in the process of testing network virtualization	8.8%	
Other	1.5%	

One conclusion that can be drawn from the data in **Table 17** is that:

There is very strong interest on the part of IT organizations to implement network virtualization.

In addition to 802.1Qbg there are a number of emerging and proposed standard protocols that are focused on optimizing the support that data center Ethernet LANs provide for server virtualization. Several of these protocols are aimed at network virtualization via the creation of multiple virtual Ethernet networks that can share a common physical infrastructure in a manner that is somewhat analogous to multiple VMs sharing a common physical server, as shown in **Figure 13**.



Most protocols for network virtualization are based on creating virtual network overlays using tunneling/encapsulation techniques. The protocols that provide network virtualization of the data center include VXLAN, NVGRE, STT, and SPB MAC-in-MAC. SPB is already a IEEE standard, while it is likely that only one of the other proposals will achieve IETF standard status.

Traditional Network Virtualization

One-to-many virtualization of network entities is not a new concept. The most common traditional applications of the virtualization concept to networks are VLANs and Virtual Routing and Forwarding (VRF) instances.

VLANs partition the Ethernet network into as many as 4,094 broadcast domains as designated by a 12 bit VLAN ID tag in the Ethernet header. VLANs have been a convenient means of isolating different types of traffic that share the same switched LAN infrastructure. In data centers making extensive use of server virtualization, the limited number of VLANs can present problems, especially in cases where a large number of tenants need to be supported, each of whom requires multiple VLANs. Extending VLANs across the data center via 802.1Q trunks to support VM mobility adds operational cost and complexity. In data centers based on Layer 2 server-to-server connectivity, large numbers of VMs, each with its own MAC address, can also place a burden on the forwarding tables capacities of Layer 2 switches.

VRF is a form of Layer 3 network virtualization in which a physical router supports multiple virtual router instances, each running its own routing protocol instance and maintaining its own forwarding table. Unlike VLANs, VRF do not use a tag in the packet header to designate the specific VRF to which a packet belongs. The appropriate VRF is derived at each hop based on the incoming interface and information in the frame. An additional requirement is that each intermediate router on the end-to-end path followed by a packet needs to be configured with a VRF instance that can forward that packet.

Network Virtualization with Overlays

Due to the shortcomings of the traditional VLAN or VRF models, a number of new techniques for creating virtual networks have emerged over recent years and months. Most of these network virtualization techniques are based on tunneling/encapsulation to construct multiple virtual network topologies overlaid on a common physical network. A virtual network can be a Layer 2 network or a Layer 3 network, while the physical network can be Layer 2, Layer 3 or a combination depending on the overlay technology. With overlays, the outer (encapsulating) header includes a field (generally 24 bits wide) that carries a virtual network instance ID (VNID) that specifies the virtual network designated to forward the packet.

Virtual network overlays can provide a wide range of benefits, including:

- Support for essentially unlimited numbers of virtual networks (24 bits equates to 16 million virtual networks)
- Decoupling of the virtual network topology, service category (L2 or L3), and addressing from those of the physical network. The decoupling avoids issues such as MAC table size in physical switches.
- Support for virtual machine mobility independent of the physical network. If a VM changes location, even to a new subnet, the switches at the edge of the overlay simply update their mapping tables to reflect the new location of the VM. The network for a new VM can be provisioned entirely at the edge of the network.

- Ability to manage overlapping IP addresses between multiple tenants.
- Support for multi-path forwarding within virtual networks

The main difference between the various overlay protocols lies in their encapsulation formats and the control plane functionality that allows ingress (encapsulating) devices to map a frame to the appropriate egress (decapsulating) device.

VXLAN

Virtual eXtensible LAN (VXLAN)²⁴ virtualizes the network by creating a Layer 2 overlay on a Layer 3 network via MAC-in-UDP encapsulation. The VXLAN segment is a Layer 3 construct that replaces the VLAN as the mechanism that segments the data center LAN for VMs. Therefore, a VM can only communicate or migrate within a VXLAN segment. The VXLAN segment has a 24 bit VXLAN Network identifier. VXLAN is transparent to the VM, which still communicates using MAC addresses. The VXLAN encapsulation is performed through a function known as the VXLAN Tunnel End Point (VTEP), typically a hypervisor switch or a possibly a physical access switch. The encapsulation allows Layer 2 communications with any end points that are within the same VXLAN segment even if these end points are in a different IP subnet. This allows live migrations to transcend Layer 3 boundaries. Since MAC frames are encapsulated within IP packets, there is no need for the individual Layer 2 switches to learn MAC addresses. This alleviates MAC table hardware capacity issues on these switches. Overlapping IP and MAC addresses are handled by the VXLAN ID, which acts as a qualifier/identifier for the specific VXLAN segment within which those addresses are valid. The VXLAN control solution uses flooding based on Any Source Multicast (ASM) to disseminate end system location information.

As noted, VXLANs uses a MAC-in-UDP encapsulation. One of the reasons for this is that modern Layer 3 devices parse the 5-tuple (including Layer 4 source and destination ports). While VXLAN uses a well-known destination UDP port, the source UDP port can be any value. As a result, a VTEP can spread all the flows from a single VM across many UDP source ports. This allows for efficient load balancing across LAGs and intermediate multi-pathing fabrics even in the case of multiple flows between only two VMs.

Where VXLAN nodes on a VXLAN overlay network need to communicate with nodes on a legacy (i.e., VLAN) portion of the network, a VXLAN gateway can be used to perform the required tunnel termination functions including encapsulation/decapsulation. The gateway functionality could be implemented in either hardware or software.

VXLAN is the subject of a IETF draft supported by VMware, Cisco, Arista Networks, Broadcom, Red Hat and Citrix. VXLAN is also supported by IBM. Pre-standard implementations in hypervisor vSwitches and physical switches are beginning to emerge.

NVGRE

Network Virtualization using Generic Router Encapsulation (NVGRE) uses the GRE tunneling protocol defined by RFC 2784 and RFC 2890. NVGRE is similar in most respects to VXLAN with two major exceptions. While GRE encapsulation is not new, most network devices do not

²⁴ <http://searchservvirtualization.techtarget.com/news/2240074318/VMware-Cisco-propose-VXLAN-for-VM-mobility>

parse GRE headers in hardware, which may lead to performance issues and issues with 5-tuple hashes for traffic distribution in multi-path data center LANs. The other exception is that the current IETF NVGRE draft does not address the control plane question, leaving that for a future draft or possibly as something to be addressed by (Software Defined Networking) SDN controllers. Some of the sponsors of NVGRE (Microsoft and Emulex) expect that some of the performance issues can be addressed by intelligent NICs that offload NVGRE endpoint processing from the hypervisor vSwitch. The intelligent NICs would also have API interfaces for integration with overlay controllers and hypervisor management systems. Emulex has also demoed intelligent NICs that offload VXLAN processing from the VMware Distributed Switches.

STT

Stateless Transport Tunneling (STT) is a third overlay technology for creating Layer 2 virtual networks over a Layer 2/3 physical network within the data center. Conceptually, there are a number of similarities between VXLAN and STT. The tunnel endpoints are typically provided by hypervisor vSwitches, the VNID is 24 bits wide, and the transport source header is manipulated to take advantage of multipathing. STT encapsulation differs from NVGRE and VXLAN in two ways. First, it uses a stateless TCP-like header inside the IP header which allows tunnel endpoints within end systems to take advantage of TCP segmentation offload (TSO) capabilities of existing TOE server NICs. The benefits to the host include lower CPU utilization and higher utilization of 10 GbE access links. STT generates a source port number based on hashing the header fields of the inner packet to ensure efficient load balancing over LAGs and multi-pathing fabrics. STT also allocates more header space to the per-packet metadata, which provides added flexibility for the virtual network control plane. With these features, STT is optimized for hypervisor vSwitches as the encapsulation/decapsulation tunnel endpoints.

The STT IETF draft sponsored by Nicira does not specify a control plane solution. However, the Nicira network virtualization solution includes OpenFlow-like hypervisor vSwitches and a control plane based on a centralized network virtualization controller that facilitates management of virtual networks.

Shortest Path Bridging MAC-in-MAC (SPBM)

IEEE 802.1aq SPBM uses IEEE 802.1ah MAC-in-MAC encapsulation and the IS-IS routing protocol to provide Layer 2 network virtualization and VLAN extension in addition to the loop-free equal cost multi-path Layer 2 forwarding functionality normally associated with SPB. VLAN extension is enabled by the 24 bit Virtual Service Network (VSN) Instance Service IDs (I-SID) that are part of the outer MAC encapsulation. Unlike other network virtualization solutions, no changes are required in the hypervisor vSwitches or NICs and switching hardware already exists that supports IEEE 802.1ah MAC-in-MAC encapsulation. For SPBM, the control plane is provided by the IS-IS routing protocol.

SPBM can also be extended to support Layer 3 forwarding and Layer 3 virtualization as described in the IP/SPB IETF draft using IP encapsulated in the outer SPBM MAC. This draft specifies how SPBM nodes can perform Inter-ISID or inter-VLAN routing. In addition, IP/SPB also provides for Layer 3 VSNs by extending Virtual Routing and Forwarding (VRF) instances at the edge of the network across the SPBM network without requiring that the core switches also support VRF instances. VLAN-extension VSNs and VRF-extension VSNs can run in parallel on the same SPB network to provide isolation of both Layer 2 and Layer 3 traffic for multi-tenant environments. With SPBM, all the core switches (starting at the access or aggregation switches

that define the SPBM boundary) need to be SPBM-capable. SPBM hardware switches are currently available from Avaya, Huawei, and Alcatel-Lucent.

A discussion of network virtualization would not be complete without at least a mention of two Cisco protocols: Overlay Transport Virtualization (OTV) and Locator/ID Separation Protocol (LISP). OTV is optimized for inter-data center VLAN extension over the WAN or Internet using MAC-in-IP encapsulation. It prevents flooding of unknown destinations across the WAN by advertising MAC address reachability using IS-IS routing protocol extensions. LISP is an encapsulating IP-in-IP technology that allows end systems to keep their IP address (ID) even as they move to a different subnet within the network (Location). By using LISP VM-Mobility, IP endpoints such as VMs can be relocated anywhere regardless of their IP addresses while maintaining direct path routing of client traffic. LISP also supports multi-tenant environments with Layer 3 virtual networks created by mapping VRFs to LISP instance-IDs.

Another way to implement network virtualization is by implementing a Software Defined Network (SDN). SDN is the subject of a subsequent section of The Report.

Network Convergence and Fabric Unification

In contrast to Second Generation Data Center LANs:

A possible 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 a reduction 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 such as TCP to manage congestion and to recover lost packets through re-transmissions. In order to emulate the lossless behavior of a Fibre Channel (FC) SAN, Ethernet needs enhanced flow control mechanisms that eliminate buffer overflows for high priority traffic flows, such as storage access flows. Lossless Ethernet is based on 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 be 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

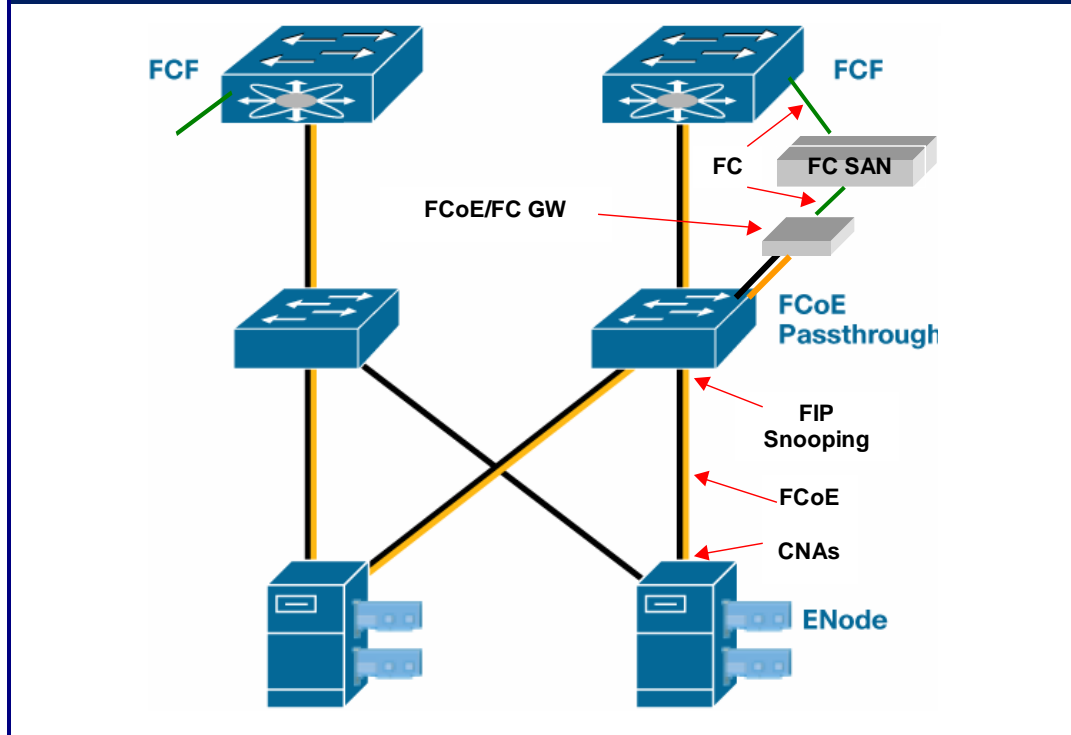
cascaded Layer 2 switches, such as those typically found in larger data centers for server interconnect, cluster interconnect and to support extensive SAN fabrics.

- **IEEE 802.1Qaz Enhanced Transmission Selection (ETS)** specifies 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. **The Data Center Bridging Exchange (DCBX)** protocol is also defined in the 802.1Qaz standard. The 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.

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 2/4/8/16 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.

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 preserves 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 (e.g., 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 14](#).



As shown in **Figure 14**, End Nodes (servers) don't 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 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 don't 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-only vendors are considerably more enthusiastic about iSCSI SANs over 10/40/100 GbE than they are about FCoE.

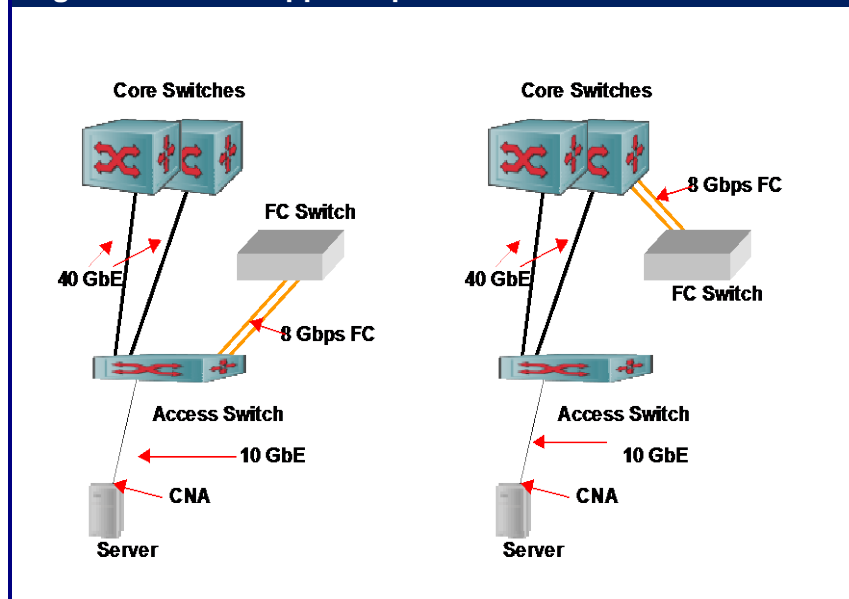
The primary drivers of FCoE are the vendors that offer both Ethernet and FC products.

These are the vendors that are already shipping 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 15**.

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.

Figure 15: FCF Support Options



Security Services in Virtualized Data Centers

As pointed out in the first section of The Report, security is generally considered by enterprise IT departments to be the primary concern in today's highly virtualized data centers and in the implementation of private or public cloud computing environments. In the traditional data center, internal security has generally been implemented by deploying dedicated physical security appliances at the Aggregation layer of a 3-tier or 2-tier network. This reduces the number of physical devices required and allows firewalls to filter traffic flowing from one access

VLAN to another. This approach has been successful in relatively static non-virtualized environments that require infrequent changes to the location and configuration of both servers and physical security appliances. This traditional model does not address inspection of inter-VM traffic within a single physical server.

With the advent of server virtualization and the dynamic migration of workloads within and between data centers, there is a growing need to make the workload's complete security environment as easily provisioned and migrated as the VMs themselves. In addition to being dynamic and virtualization-aware, the security solution needs to be both scalable and automated to the degree possible.

For enterprise data centers and Private Cloud Networking, the prevalent traffic isolation solution has been to make extensive uses of VLANs to isolate VMs performing different workloads or different aspects of the workload (e.g., web, application, and database tiers). In addition, firewalls, Intrusion Prevention Systems and other security appliances are generally required to filter and monitor inter-VM and inter-VLAN traffic in order to provide an additional layer of security for critical workloads and data resources.

In multi-tenant environments, it is highly desirable to be able to secure traffic within the tenant network as well as firewalling traffic at the tenant edge. The problem is most significant in highly virtualized IaaS data centers where a physical server may host VMs from multiple clients. In order to meet the demand for highly dynamic provisioning of resources IaaS service providers will focus on maximizing the use of virtual security appliances rather than physical devices. Traffic isolation in multi-tenant environments will be increasingly based on network virtualization based on either overlays or OpenFlow or possibly a combination of these techniques.

One approach for securing highly virtualized server environments is to use virtual security appliances on the same servers as the virtualized applications. Virtual appliances can be dynamically provisioned and migrated along with application VMs. Some virtual security appliances can support multiple security functions in a single VM. A virtual security appliance integrated with the hypervisor vNICs can provide security services for all the VMs on a host, inspecting both inter-VM traffic and traffic from external sources. Where the virtual security appliance also supports routing functionality, it can also inspect inter-VLAN traffic on the same host. When the VMs and the virtual security appliances are on separate VLANs and on separate hosts, traffic between them is typically switched at the Layer 3 tier of the physical network (typically at the aggregation layer). This means that a significant volume of security traffic may have to make a rather inefficient round trip through the physical network even if the application VM and the virtual security appliance are in the same POD or even on the same physical server (i.e., where the virtual security appliance doesn't support routing).

A second approach, more applicable in enterprise data centers because it does not involve virtual appliances on the servers, is to deploy a virtualized physical security appliance that can support a large number of instances of virtual security devices, such as firewalls, IDS/IPS, WAF, etc. Potentially, these instances could be implemented as VMs running on the security device's hypervisor. This type of integrated security device can also include its own physical Layer 2 and Layer 3 switching functionality, which allows the device to be installed in line between the access and aggregation layers of the physical data center LAN. The VLANs used by the virtualized servers are trunked to the virtualized security appliance via the hypervisor vSwitches and the physical access switches. There are a number of benefits of the integrated virtualized security appliance including:

- Specialized or dedicated hardware support for a number of security functions
- Ability to flexibly serialize different security services (firewall, IPS, etc) without having to change switch configurations or install additional physical security appliances
- Support for dynamic changes to security configurations for traffic among VLANs
- Ability to switch inter-POD security traffic without involving the aggregation layer switches

With the advent of DVSs and Layer 2 network virtualization using overlays, network partitioning can be based on virtual Ethernet overlay networks rather than simple VLANs. This vastly increases the number of virtual networks that can co-exist in the enterprise or multi-tenant IaaS data center and provides support for overlapping IP addresses among multiple tenants. Also, because a virtual network can span Layer 3 boundaries, VMs on the same physical server can communicate with each other across subnet boundaries via the DVS without involving Layer 3 switching in upstream physical devices. This can optimize securing local communication between co-resident VMs running different applications on separate subnets or VMs accessing the security services provided by co-resident virtual appliances on separate subnets. The overlay tunnels eliminate the need for inline security services and makes it possible to direct traffic to security services provided by virtual or physical security devices anywhere in the network.

As noted earlier, another potential approach to network virtualization is based on OpenFlow. The OpenFlow network can potentially be partitioned into multiple virtual networks based on certain characteristics of the 12-tuple used to differentiate flows. Each of the OpenFlow virtual networks can have its own independent OpenFlow controller, providing isolation of virtual networks at the control plane as well as the data plane. OpenFlow also provides a high degree of flexibility where the controller can direct flows to either physical security devices or virtual security appliances. It is also possible that the OpenFlow controller itself would provide some of the security services required.

Summary of Third Generation Data Center LAN Technologies

The data center LAN is still in the throes of rather dramatic technology developments, summarized in **Table 18**. As shown in the table, a number of standards have been completed in the last year or so, creating the expectation that more products supporting these standards will be announced in the near future.

Table 18: Status of Data Center Technology Evolution	
Technology Development	Status
Two-tier networks with Layer 2 connectivity extending VLANs across the data center.	On-going deployment
Standardized edge virtualization automating Layer 2 configuration for VM creation and mobility. Possible changing role for the hypervisor vSwitch as a port aggregator (VEPA) for EVB, potentially eliminating the vSwitch tier.	The 802.1Qbg standard is in place and some implementations are available.
Reduced role for blade switches to eliminate switch tier proliferation.	On-going with proprietary fabric extenders. Work on the IEEE802.1BR standard is in progress

Table 18: Status of Data Center Technology Evolution

Technology Development	Status
Multi-chassis LAG and switch virtualization technology to address STP issues and provide active-active redundant server connectivity.	On-going deployment
Multi-core servers with notably more VMs per server and 10 GbE connectivity to the LAN.	Adoption stage.
40 GbE and 100 GbE uplinks and core switches.	A standard has been in place for some time: 40 GbE is becoming widely available on access and core switches 100 GbE is becoming available. But adopted primarily by service providers due to economic considerations
TRILL enabling new Layer 2 data center LAN topologies; e.g., fully meshed, fat tree with equal cost multi-path forwarding	The TRILL standard RFC 6325 has also been approved. Enhancement being proposed to IETF. Pre-standard switch implementations of TRILL with proprietary extensions are available. No standard TRILL yet.
SPB enabling new Layer 2 data center LAN topologies; e.g., fully meshed, fat tree with equal cost multi-path forwarding	SPB (IEEE 802.1aq) has been finalized and switch products are available.
SPB Network Virtualization	Layer 2 virtualization covered in IEEE 802.1aq. Products are available. Layer 3 virtualization is the subject of an Internet draft and implemented by Avaya in its SPB switches
VXLAN Network Virtualization	A draft was recently submitted to the IETF. Pre-standard implementations are available in vSwitches and some access switches
NVGRE and STT Network Virtualization	Drafts were recently submitted to the IETF. STT is implemented by Nicira
SDN	Vendors are beginning to offer SDN solutions based on OpenFlow. ONF standards are limited to OpenFlow
OpenFlow	OF V1.0 hybrid switches and controllers are available from multiple vendors OF V1.3 spec has been released
DCB delivering lossless Ethernet for 10 GbE and higher speed Ethernet	Standards are in place. Switches with DCB are available.
10 GbE FCoE approach to fabric unification	FCoE standard is in place and products are available
10 GbE iSCSI approach to fabric unification	Early implementations
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.

Software Defined Networking (SDN)

In the current environment vendors tend to have different definitions of SDN. The three most common ways that vendors use the phrase *software defined networks* are discussed below.

1. Programmability of switch control planes whether or not the control plane is segregated and centralized

This approach to SDN is based on having direct programmatic interfaces into network devices, which are broadly defined to include all L2 - L7 functionality. In this approach, the control and forwarding planes are not separated, nor is the control plane centralized. Providing direct programmatic interfaces into networking devices is not new, as multiple vendors have supported this functionality for several years.

One advantage of this approach is that it enables very detailed access into, and control over, network elements. However, it doesn't provide a central point of control and is vendor specific. While some network service providers may adopt the approach of directly accessing network platforms, it is unlikely to gain much traction in the enterprise market in at least the near term.

2. Distributed Virtual Switching with segregation of control and data planes

In this approach to SDN the control and forwarding planes are separated. This approach is based on leveraging a virtual switch (vSwitch) and having the vSwitch function as a forwarding engine that is programmed by a device that is separate from the vSwitch. This functionality is used as part of an overlay network that rides on top of the existing network infrastructure using protocols such as VXLAN or NVGRE. As was the case with the approach to SDN discussed above, multiple vendors have supported this approach to SDN for several years.

3. An architecture similar to the one shown in [Figure 16](#)

This is the most common way that vendors define SDN. Based on this definition, SDN is positioned as an emerging network paradigm that is based on multiple levels of abstraction. These levels of abstraction allow network services to be defined, programmatically implemented, and managed centrally without requiring network operations personnel to interface directly with the control and management planes of each individual network element that is involved in delivering the service. Instead, the SDN operator can deal with a pool of devices as a single entity.

There are a couple of important options for how the architecture shown in [Figure 16](#) could be implemented. One key option is the protocol that is used to communicate between the switch and the controller. The most commonly discussed such protocol is OpenFlow, which is described in a subsequent section of this document. Alternative ways to communicate between the controller and the switch include the Extensible Messaging and Presence Protocol, the Network Configuration Protocol and OpenStack. The other key option is the amount of intelligence in the switch. In one alternative, referred to as a pure SDN switch, the intelligence in the SDN switch is limited to just

what is needed for data plane packet forwarding. In the other alternative, referred to as a hybrid SDN switch, some of the traditional control plane functionality may be centralized and the remaining functionality remains distributed within switches. Depending upon how much of the control functionality is centralized, this scenario may not result in switches with significantly less functionality and in fact may result in switches that require additional functionality.

Unless specifically mentioned, throughout the rest of this publication the definition of SDN that will be used is the third one in the preceding list. In addition, unless specifically mentioned, it will be assumed that OpenFlow is used to communicate between the controller and the switch and that the only intelligence in the switch is just what is needed for data plane packet forwarding.

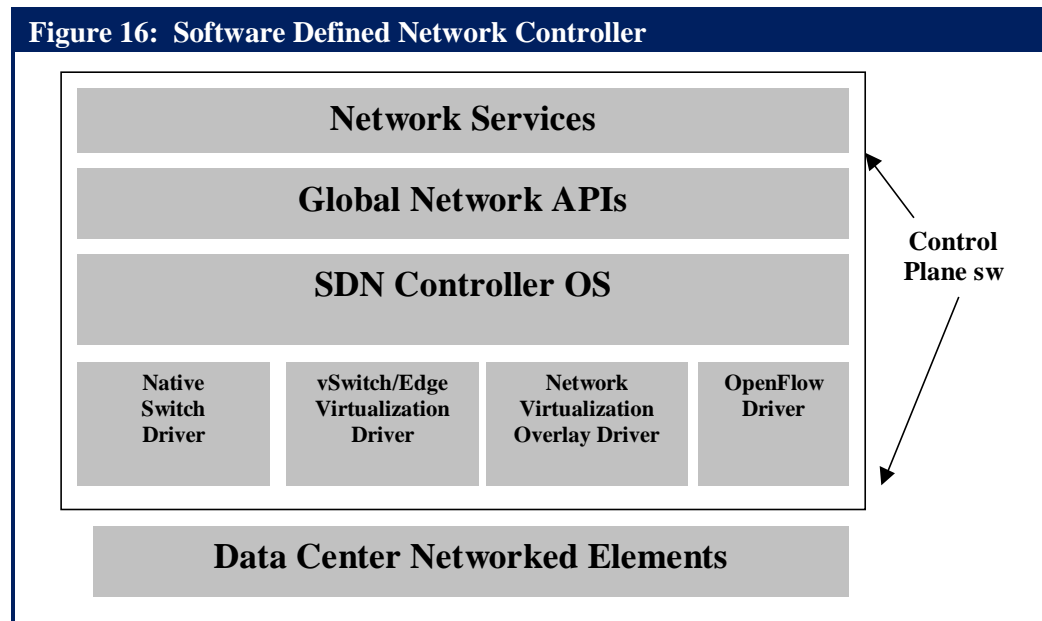
With this definition of SDN, network flows are controlled at the level of the global network abstraction with the aid of the OpenFlow protocol, rather than at the level of the individual devices. Global control of the network is achieved by logical centralization of the control plane function. Based on these characteristics, a well-designed SDN offers the potential advantage of greatly improved flexibility, highly reduced operational complexity, and a high degree of agility in responding to dynamic changes in the demand for network resources.

Another aspect of SDN that is of interest for cloud computing is the automated provisioning of networks as a complement to the automated provisioning of servers and storage. An SDN can provide this capability via interfaces with cloud controller orchestration software, such as the open source OpenStack controller and its "Quantum" virtual network interface.

Most of the networking industry that supports the SDN movement believes that SDNs should be based on industry standards and open source code to the degree possible. The open development model is the preferred model for timely adoption of new SDN standards that support multi-vendor interoperability and the creation of a large ecosystem of vendors providing a range of SDN components and functionality needed to span a variety of SDN use cases.

The SDN Network Architecture

A layered architecture for SDN is shown in **Figure 16**. In **Figure 16**, the control plane function is centralized in SDN Controller software that is installed on a server or on a redundant cluster of servers for higher availability and performance.



Below is a description of the primary components of the network model in **Figure 16**.

- **Network Services**
These are written to a set of Global Network APIs provided by the SDN Controller's operating system (OS). Network Services might include SAN services, Security services, Multi-tenant services, and Multi-path load balancing services provided by the SDN Controller vendor, as well as other services provided by an eco-system of ISVs and third parties writing applications to a set of published APIs.
- **The SDN Controller's Operating System**
This supports a number of drivers that distribute state in order to control the behavior of the underlying network elements so that the network will provide the desired network services. Below is an overview of these lower level control elements.
- **Virtual Switch /Edge Virtualization Drivers**
These enable SDNs to address some of the special networking requirements imposed by server virtualization, including control of the edge virtualization capabilities of hypervisor-based distributed virtual switches (DVSs) and/or access switches. With standards-based edge virtualization both the hypervisor DVS and the access switch can support the IEEE 802.1Qbg standard²⁵, which enables edge virtual bridging.

²⁵ <http://www.ieee802.org/1/pages/802.1bg.html>

- **Network Virtualization Overlay Drivers**

These interface with edge switches to provide network virtualization by overlaying a virtual Layer 2 Ethernet network over a Layer 2/ Layer 3 physical network. The overlay is generally implemented using some form of encapsulation/ tunneling that may be performed by an SDN controlled vSwitch, virtual appliance, or physical access switch.

- **OpenFlow Networking Drivers**

These interface with OpenFlow-enabled switches.

At the present time, there are a number of OpenFlow switches and SDN controllers available in the marketplace. In addition, a number of vendors, including controller vendors, switch vendors and application delivery controller vendors, have announced network services that are layered on the controller.

Open Networking Foundation

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

As part of their stewardship of the OpenFlow protocol, in March 2012 the ONF sponsored an interoperability event that was open to all of the members of the ONF²⁶. A total of fourteen companies and two research institutions participated in the event which focused on the OpenFlow v1.0 standard. According to the ONF, the majority of its members have implemented v1.0. The ONF has also stated that many of its members are not going to implement v1.1 but will move forward and implement v1.2 and v1.3.

The interoperability event tested the following capabilities:

1. Discovering the network using the Link Layer Discovery Protocol (LLDP)
2. Dynamically provisioning point-to-point Layer 2 paths across an OpenFlow network
3. Learning the Layer 3 (IP) network and responding to a failed link
4. Performing load balancing on flows
5. Slicing the network with FlowVisor, which is a special purpose OpenFlow controller that acts as a transparent proxy between OpenFlow switches and multiple OpenFlow controllers

Additional information on the testing and the lessons learned can be found at ONF Interoperability Event White Paper²⁷.

One of the criticisms of the ONF is that it is focused just on OpenFlow-based SDNs and that as previously mentioned in this report; there are other ways to implement an SDN. While there is some validity to that criticism, one of the other approaches to implementing an SDN, providing direct access to switches and routers, is by its nature vendor specific and hence not subject to standardization by the ONF or any other organization. The other approach, the use of vSwitches and overlay networks, encroaches on the domain of the IETF, which is currently working on overlay protocols including VXLAN and NVGRE.

Another criticism is that the ONF has been too focused on enabling L2 and L3 functionality and has had too little focus on enabling L4 – L7 functionality. There is also some validity to that criticism. However, the success of either developing or adopting a new technology is predicated in part on being able to have a broad enough scope so that the technology does indeed add significant value, but not so broad as to cause undue delay or organizational barriers. For example, a network organization that is considering implementing SDN could advocate that by so doing, it would improve L2 and L3 networking functions and would also significantly improve L4 – L7 functions such as load balancing and security. The problem with taking that broad of an

²⁶ The ONF sponsored a similar event in October 2012, the results of which were not made public prior to the publication of this document.

²⁷ <https://www.opennetworking.org/membership/onf-documents>

approach to SDN deployment is that it will likely mean that multiple groups within the IT organization would all have to agree to the deployment of SDN and that a level of consensus would have to be reached relative to how it would be deployed. In most IT organizations, getting the participation and buy-in from multiple groups prior to the deployment of SDN would result in a significant delay in the implementation of the technology. An approach that is more likely to succeed is for the networking organization to implement SDN for purely networking reasons and hence not need the approval and buy-in of other groups in the IT organization. Then, at some appropriate time in the future, the network organization can encourage other IT groups to leverage their SDN deployment. A related consideration is that over time the deployment of SDN may encourage significant changes in the roles, culture and structure of IT organization. However, in the vast majority of cases, any approach to SDN deployment that requires significant changes in the roles, culture and structure of IT organization prior to implementation is DOA. Similar to the need for network organizations to focus initially on L2 and L3 functionality, if the ONF had adopted too broad of a focus early on, it ran the risk of making little if any progress.

In August 2012 the ONF announced four new initiatives that have less of a focus on OpenFlow than has been typical of past ONF initiatives. These four new initiatives, which are described below, have the potential to significantly accelerate SDN adoption. These new initiatives are:

1. Architecture and Framework

This initiative will look at upper layer orchestration of the network with the goal of exposing the various interfaces and elements of an SDN and identifying how these interfaces and elements relate both to each other and to legacy networking. To the degree that this initiative is successful, it will mitigate one of the challenges that is associated with the adoption of SDN, which is how to integrate an SDN into an existing production network. This initiative is also intended to develop what the ONF refers to as “network solution elements”, which refers to entities such as APIs and data models, and to enable these network solution elements to “work well together”. While the ONF did not define what they meant by “work well together”, the goal is to foster greater automation of the network and reduce the amount of manual tasks that are currently required.

2. New Transport

The ONF new transport initiative is intended to accelerate the deployment of OpenFlow and SDN in carrier networks, optical networks, and wireless networks by defining the requirements and use cases necessary to deploy SDN. According to the ONF, the initiative will investigate how to use OpenFlow and switches not just between Ethernet ports, but also between fibers, wavelengths, wireless channels and circuits. The goal of this initiative is for network operators and users to gain both economies of scale and more system-wide consistency in applying policy and security across a broader reach.

3. Northbound API

This initiative will survey and catalog the APIs that exist, define how to characterize them, outline what they are intended to be used for, and how they interact with the network. The ONF stated their belief that cataloging and characterizing the APIs will offer a clear understanding of what functions the market views as important and the common thread for application scenarios. They also stated their belief that this work will aid software developers to better program and virtualize the network, and enable network operators to translate network capabilities into lucrative services.

4. **Forwarding Abstractions**

The forwarding abstractions initiative will focus on the development of next generation forwarding plane models, with a particular interest in terms of how to exploit and differentiate the capabilities of OpenFlow based hardware switches. The ONF stated their belief that one of the key benefits of SDN is the ability to take advantage of merchant silicon to drive better price and performance in the data center. The ONF also stated their belief that this initiative will foster a competitive marketplace for high performance hardware that meets the needs of demanding customers and that network operators, including enterprises, will be able to reap the benefits of OpenFlow in the core of their networks, not just the edge.

OpenFlow

OpenFlow is an open protocol between a central SDN/OpenFlow controller and an OpenFlow switch that can be used to program the forwarding behavior of the switch. Using pure OpenFlow switches, a single central controller can program all the physical and virtual switches in the network. All of the control functions of a traditional switch (e.g. routing protocols that are used to build forwarding information bases (FIBs)) are run in the central controller. As a result, the switching functionality of the OpenFlow switch is restricted entirely to the data plane,

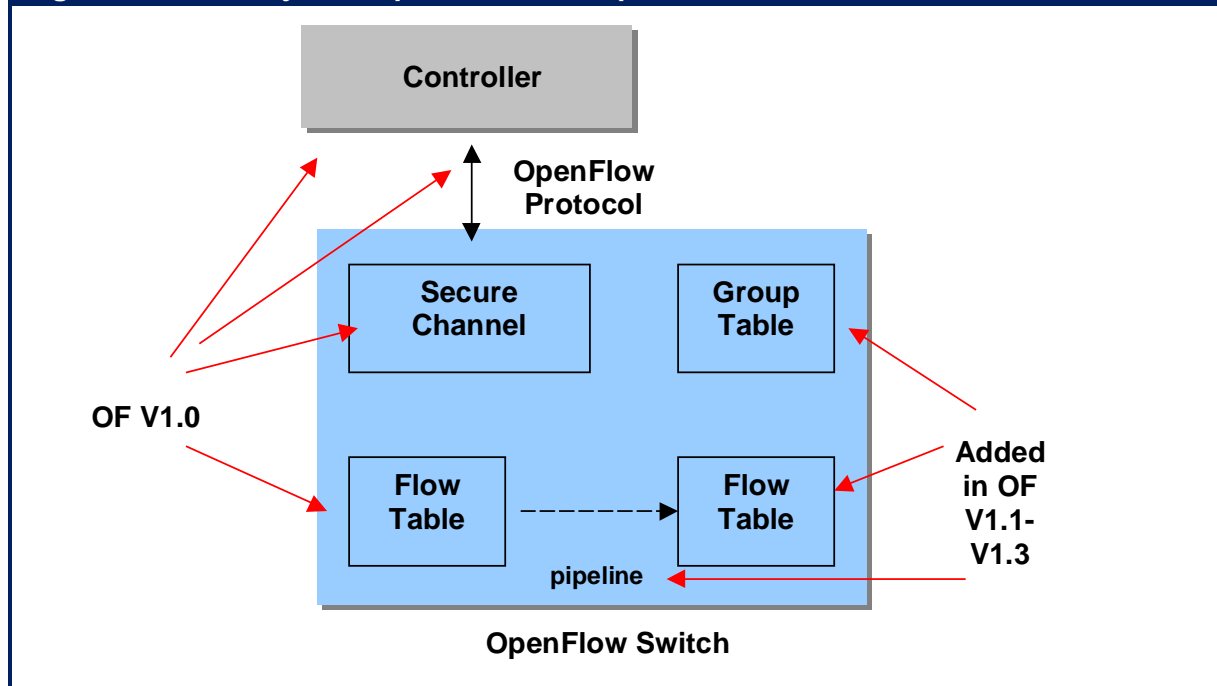
Most modern Ethernet switches and routers contain flow-tables, typically supported by TCAMs that run at line-rate to perform forwarding functions based on Layer 2, 3, and 4 packet headers. While each vendor's flow-table is different, there is a common set of functions supported by a wide variety of switches and routers. It is this common set of functions that is exploited by the OpenFlow protocol.

Many existing high functionality Layer 2/3 switches can be converted to be OpenFlow-hybrid switches by the relatively simple addition of an OpenFlow agent in firmware supported by the native switch Network Operating System (NOS). As previously discussed, an alternative to adapting an existing switch to support OpenFlow would be to build an OpenFlow-only switch that, by definition, is dedicated to supporting only OpenFlow forwarding. In theory at least, an OpenFlow-only switch would be extremely simple and inexpensive to build because it would have very little resident software and would not require a powerful CPU or large memory to support the extensive control functionality typically packaged in a traditional network operating system (NOS). The ability to build a highly scalable, low cost, OpenFlow-only switch is currently limited by the ability of the merchant silicon vendors to supply the necessary functionality. That is a large part of the motivation for the previously discussed ONF initiative on forwarding abstractions.

The basic elements of an OpenFlow V1.0 network are shown on the left hand side of Figure 2. Most existing Open Flow Switches have been built to the V1.0 spec (12/2009). This spec has been enhanced three times in V1.1 (2/2011), V1.2 (12/2011), and V1.3 (6/2012) to add functionality including additional components as indicated on the right hand side of the figure.

As shown in **Figure 17**, the central controller communicates with the switch's OpenFlow agent over a secure TLS channel. This channel could be either in-band or out-of-band. The OpenFlow agent on the switch populates the flow table as directed by the controller.

Figure 17: The Major Components of an OpenFlow Switch V1.0-V1.3



The data path of an OpenFlow V1.0 switch is comprised of two entities. One entity is a single Flow Table that includes the rules for matching flows to table entries. The second entity consists of counters that record the number of packets and bytes received per flow and other port and table statistics. **Figure 18** shows the 12-tuple of header fields that are used to match flows in the flow table.

Figure 18: The OpenFlow V1.0 Flow Table Fields

Ingress Port	Ether Src	Ether Dest	Ether Type	VLAN ID	VLAN Prior	IP Src	IP Dest	IP Proto	IP TOS	Src Port	Dest Port
--------------	-----------	------------	------------	---------	------------	--------	---------	----------	--------	----------	-----------

OpenFlow switches are required to support two basic types of actions: Forward and Drop. Forwarding is either directed to a physical port or to one of the following virtual ports:

- ALL: Send the packet out all interfaces, not including the incoming interface.
- CONTROLLER: Encapsulate and send the packet to the controller.
- LOCAL: Send the packet to the switch's local networking stack.
- TABLE: Perform actions in the flow table. Applies for only packet-out messages.
- IN PORT: Send the packet out the input port.

For OpenFlow V1.0 there are also a number of optional/recommended actions:

- NORMAL: Process the packet using the traditional forwarding path supported by the switch (for OpenFlow-hybrid switches)
- FLOOD: Flood the packet along the spanning tree.
- ENQUEUE: Forward a packet through a specific port queue to provide QoS.

- **MODIFY FIELD:** Change the content of header fields, including set VLAN ID and priority, strip VLAN, modify Ethernet or IPV4 source and destination addresses, modify IPV4 TOS, modify transport source and destination ports.

When a packet arrives at the OpenFlow V1.0 switch, the header fields are compared to flow table entries. If a match is found, the packet is either forwarded to specified port(s) or dropped depending on the action stored in the flow table. When an OpenFlow Switch receives a packet that does not match the flow table entries, it encapsulates the packet and sends it to the controller. The controller then decides how the packet should be handled and notifies the switch to either drop the packet or make a new entry in the flow table to support the new flow.

Over the last year and a half extensive enhancements have been made to the OpenFlow specification under of the auspices of the Open Networking Foundation. A complete listing of the enhancements included in OpenFlow V1.1-V1.3 is beyond the scope of this document. However, some of the major changes include:

- Additional components of a flow entry in the flow table. In addition to the match fields, the following fields are included in the entry:
 - **PRIORITY:** matching precedence of the flow entry
 - **COUNTERS:** to update for matching packets
 - **INSTRUCTIONS:** to modify the action set or pipeline processing
 - **TIMEOUTS:** maximum amount of time or idle time before flow expiration
 - **COOKIE:** opaque data value chosen and used by the controller to process flows
- Flexible pipeline processing through multiple flow tables, as shown in the right hand side of Figure 2. As a packet is processed through the pipeline, it is associated with a set of accumulating actions and metadata. The action set is resolved and applied at the end of the pipeline. The metadata allows a limited amount of state to be passed down the pipeline.
- The new group table abstraction and group action enable OpenFlow to represent a set of ports as a single entity for forwarding packets. Different types of groups are provided, to represent different forwarding abstractions, such as multicasting or multi-pathing.
- Improved tag handling includes support for Q-in-Q plus adding, modifying and removing VLAN headers and MPLS shim headers.
- Support for virtual ports, which can represent complex forwarding abstractions such as LAGs or tunnels.
- OpenFlow Extensible Match (OXM) uses a TLV (Type Link Value) structure to give a unique type to each header field increasing the flexibility of the match process.
- Basic support for IPv6 match and header rewrite has been added, via OXM.
- Support for multiple controllers to improve reliability.

Potential Benefits of OpenFlow

There are a number of possible ways for the control centralization, programmability, and flow forwarding characteristics of OpenFlow to be exploited by innovative users and vendors of network devices and software. This includes:

- **Centralized FIB**

One of the primary benefits of OpenFlow is the centralized nature of the Forwarding Information Base (FIB). Centralization allows optimum routes to be calculated deterministically for each flow leveraging a complete model of the end-to-end topology of the network. This model can be built using a discovery protocol, such as the Link Layer Discovery Protocol (LLDP). Based on an understanding of the service levels required for each type of flow, the centralized OpenFlow controller can apply traffic engineering principles to ensure each flow is properly serviced. The result can be much better utilization of the network without sacrificing service quality. Centralized route processing also allows the pre-computation of a set of fail-over routes for each possible link or node failure. Centralized processing also can take advantage of virtually unlimited processing power or multi-core processors and cluster computing for calculating routes and processing new flows.

- **The [Google G-Scale WAN](#) Backbone**

This is the WAN that links Google's various global data centers. As is mentioned below, the most common discussion of implementing SDN focuses on the data center. However, the G-Scale WAN is a prime example of a production OpenFlow Layer 3 WAN that is realizing the benefits of FIB centralization. The G-Scale control plane is based on BGP and IS-to-IS and the OpenFlow-only switches are very simple 128 port 10 GbE switches that were built by Google using merchant silicon. It is important to note that when Google built these switches, 128 port 10 GbE switches had not yet been introduced in the commercial market. The Google G-Scale WAN is discussed in more detail in the next section of The Report.

- **OpenFlow Virtual Networking**

As described in a preceding section of The Report, there are a number of approaches to network virtualization including simple VLANs and network overlays based on various MAC-in-MAC, MAC-in-IP or UDP encapsulations. Future versions of OpenFlow specs will undoubtedly support standards-based overlays. In the interim, OpenFlow can potentially provide another type of virtualization for isolating network traffic based on segregating flows. One very simple way to do this is to isolate sets of MAC addresses without relying on VLANs by adding a filtering layer to the OpenFlow controller. This type of functionality is available in v0.85 of the Floodlight controller. Floodlight's [VirtualNetworkFilter](#) module also implements the OpenStack Quantum API. This provides the option of automatically provisioning OpenFlow virtual networks from the OpenStack cloud management system in conjunction with provisioning virtual servers and storage resources via the OpenStack Nova and Swift capabilities.

- **OpenFlow Multi-Pathing**

Most networking vendors offer data center fabric solutions featuring some form of Layer 2 multi-pathing to improve the network's capacity to handle "east-west" traffic flow which is characteristic of server virtualization, converged storage networking, and cluster computing. OpenFlow offers another approach to multi-pathing that does not rely on

standards such as TRILL or SPB. As noted earlier, the OpenFlow Controller (OFC) can use LLDP to discover the entire network topology via discovering switches and switch adjacencies. Using this topological model, the OFC can compute all the parallel physical paths, including paths that share some network nodes and other paths that are entirely disjoint - and therefore offer higher reliability. The OFC can then assign each flow across the network fabric to a specific path and configure the OpenFlow switches' flow tables accordingly. The OFC can then offer shared and disjoint multi-pathing as network services that can be delivered to applications. With appropriate processing power, the OFC can support very large-scale networks and high availability via path redundancy and fast convergence following link or node failures.

- **OpenFlow Firewalls and Load Balancers**

By virtue of Layer 2-4 flow matching capability OpenFlow access switches can perform filtering of packets as they enter the network, acting as simple firewalls at the edge. With OpenFlow switches that support modification of packet headers, the SDN/OF Controller will also be able to have the switch redirect certain suspicious traffic flows to higher-layer security controls, such as IDS/IPS systems, application firewalls, and Data Loss Prevention (DLP) devices. Another possible security application of OpenFlow would be in Network Access Control (NAC).

OpenFlow with packet header modification will also allow the switch to function as a simple, cost-effective load-balancing device. With modification functionality, a new flow can result in a new flow table entry that includes an action to modify the destination MAC and IP addresses. The modified address can be used to direct traffic to the server selected by the controller.

The Marketplace Reality

In July and August of 2012, Ashton, Metzler & Associates and Information Week conducted extensive market research into SDN. This included a survey that was completed by 250 qualified Information Week subscribers. It also included interviews that were conducted with both enterprise IT organizations as well as with vendors. This sub-section of The Report will discuss some of the key findings of that market research.

One key finding was that:

Most enterprise IT organizations have little if any knowledge of SDN.

That conclusion follows because over a third of the 393 IT professionals who received the screener for the SDN survey indicated that they had no familiarity with SDN and roughly half of the respondents who did have some familiarity with SDN indicated that they were only somewhat familiar with it.

Of the Information Week Respondents who were familiar with SDN, there was a high degree of familiarity with OpenFlow. However, in spite of the fact that, as previously mentioned, it is possible to implement an SDN and not use OpenFlow:

The vast majority of IT organizations believe that OpenFlow is an important component of an SDN.

The fact that OpenFlow is perceived as being so important to SDN could be another indication that the overall awareness of what SDN somewhat lags the reality. Alternatively, it could reflect a feeling on the part of IT organizations that while there are other ways to create an SDN, that OpenFlow provides distinct advantages that they deem to be critical.

Relative to the question of whether or not SDN switches will be just dumb forwarding engines or more highly functional hybrid SDN switches, the Information Week Respondents were asked “Do you believe that SDN will relegate switches and routers to being just relatively dumb forwarding engines?” They were given three possible answers: Yes; No; Don’t Know. The 250 responses were almost equally split across the three answers.

There is not a consensus amongst IT organizations about whether or not SDN will relegate switches and routers to be just dumb forwarding engines.

While SDN can be applied in a variety of places within the network, including the WAN, most of the current discussion of SDN focuses on implementing SDN in the data center LAN. With that in mind, the Information Week Respondents were given a set of fourteen challenges that are associated with data center LANs. They were asked to indicate which three challenges they thought SDN would be most helpful in resolving. Their responses are shown in [Table 19](#).

Table 19: LAN Challenges Mitigated by SDN	
Challenge	Percentage of Respondents
Improve network utilization and efficiency	42%
Automate more provisioning and management	35%
Improve security	32%
Implement network-wide policies	31%
Reduce cost	29%
Get more visibility into applications that are using the network	25%
Reduce complexity	23%
Increase scalability	20%
Reduce reliance on proprietary protocols or proprietary extensions of standards-based protocols	12%
Support creation of a private or hybrid cloud	10%
Support creation and dynamic movement of virtual machines	8%
Reduce reliance on vendor's product life cycles	4%
Support more east-west traffic	1%
Other	1%
Source: Information Week and AM&A	

The top five rows in **Table 19** demonstrate that:

IT organizations believe the primary value that SDN offers in the data center is that it can help IT organizations to reduce costs, automate management, and enforce security policies.

When discussing SDN, it is common for the trade press and industry analysts to talk about the ability of an SDN to better support the adoption of private and/or hybrid cloud computing. The data in **Table 19** indicates that that capability is not currently a strong driver of enterprise adoption of SDN.

It is common to have technology adoption driven by different factors at different points in the adoption cycle. For example, the initial driver of server virtualization was cost savings. However, once IT organizations began to implement server virtualization, most of them found that the agility that virtualized servers provided became as important to them as the cost savings. In similar fashion, IT organizations may well implement a SDN initially for cost savings or added security and later expand that implementation because it provides other capabilities, such as making it easier to support cloud computing.

As previously mentioned, a number of vendors, including controller vendors, switch vendors and application delivery controller vendors, have announced network services that are layered on the controller. Those network services include:

- Network virtualization
- Load balancing
- Firewalls
- DDOS prevention
- Traffic engineering
- Disaster recovery
- Application acceleration via techniques such as SSL offload
- Web optimization
- Network analysis whereby management data is filtered from network elements and sent to a central site for analysis.

In the near term, SDN applications will come primarily from current infrastructure players. While infrastructure players will likely continue to develop SDN applications:

One of the key promises of SDN is that developer communities will be created and that these communities will develop a wide range of applications.

While cost savings can drive the adoption of technology or new ways of implementing technology, a key factor that needs to be considered is how those changes impact security. The Information Week Respondents were asked about the impact of SDN on security. Their answers indicated that only a small minority of IT organizations thinks that implementing SDN will make networks less secure. In contrast:

The majority of IT organizations believe that implementing SDN will make networks more secure.

A previous section discussed some of the ways that SDN could provide more security functionality; e.g., by providing simple firewalls at the edge of the network. The primary ways that The Information Week Respondents believe that SDN will increase security is that it will:

- Make it easier to apply a unified security policy
- Make it easier to encrypt data
- Enable access control that is more granular and more integrated
- Provide additional points where security controls can be placed
- Make it easier to inspect and firewall VM to VM traffic on the same physical server

In order to understand the resistance to implementing SDN, the Information Week Respondents were given a set of fourteen potential impediments to SDN adoption. They were asked to indicate which the three top impediments to their company adopting SDN in the next two years. Their responses are shown in [Table 20](#).

Table 20: Inhibitors to SDN Deployment	
Challenge	Percentage of Respondents
Immaturity of current products	41%
Confusion and lack of definition in terms of vendor's strategies	32%
Immaturity of enabling technologies	25%
Other technology or business priorities	24%
Lack of resources to evaluate SDN	23%
Concern that the technology will not scale to support enterprise-class networks	22%
Worry that the cost to implement will exceed ROI	18%
We don't see a compelling value proposition	18%
Lack of a critical mass of organizations that have deployed SDN	14%
Concern that major networking vendors will derail SDN by adding proprietary features	13%
Not scheduled to have a technology refresh in that time frame	11%
No inhibitors to implementing SDN	4%
We've already implemented SDN	2%
Other	2%
Source: Information Week and AM&A	

The data in **Table 20** demonstrates that:

The primary inhibitor to SDN adoption is the overall confusion in the market and the immaturity of products and vendor strategies.

The Information Week Respondents were asked when they expected to have SDN in production. Their answers are shown in **Table 21**.

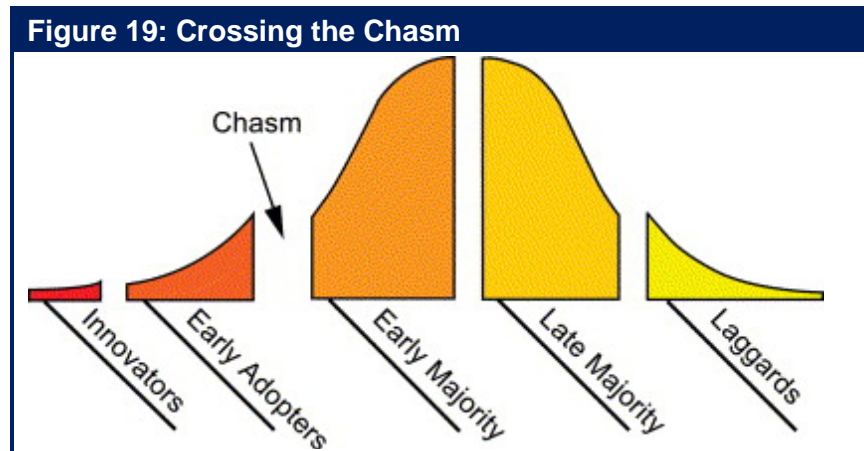
Table 21: SDN Production Timeline	
Timeframe for Production	Percentage of Respondents
SDN in production now	4%
Less than six months	5%
Six to twelve months	9%
More than twelve months but less than twenty four months	17%
More than twenty four months	11%
No plans to implement SDN	37%
Don't know	17%
Source: Information Week and AM&A	

As shown in **Table 21**, currently 4% of IT organizations have SDN already in production networks and an additional 14% expect that they will within a year. If that data is completely accurate, then 18% of IT organizations will have SDN in production within the next year. However, survey data about the planned deployment of technology is seldom completely accurate. For example, an IT organization that indicates that it has no plans to implement a new technology in the next year is more likely accurate than one that says they do. That follows because if the IT organization has not yet started the planning and lined up the resources to test and implement the technology, it is highly unlikely that they will be able to turn that around and implement the technology in the next six to twelve months. However, a company may have every intention of trialing and implementing a new technology in the next six to twelve months, but priorities can change in that time frame. As a result, it is highly likely that somewhat less than 18% of IT organizations will have implemented SDN in a production network within the next year.

Crossing the Chasm

In 1991 Geoffrey Moore wrote *Crossing the Chasm: Marketing and Selling High-Tech Products to Mainstream Customers*. In the book, Moore argues that there is a chasm (Figure 19) between the early adopters of a technology and the early majority of pragmatists and that these two groups approach the adoption of technology very differently. For example, the early

adopters of a technology are typically the organizations who identify the primary use cases of a technology and who have both the capability and the orientation to work through the issues that are associated with implementing early stage technologies. In contrast, the early majority typically adopts a technology once the use cases have been identified and validated and once the solutions are stable. While there is a chasm, or discontinuity, between the early adopters and the mainstream adopters, there is typically a continuum of risks and rewards that then separates the early majority from the late majority from the laggards.



At the current point in time, SDN is appropriate only for early adopters. The market research previously presented indicates that 18% of IT organizations intend to have SDN in production within a year. Some market adoption studies²⁸ indicate that the innovators and early adopters are roughly 16% of the total number of companies. Hence, if that market research is close to 100% accurate, then one could argue that SDN will cross the chasm and become a mainstream approach to networking in roughly a year. However, as was previously discussed, survey respondents tend to be optimistic relative to their adoption of technology. In addition, below is a list of factors that will influence the rate of adoption of SDN and hence will either increase or decrease the amount of time it will take for SDN to cross the chasm and become a mainstream approach to networking.

- The development and validation of compelling use cases.
- The stability of the OpenFlow protocol.
- The stability of the north bound APIs.
- Broad interoperability of products.
- The creation of an application developer community.
- The development of strong partnerships amongst members of the SDN ecosystem.
- Ongoing mergers and acquisitions.
- The lack of a major issue such as the inability of SDN solutions to scale or a major security incident that was the result of deploying SDN.

The bottom line is that SDN will likely cross the chasm in the next year or two.

²⁸ <http://www.zonalatina.com/Zldata99.htm>

A Plan for SDN

Given that there is a high probability that SDN will have a major positive impact on networking, IT organizations need to break through the cloud of confusion that surrounds SDN in order to better understand it and to establish an SDN strategy – even if that strategy ends up being that the IT organization decides to do nothing relative to SDN for the foreseeable future. Some of the components of that strategy are:

- A firm definition of what SDN means to the organization. This includes taking a position relative to whether or not they want to implement an SDN that features:
 - The direct programmability of switches and routers, which in most cases will be accomplished by leveraging software created by a third party.
 - The separation of the control and forwarding planes and use OpenFlow for communications between them.
 - The separation of the control and forwarding planes and use something other than OpenFlow for communications between them.
 - An overlay network.
 - Other approaches and technologies.
- The use cases that justify deploying SDN, whether that is to solve problems or to add value. Included in this component of the strategy is an analysis of alternative ways to solve those problems or add that value and the recognition that the use cases may change over time.
- An ongoing analysis of the progress that SDN is making relative to crossing the chasm. This includes analyzing the items mentioned in the preceding section; e.g., the stability of OpenFlow and of the northbound APIs.
- The identification of how extensive the implementation of SDN will be both initially and over the first couple of years of deployment. For example, will the implementation just include top of rack switches or will it also include some core switches? Will it include L4 – L7 functionality, such as load balancing or protection against DOS attacks?
- A decision on whether any of the control functions that have historically been done in switches and routers will be done in SDN controllers.
- An analysis of how the deployment of SDN fits in with both the existing infrastructure as well as with other IT initiatives that are in progress.
- An analysis of the SDN strategies and offerings of various vendors and the identification of one or more viable SDN designs. This includes an analysis of the risks and rewards of acquiring pieces of the SDN from disparate vendors vs. trying to acquire all or most of the solution from a single vendor.

- The identification as to whether or not the IT organization will write applications itself to take advantage of SDN and if so, what has to happen within the organization to enable that capability.
- The identification and analysis of the commercially available applications that take advantage of SDN.
- An evaluation of the availability and scalability characteristics of the particular SDN designs that are under consideration.
- An analysis how the IT organization can provide a sufficient level of security for the controllers.
- Assuming that the IT organization is interested in OpenFlow: An analysis of whether to implement OpenFlow only switches or hybrid switches that support OpenFlow and traditional networking.
- The identification of how the IT organization will manage and troubleshoot their SDN deployment.
- An evaluation of the publicly available reports on interoperability testing.
- A plan for testing the SDN designs and use cases that are under consideration.
- An analysis of how the intended implementation of SDN would impact the current networks.
- A plan for how the IT organization will minimize and mitigate the risks that are associated with implementing SDN.
- A program for getting management buy-in. This includes getting funding as well as the buy-in from any other organization that will be directly impacted by the deployment of SDN.

The Wide Area Network (WAN)

Introduction

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. While the early use of the Internet was strictly for academic and research purposes, the use of the Internet for commercial purposes started in the early 1990s with the development of the World Wide Web.

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 or private WAN technologies²⁹. 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 that were discussed in a preceding section of The Report.

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

Today there is not a fundamentally new generation of technology under development that is focused on the WAN.

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)³⁰. As described later in this section of the report, 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.

²⁹ An enterprise or private 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.

³⁰ <http://vlt.me/vpls-0810>

Contrasting the LAN and the WAN

The WAN is notably different than the data center LAN. These differences include the fact that:

- After a lengthy period in which there was little if any fundamental innovation, the LAN is experiencing broad fundamental change. In contrast, after a lengthy period in which the WAN underwent repeated fundamental change, there are currently no fundamentally new WAN specific technologies under development.
- While there are no fundamentally new WAN specific technologies under development, there are new WAN architectures being developed and implemented.
- In the vast majority of instances, the latency, jitter and packet loss that the LAN exhibits doesn't have an appreciable impact on application performance. In many instances, the latency, jitter and packet loss that public and private WANs exhibit has an appreciable impact on application performance. This is particularly true of 3G/4G networks.
- One of the primary design criteria for designing a data center LAN is scalability. A manifestation of the ongoing improvements in LAN scalability is that over the last fifteen years the speed of a data center LAN has increased from 10 Mbps to 10 Gbps – which is a factor of a thousand. In contrast, in many cases the primary design criterion for designing a WAN is to minimize cost. For example, in many parts of the world it is possible to get high-speed WAN links such as an OC-192 link. These links, however, are usually not affordable.
- The LAN follows Moore's Law. In contrast, the price/performance of enterprise WAN services such as MPLS doesn't come close to doubling every two years.

The WAN doesn't follow Moore's Law.

WAN Budgets

Both in 2011 and again in 2012, The Webtorials Respondents were asked how their budget for the forthcoming year for all WAN services compares to what it is in the current year. Their responses are contained in **Table 22**.

Table 22: WAN Budget Increases		
	Responses in 2011	Responses in 2012
REDUCED BY MORE THAN 10%	3.2%	7.2%
Reduced by 1% to 10%	11.1%	18.5%
Basically static	34.9%	40.5%
Increased by 1% to 10%	32.8%	21.2%
Increased by more than 10%	18.0%	12.6%

The change in the budget for WAN services in 2012 is notably different than what the corresponding change was in 2011. For example, in 2011 half of the **Survey Respondents** indicated that their WAN budget was increasing while in 2012, only a third did.

WAN budgets are notably more constrained than they were a year ago.

As is explained in the next subsection, the adoption of cloud computing will increase the rate of growth in the amount of traffic that transits the WAN. As such:

IT organizations need to make changes relative to how they use WAN services in order to support a significant increase in WAN traffic while experiencing a highly constrained WAN budget.

Drivers of Change

As mentioned in the section of this report entitled [The Emergence of Cloud Computing and Cloud Networking](#), one of the characteristics of cloud computing is increased reliance on the network. The increased reliance on the WAN in particular stems from the fact that the resources that support cloud computing solutions are centralized in a small number of data centers and the vast majority of users access these solutions over the WAN. Hence, the more use that organizations make of cloud computing in general, the more traffic transits both the Internet and enterprise WAN services. When looking just at public or hybrid cloud services, the adoption of these services primarily results in more traffic transiting the Internet.

Below are some of the specific factors that are putting more traffic onto the WAN and hence, driving the need for IT organizations to change their approach to wide area networking.

Virtual Machine Migration

The section of this report entitled *The Emerging Data Center LAN* quantified the great interest that IT organizations have in server virtualization in general and in moving virtual machines (VMs) between data centers in particular. That section of the report also discussed the fact that one of the requirements associated with moving VMs between data centers is that the data storage location, including the boot device used by the VM being migrated, must be accessible

by both the source and destination physical servers at all times. If the servers are at two distinct locations and the data is replicated at the second site, the two data sets must be identical. One approach to enabling data access is to extend the SAN to the two sites and to maintain a single data source. Another option is migrate the data along with the VM to the secondary site. In either case, it is necessary to coordinate VM and storage migrations and to be able to move large data sets efficiently between data centers, which will have a significant impact on the WAN.

Virtual Desktops

Another form of virtualization that will drive a further increase in WAN traffic is desktop virtualization. In order to quantify the interest that IT organizations have in desktop virtualization, The **Survey Respondents** were asked to indicate the percentage of their company's desktops that have either already been virtualized or that they expected would be virtualized within the next year. Their responses are shown in **Table 23**.

Table 23: Deployment of Virtualized Desktops					
	None	1% to 25%	26% to 50%	51% to 75%	76% to 100%
Have already been virtualized	44%	49%	6%	1%	0%
Expect to be virtualized within a year	24%	53%	20%	1%	1%

The data in **Table 23** indicates the growing interest that IT organizations have in desktop virtualization. For example:

Over the next year, the percentage of IT organizations that have not implemented any desktop virtualization will be cut roughly in half.

Part of the challenge in supporting virtualized desktops is that the implementation of virtualized desktops puts more traffic on the WAN, which typically leads to the need for more bandwidth. In addition to the bandwidth challenges, as explained in [The 2012 Application and Service Delivery Handbook](#)³¹, there are performance challenges associated with each of the two primary form of desktop virtualization; e.g., client side (a.k.a., streamed desktops) and server side (a.k.a., hosted desktops).

Collaboration

As was described in the section of this report that is entitled *The Emergence of Cloud Computing and Cloud Networking*, many organizations are beginning to acquire services such as collaboration from a cloud computing service provider (CCSP). Independent of whether the collaboration service is provided by a CCSP or by the IT organization, it stresses the WAN. This stress comes in part from the fact that the performance of applications such as video and telepresence is very sensitive to delay, jitter and packet loss. The stress also comes in part because video and telepresence consume considerable WAN bandwidth. It is common, for example, to allocate several megabits per second of WAN bandwidth to a single telepresence session.

³¹ <http://www.webtorials.com/content/2012/07/2012-application-service-delivery-handbook-1.html>

Mobile Workers

In the last few years there has been an explosive growth in the number of mobile workers. There are a number of concerns relative to supporting mobile workers. One such concern is that up through 2010, the most common device used by a mobile worker was a PC. In 2011, however, more tablets and smartphones shipped than PCs³². Related to the dramatic shift in the number and types of mobile devices that are being shipped, many companies have adopted the BYOD (Bring Your Own Device to Work) concept whereby employees use their own devices to access applications.

The **Survey Respondents** were asked to indicate the types of employee owned devices that their organization allows to connect to their branch office networks and which of these devices is actively supported. Their responses are shown in **Table 24**.

Table 24: Support for Employee Owned Devices			
	Not Allowed	Allowed but not Supported	Allowed and Supported
Company managed, employee owned laptop	22%	24%	54%
Employee owned and managed laptop	38%	38%	25%
Blackberry	17%	24%	58%
Apple iPhone	14%	30%	55%
Android phone	19%	33%	48%
Windows mobile phone	26%	40%	34%
Apple iPad	18%	40%	52%
Android based tablet	28%	37%	35%
Windows based tablet	28%	36%	37%

The data in **Table 24** indicates that there is wide acceptance BYOD. In particular:

IT organizations are required to support a wide range of end user devices.

As a result of the movement to adopt BYOD, the typical branch office network now contains three types of end user devices that are all accessing business critical applications and services. This includes PCs as well as the new generation of mobile devices; i.e., smartphones and tablet computers. Because of their small size, this new generation of mobile devices doesn't usually have wired Ethernet ports and so they are typically connected via what is hopefully a secure WiFi network in the branch office or a 3G/4G service when WiFi isn't available.

Another key concern relative to supporting mobile workers is how the applications that these workers access has changed. At one time, mobile workers tended to primarily access either recreational applications or applications that were not delay sensitive; e.g., email. However, in the current environment mobile workers also need to access a wide range of business critical

³² <http://gizmodo.com/5882172/the-world-now-buys-more-smartphones-than-computers>

applications, many of which are delay sensitive. This shift in the applications accessed by mobile workers was highlighted by SAP's announcement³³ that it will leverage its Sybase acquisition to offer access to its business applications to mobile workers.

One of the technical issues associated with supporting mobile workers' access to delay sensitive, business critical applications is that because of the way that TCP functions, even the small amount of packet loss that is often associated with wireless networks results in a dramatic reduction in throughput. A related issue is that typically there is a large amount of delay associated with 3G and 4G networks.

WAN Requirements

This subsection of The Report will summarize some of the emerging requirements that WANs must satisfy on a going forward basis. For example, whether providing connectivity from the branch office to the corporate data center or from the branch office to a CCSP's facility, the WAN must be able to prioritize applications in accordance with business priorities. MPLS provides built-in Class-of-Service (CoS), but the Internet does not have that type of capability. Nevertheless, in order to meet business demands the network must be able to examine, recognize and classify network traffic in a way that reflects business priorities and not just the network protocol or TCP/UDP port numbers. The ability to recognize and differentiate different applications on the network requires Deep Packet Inspection (DPI), network fingerprinting and pattern matching. Once the network traffic has been recognized, it must be placed into queues that reflect the different Quality of Service (QoS) demands that are associated with the varying traffic types and business priorities. As applications on the network change, the network must recognize these changes and adapt to them. This capability is needed whether MPLS or the Internet is being used and whether the branch office user is accessing resources in the corporate data center or in a CCSP's facilities.

Given the complexity of contemporary applications combined with the impact of poorly performing business critical applications, it is critical that IT organizations have visibility into each component of the network in order to understand both the underlying cause of poor application performance and to identify what is needed to remedy the problem and restore full application performance. In order to have this visibility, IT organizations must be able to determine what path or paths between the end users and the applications are associated with degraded application performance. IT organizations must also have the capability to isolate the performance problem to a particular network segment and then be able to use effective diagnostic tools to rapidly determine the root cause of the degradation.

As previously mentioned, in the current environment some end users are mobile and some reside in branch offices. In addition, some applications are hosted in a corporate data center and some at a CCSP's facility. These two factors mean that visibility must be provided across MPLS networks as well as across the Internet and that the visibility must extend from the end user to both the corporate data center and to the CCSP's facilities. In addition, performance data needs to be collected from various points in the network in order to create a comprehensive view of the end-to-end performance. Performance data can be collected from routers and/or dedicated appliances. Routers and dedicated appliances typically send performance data to collection and reporting systems that analyze the data, frequently using Flexible Netflow or the IETF IP Flow Information Export; i.e., IPFIX, RFC 5101, 5102 and 5103.

³³ Wall Street Journal, May 17, 2011, page B7

Traditional routing protocols can be used to reroute traffic when network segments fail, but they are not capable of identifying the optimal or best path for network traffic. In addition, when a network link is overloaded routing protocols continue to send traffic to it even though a better performing path is available. These limitations of routing protocols often result in unnecessary application degradation. With the growing adoption of cloud computing and the increasing reliance on the Internet to provide connectivity from branch offices to facilities provided by CSPs, it is critical that the best possible path (e.g., some combination of low delay, jitter and packet loss) from the users to the applications be identified in order to improve application performance. Identifying the best path from point A to point B enables IT organizations to place business critical applications on the path that exhibits the best performance while the remaining applications transit a different path.

One of the implications of the growing adoption of cloud computing is that the network must be able to dynamically provision secure and reliable connectivity between branch offices and a CCSP's facilities to support functions such as cloud bursting and failover/disaster recovery between a corporate data center and a CSP. While some applications use SSL encryption, for those that don't, network level encryption is typically needed and in many cases it is required by regulations or industry standards; e.g. HIPPA and PCI DSS. There are several technologies that are used to encrypt network traffic including the IETF's RFC 4301 IPsec and Group Encrypted Transport VPN (GET VPN).

The previously mentioned growing adoption of a new generation of mobile devices has transformed how end users access applications. It has also created challenges for application architects as well as for those responsible for the IT infrastructure and for security. One way to respond to these challenges is to take the software that previously ran on desktop and laptop PCs and which provided key network functionality, and run this software on the new generation of mobile devices. One limitation of this approach is that these devices have limited processing power – typically one tenth the processing power of traditional desktops and laptops. In addition, this new generation of mobile devices often has significant limitation on the size and resolution of their display and can have limitations on the functionality of their web browsers, such as not supporting Flash. The network must be able to adapt to these restrictions and support the new generation of mobile devices that already outsell PCs and will soon become the dominant form of end user device. This includes being able to assess the device's security posture, VPN compatibility and wherever possible, providing WAN optimization.

The last few years has seen the evolution of a new generation of very sophisticated hacker. For example, it is somewhat common for crime families, hactivists and national governments to take advantage of Internet connectivity to gain access to applications, servers and end user devices. They use this access to achieve a variety of ideological and political goals as well as to extort money. The network must be able to automatically block any and all attempted security attacks and to allow only legitimate traffic to transit the network. Network security must be effective and pervasive at all points of the network. This includes network access and egress links as well as end user devices, whether those devices are mobile or stationary. With the movement to adopt cloud computing and to allow branch office networks to connect directly to CCSPs, network security in the corporate data center is no longer sufficient. In order to be effective, network security must be distributed to branch office networks as well as to CCSPs. This means that virtualized network firewalls and network access control systems must be cost effective enough so that IT organizations can afford to deploy them in branch office networks. Network firewalls and network access support systems that are deployed at CCSPs' facilities must work together with the network security systems that reside both in enterprise branch office networks and in

corporate data centers in order to provide a comprehensive, defense in depth network security. In order to reduce or eliminate the backhauling of Internet traffic, the branch offices' network equipment must also be sophisticated enough to provide the same level of security as is traditionally provided at the corporate data center. The final section of The Report entitled *Management and Security* will discuss the varying ways that IT organizations are implementing security and will also discuss the role of cloud based security.

The broad and rapidly growing movement to adopt both cloud computing and a new generation of mobile devices makes it significantly more difficult to achieve some of the goals of effective service delivery; e.g., effective management, appropriate levels of security. That increased difficulty stems from the fact that while it will still be common for the IT organization to own and manage the IT infrastructure that supports business critical applications, on an increasing basis that infrastructure will be owned and managed by one or more CSPs. Despite being owned by the CSP, the IT organization still needs to have end-to-end visibility and to be able to direct security policies. Standards such as NetFlow v9 and IETF IPFIX are key building blocks that can be leveraged to provide that functionality.

Traditional WAN Services

Background

The **Survey Respondents** were given a set of eleven WAN services and asked to indicate the extent to which they currently utilize each WAN service. 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 Webtorials Respondents don't have any Frame Relay in their networks and almost two thirds of The Webtorials Respondents don't have any ATM in their networks. In addition, few IT organizations are increasing their use of these technologies³⁴, while many IT organizations are decreasing their use of these technologies³⁵.

One of the observations that can be drawn from the response to this survey question is that:

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

Because of the prevalence of MPLS and the Internet and the lack of development of new WAN technologies, the majority of the rest of this section of The Report will discuss how functionality is being added to MPLS and the Internet to respond to the emerging requirements. This section will also briefly discuss the possible use of software defined networking in the WAN.

WAN Design Criteria and Challenges

The **Survey Respondents** were given a list of possible concerns and were asked to indicate which two were their company's primary concerns relative to its use of MPLS and the Internet. The set of concerns that were presented to the **Survey Respondents** is shown in the left hand column of **Table 25**. The second and third columns from the left in **Table 25** show the percentage of the **Survey Respondents** who indicated that the concern is one of their company's two primary concerns with MPLS and the Internet respectively. The right hand column is the difference between the second and third columns from the left. This column will be referred to as the delta column.

The delta column contains positive and negative numbers. A positive number means that that concern was mentioned more often relative to MPLS than it was mentioned relative to the Internet. For example, the **Survey Respondents** mentioned cost as one of their primary concerns about the use of MPLS 22.1% more often than they mentioned cost as one of their primary concerns about the use of the Internet. Analogously, a negative number means that that concern was mentioned more often relative to the Internet than it was relative to MPLS. For example, the **Survey Respondents** mentioned latency as one of their primary concerns about the use of the Internet 19.3% more often than they mentioned latency as one of their primary concerns about use of MPLS.

³⁴ Roughly 2% of IT organizations are increasing their use of Frame Relay and 6% of IT organizations are increasing their use of ATM.

³⁵ Roughly 34% of IT organizations are decreasing their use of Frame Relay and 22% of IT organizations are decreasing their use of ATM.

Table 25: Concerns about MPLS			
Concern	MPLS	Internet	Delta
Cost	60.1%	38.0%	22.1%
Lead time to implement new circuits	32.2%	11.4%	20.8%
Uptime	30.1%	46.3%	-16.2%
Latency	27.0%	46.3%	-19.3%
Lead time to increase capacity on existing circuits	23.5%	13.1%	10.4%
Jitter	14.8%	18.8%	-4.0%
Packet Loss	12.2%	26.2%	-14.0%

The primary concerns that IT organizations have with the use of MPLS are cost, the lead time to implement new circuits and uptime. The primary concerns that IT organizations have with the use of the Internet are uptime, latency and cost.

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

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

As shown in **Table 25**, MPLS is regarded by the **Survey Respondents** as doing a good job at ensuring appropriate performance because it exhibits relatively small amounts of delay, jitter and packet loss. Unfortunately, MPLS is regarded poorly relative to the goal of minimizing cost. In contrast, the Internet is regarded relatively well on the goal of minimizing cost but is regarded relatively poorly on the goal of ensuring appropriate performance. In addition, the **Survey Respondents** expressed concerns about both MPLS and the Internet relative to the goal of maximizing availability.

One viable approach to WAN design is to use both the Internet and MPLS in ways that maximize the benefits of each while minimizing their deficiencies.

As was pointed out in the section of this report entitled *The Emergence of Cloud Computing and Cloud Networking*, the goal of cloud computing is to enable IT organizations to achieve a dramatic improvement in the cost effective, elastic provisioning of IT services that are *good enough*. As that section also pointed out, in order to support a small number of business critical services and applications, a cloud network that is *good enough* will have to provide the highest possible levels of availability and performance. However, in a growing number of instances, a cloud network is *good enough* if it provides a best effort level of service at a reduced price. Hence, independent of the concerns those IT organizations have about the Internet:

In a growing number of instances, Internet-based VPNs that use DSL for access are 'good enough' to be a cloud network.

Some of the concerns that IT organizations have with the use of the Internet such as uptime, stem from the fact that in many cases IT organizations access the Internet over a single DSL link. The availability of DSL is somewhat lower than the availability of access technologies such as T1/E1 links. One impact of this reduced availability is that Internet VPNs based on DSL access are 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 such as DSL and cable. One technology that addresses this issue is referred to as an *aggregated virtual WAN*.

The key concept behind an aggregated virtual WAN is that it simultaneously utilizes multiple enterprise WAN services and/or Internet connections in order to optimize reliability and minimize packet loss, latency and jitter.

Aggregated virtual WANs and other types of alternate WAN services are discussed later in this section of the report. As that discussion highlights, aggregated virtual WANs have the potential to maximize the benefits of the Internet and possibly MPLS while minimizing the negative aspects of both.

Local Access to the Internet

The traditional approach to providing Internet access to branch office employees has been to carry their 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 their Internet traffic and it 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's WAN, which adds to the cost of the WAN. Another disadvantage of this approach is that it usually adds additional delay to the Internet traffic. The fact that centralized Internet access exhibits these disadvantages is significant because as highlighted in **Table 26**, cost and delay are two of the primary concerns that IT organizations have relative to the use of the Internet.

Some of the concerns that IT organizations have about the use of the Internet are exacerbated by backhauling Internet traffic to a central site.

The **Survey 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 shown in **Table 26**.

The way to read the data in **Table 26** is that 32.1% of the **Survey Respondents** route 100% of their Internet traffic to a central site and that 17.3% of the **Survey Respondents** route between 76% and 99% of their Internet traffic to a central site

Table 26: Centralized Access to the Internet	
Percentage of Internet Traffic	Currently Routed to a Central Site
100%	32.1%
76% to 99%	17.3%
51% to 75%	15.6%
26% to 50%	13.1%
1% to 25%	12.2%
0%	9.7%

The **Survey Respondents** also indicated that driven in part to save money and in part to improve application performance that:

Over the next year, IT organizations will make an increased use of distributed access to the Internet from their branch offices.

Cloud Networking Without the Internet

There is a temptation to associate the WAN component of *cloud networking* either exclusively or primarily with the traditional Internet³⁶. However, due to a variety of well-known issues, such as 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 **Survey Respondents** were asked to indicate which WAN service their users would most likely use when accessing public and private cloud computing services over the next year. Their responses are shown in **Table 27**.

Table 27: WAN Services to Access Cloud Computing Services				
	The Internet	An Internet Overlay	A traditional WAN service such as MPLS	WAN Optimization combined with a traditional WAN service; e.g. MPLS
Public Cloud Computing Services	61.2%	4.9%	18.8%	15.1%
Private Cloud Computing Services	35.3%	1.0%	36.7%	27.0%

The data in **Table 27** indicates that IT organizations understand the limitations of the traditional Internet relative to supporting cloud computing. In particular:

In roughly forty percent of the instances that business users are accessing public cloud computing services, the primary WAN service that they intend to use is not the traditional Internet.

In almost two thirds of the instances that business users are accessing private cloud computing services, the primary WAN service that they intend to use is not the traditional Internet.

Techniques that IT organizations can use to mitigate their concerns about the use of the Internet are discussed later in this section of the report.

³⁶ Throughout this report, the phrase "traditional Internet" will refer to the use of the Internet without the use of any optimization functionality.

Service Level Agreements

As previously stated, 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 varying types of applications that transit a WAN. For example, 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 MPLS 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 **Survey 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 20**.

The fact that two thirds of the **Survey 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 network service providers as being effective.

Figure 20: The Effectiveness of SLAs




Optimizing the Performance of IT Resources

Background

This subsection of The Report will discuss techniques that IT organizations can implement to overcome the limitations of protocols and applications and to optimize the use of their servers. The focus of this subsection is on how these techniques enable IT organizations to ensure acceptable application and service delivery over a WAN. The discussion in this subsection will focus on two classes of products: WAN Optimization Controllers (WOCs) and Application Delivery Controllers (ADCs).

The introduction to this section of The Report discussed how the adoption of cloud computing in general is impacting the WAN and also discussed some of the specific factors that are driving change in the WAN. These factors included both the increasing number of mobile workers and the impact of multiple forms of virtualization. In order to gauge the effect that these factors have on the ability of an IT organizations to ensure acceptable application and service delivery, The **Survey Respondents** were asked “How will the ongoing adoption of mobile workers, virtualization and cloud computing impact the difficulty that your organization has with ensuring acceptable application performance?” Their responses are shown in **Figure 21**.



One conclusion that can be drawn from **Figure 21** is that:

The majority of IT organizations believe that factors such as the growth in the number of mobile workers and the increase in the use of virtualization and cloud computing will make ensuring acceptable service and application delivery either harder or notably harder.

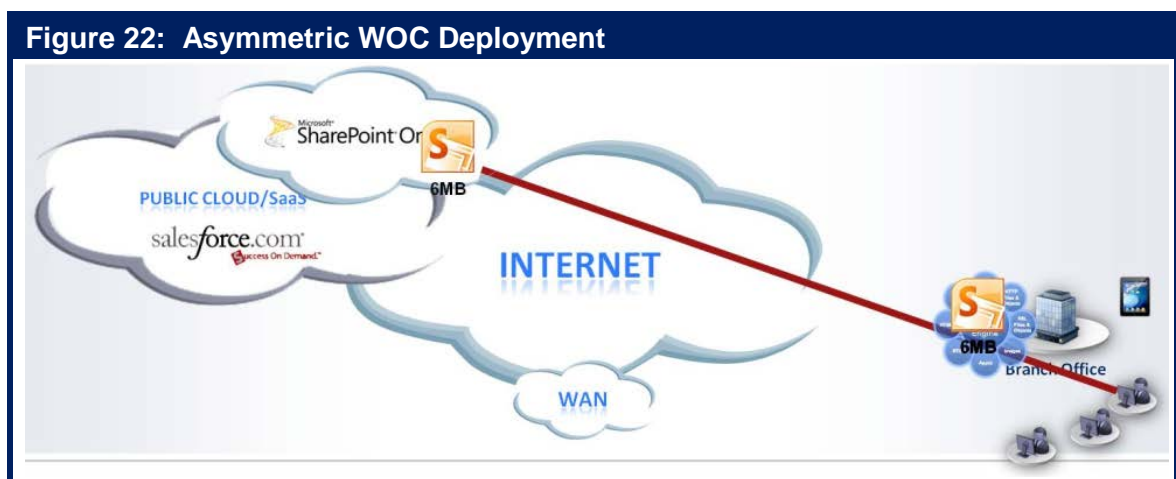
WAN Optimization Controllers (WOCs)

Goals of a WOC

The goal of a WOC is to improve the performance of applications and services that are delivered across a WAN from the data center either to a branch office, a home office or directly to a mobile user. In some cases the data center is owned and managed by the enterprise IT organization and in other cases it is owned and managed by a CCSP. The WOC accomplishes this goal by implementing techniques to overcome the limitations of the WAN such as constrained bandwidth, delay and packet loss.

WOCs are often referred to as *symmetric solutions* because they usually require complementary functionality at both ends of the connection; i.e., a WOC in the data center and another WOC at the branch office. However, the requirement to improve the performance of applications and services acquired from a CCSP has been the impetus for the deployment of WOCs in an asymmetric fashion. One of the advantages of an asymmetric deployment of a WOC is shown in **Figure 22**. As shown in the figure, in an asymmetric deployment of a WOC content is downloaded from a CCSP to a WOC in a branch office. Once the content is stored in the WOC's cache for a single user, subsequent users who want to access the same content will experience accelerated application delivery. Caching can be optimized for a range of cloud content, including Web applications, streaming video (e.g., delivered via Flash/RTMP or RTSP) and dynamic Web 2.0 content.

As previously described, IT organizations are moving away from a WAN design in which they backhaul their Internet traffic from their branch offices to a central site prior to handing it off to the Internet. Also, as is described in the next section of this report, there are a variety of techniques that enable IT organizations to improve both the price-performance and the availability of distributed Internet access. As a result of these factors, asymmetric WOC deployment as described in the preceding paragraph will increasingly be utilized as part of a network design that features distributed Internet access. However, for this network design to be effective, IT organizations need to ensure that the design includes appropriate security functionality.



Modeling Application Response Time

A model is helpful to illustrate how the performance of a WAN can impact the performance of an application and it also serves to illustrate how a WOC can improve application performance. The following model (**Figure 23**) is a variation of the application response time model created by Sevcik and Wetzel³⁷. Like all mathematical models, the following is only an approximation. For example, the model shown in **Figure 23** doesn't account for the impact of packet loss.

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).

Figure 23: Application Response Time Model

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

In order to improve the performance of applications that are delivered over the WAN, WOCs implement a variety of techniques. For example, to mitigate the impact of a large payload, WOCs implement techniques such as compression and de-duplication. These techniques are explained in detail in [The 2012 Application Delivery Handbook](#). The handbook also details criteria that IT organizations can use to evaluate WOCs as well as specific techniques that WOCs need to support in order to optimize:

- The rapidly growing amount of traffic that goes between data centers
- Desktop virtualization
- Delay sensitive applications such as voice, video and telepresence

The **2012 Application Delivery Handbook** also describes techniques that can optimize the delivery of applications to mobile workers. Many IT organizations, however, resist putting any additional software on the user's device. In addition, many users resent having multiple clients (e.g., WOC, SSL VPN, IPsec VPN, wireless/cellular access) on their access device that are not integrated. One option for IT organizations on a going forward basis is to implement WOC software on mobile devices that is integrated with the other clients used by mobile workers. As is explained below, an alternative way that IT organizations can improve the performance of applications and services delivered to mobile users is to utilize an optimization service from a CCSP.

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 a SLB, the ADC has assumed, and will most likely continue to assume, a wide range of sophisticated roles

³⁷ Why SAP Performance Needs Help, NetForecast Report 5084, <http://www.netforecast.com/ReportsFrameset.htm>

that enhance server efficiency and security and which provides 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 23](#), 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 [2012 Application Delivery Handbook](#) describes the primary techniques implemented by ADCs and identifies criteria that IT organizations can use to evaluate ADCs

Virtual Appliances

The section of this report entitled *The Emerging Data Center LAN* 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 a variety of fundamentally different ways. For example, two or more appliances, such as ADCs, can be combined to appear as a single logical ADC. Alternatively, a single physical ADC can be partitioned into a number of logical ADCs or ADC contexts. Each logical ADC can be configured individually to meet the server-load balancing, acceleration and security requirements of a single application or a cluster of applications.

However, the most 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. As explained in the next subsection of this report, virtual appliances make it easier for an IT organization to deploy network and application optimization functionality at a CCSP's data center. That, however, is not the only advantage of a virtualized appliance.

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 acquisition cost of a software-based appliance can be a third less than the cost of a hardware-based appliance³⁸. 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 appliance³⁹.

³⁸ The actual price difference between a hardware-based appliance and a software-based appliance will differ by vendor.

³⁹ This statement makes a number of assumptions, including the assumption that the vendor does not charge for the backup software-based appliance.

In addition, many IT organizations choose to implement a proof-of-concept (POC) trial prior to acquiring an appliance such as a WOC or an ADC. Using WOCs as an example, the purpose of these trials is to enable the IT organization to quantify the performance improvements provided by the WOCs and to understand related issues such as the manageability and transparency of the WOCs. While it is possible to conduct a POC using a hardware-based WOC, it is easier to do so with a virtual WOC. This follows in part because a virtual WOC can be downloaded in a matter of minutes, whereas it typically takes a few days to ship a hardware-based WOC. The value of the ease of downloading a virtual appliance is magnified in those cases in which the appliance is being delivered to a country where it takes a long time to get through customs.

Virtual appliances make it easier to conduct a proof of concept trial.

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


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

A virtualized ADC also 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 virtual 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 latter is a viable option from an organizational perspective because any actions taken by the application group relative to their virtual ADC will only impact their application.

A virtual firewall appliance can also 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. Through tight integration with the virtual server management system, virtual 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 Survey Respondents were asked whether or not their company had deployed virtual functionality in their branch office networks. Fifty-five percent indicated that they had and those respondents were then asked to indicate the type of virtual functionality their organization had implemented. Their responses are shown in [Figure 24](#).

Figure 24: Virtual Functionality in Branch Offices



Optimizing Access to Public Cloud Computing Solutions

The conventional wisdom in the IT industry is that one of the key challenges facing IT organizations that use public cloud based solutions is improving the performance of those solutions. In order to understand how IT organizations intend to optimize the performance of services that they acquire from CCSPs, the **Survey Respondents** were given the following question:

If your company either currently acquires services from an Infrastructure-as-a-Service (IaaS) provider or you expect that it will within the next year, which of the following best describes the primary approach that your company will take to optimize the performance of those services?

Their responses are shown in **Table 28**.

Table 28: Optimizing IaaS Services	
TECHNIQUE	Percentage of Respondents
Don't know	35.3%
We will leverage optimization functionality provided by the IaaS provider	27.8%
We will not do anything	18.2%
We will place a WAN optimization controller on the service provider's site and on our site	16.0%
We will use an optimization service from a company such as Akamai or Aryaka	2.7%

One conclusion that can be drawn from the data in **Table 28** is:

The majority of IT organizations are either undecided about how they will optimize the performance of IaaS services or they intend to do nothing.

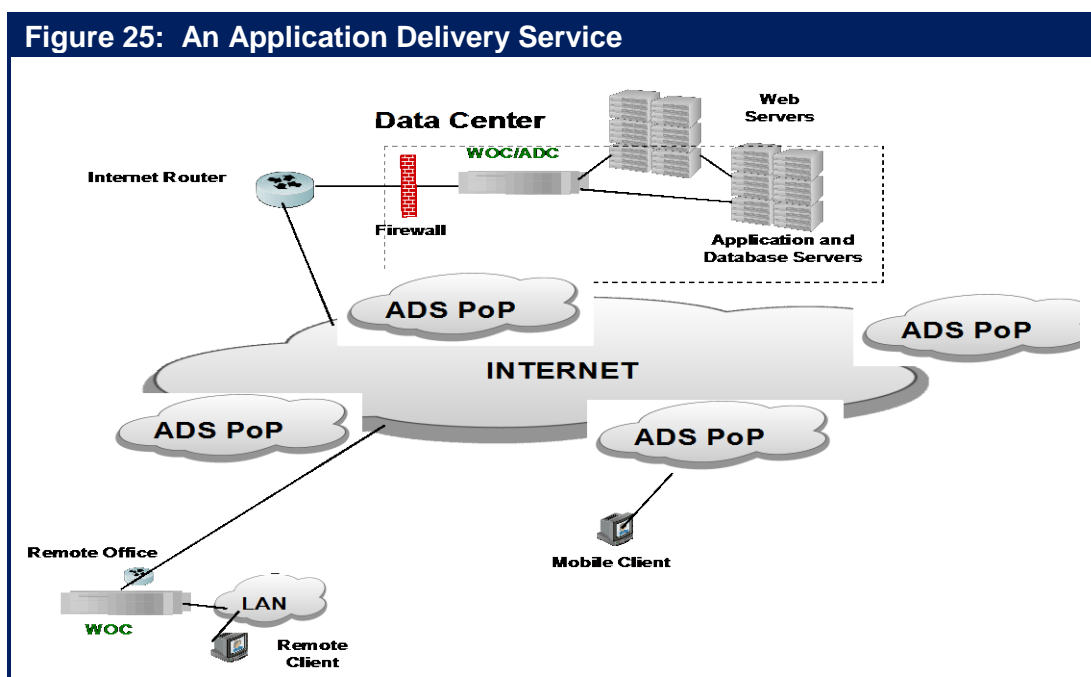
The data in **Table 28** also shows some interest on the part of IT organizations to place a WOC on premise at an IaaS provider's data center. Referring back to the discussion in the previous subsection, IT organization will have a notably easier time placing an optimization device, whether that is a WOC or an ADC, at an IaaS provider's data center if the device is virtualized. That follows because if the device is virtualized, the IT organization can control the deployment of the functionality. If the device is physical, then the IT organization needs to get the IaaS provider to offer space for the device and to install it.

Alternative WAN Services

As noted, there isn't a new generation of fundamentally new technology focused on the WAN that is 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. A number of these alternatives are either complementary to the WAN optimization technologies previously discussed or they depend partially on WAN optimization technologies to deliver acceptable levels of service quality.

An Internet Overlay

As described in the preceding subsection, IT organizations often implement WOCs and ADCs in order to improve network and application performance. However, these solutions make the assumption that the 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 doesn't 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 25**, 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) that is close to users of the application, typically within a single network hop. This edge server that is close to the users then optimizes the traffic flow to the server closest to the data center's origin server. Throughout this section, the Internet overlay that is depicted in **Figure 25** will be referred to as an Application Delivery Network (ADN).



An ADN provides a variety of optimization functions that generally complements the functionality provided by WOCs and ADCs. One such function that is often provided by an ADN 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 that are close to the users. Because the content is close to the users, 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 functionality that is often associated with an ADN includes:

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

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

One use case for an ADN that is growing in importance stems from that fact that not all CCSPs will support virtual WOC instances in their data centers. This is particularly true of SaaS providers. Access to services provided by a CCSP can be accelerated via an ADN.

An Integrated Private-Public Solution

In almost all instances when a user accesses a CCSP-provided application or service they do that over the Internet and not over a private WAN service such as MPLS. That follows in large part because from the perspective of the CCSP, one or two high-speed Internet connections are much simpler and more economical to provision and manage than are connections to the varying private WAN services offered by multiple network service providers. In addition, the high fixed costs of these private WAN services can detract significantly from the overall cost-effectiveness of providing SaaS-based applications.

As previously discussed it is quite common for IT organizations to provide Internet access to branch office employees by carrying their Internet traffic on a private network (e.g., an MPLS network) to a central site where the traffic is handed off to the Internet. As was also previously discussed, many IT organizations have implemented WOCs in order to overcome the performance challenges that are associated with private WAN services. This means that the existing WOCs can utilize technology to overcome performance challenges such as TCP's retransmission timeout and the TCP slow start algorithm over the private WAN that connects a branch office to a central site. However, in the traditional scenario, once that traffic is handed off to the Internet, the performance of the application is negatively impacted by the limitations of TCP and by the transmission impairments (e.g., delay, jitter, packet loss) within the Internet.

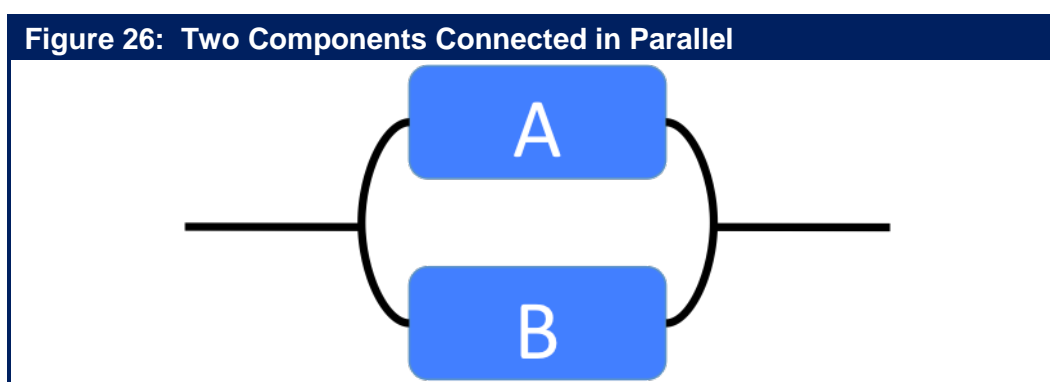
It is possible to mitigate the impact of the performance impairments that are associated with employees in branch office using both a private network and the Internet to access a CCSP-provided applications and services by implementing an end-to-end approach to network and application optimization. A key component of the end-to-end approach is to integrate the optimization that is in place for private WANs with the performance gains that are provided by an ADN. As part of this integration, key functionality that is part of the ADN must be integrated

into the WOC that sits in the enterprise data center. In addition, WOCs have to be distributed to the PoPs that support the ADN. The integration ensures a seamless handoff of functionality such as TCP optimization between the WOC in the data center and the ADN.

Dual ISP Internet VPN with Policy Based Routing

A preceding section of this report identified the concerns that IT organizations have with the use of the Internet. The two primary concerns are uptime and latency. Another approach to overcoming the limitations of the Internet is to connect each enterprise site to two ISPs. Having dual connections can enable IT organizations to add inexpensive WAN bandwidth and can dramatically improve the reliability and availability of the WAN.

For example, **Figure 26** depicts a system that is composed of two components that are connected in parallel.



The system depicted in **Figure 26** 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 time⁴⁰. This level of availability is equal to or exceeds the availability of most MPLS networks.

Traffic can be shared by the two connections by using Policy Based Routing (PBR). 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, policy-based routing allows network administrators to create routing policies to select the path for each packet based on factors such as the identity of a particular end system, the protocol or the application.

Perhaps the biggest limitation of the PBR approach is 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 nature of the policies means that, unless there is an outage of one of the links, a given class of traffic will always be allocated to the same network connection.

⁴⁰ If, as described later, 4G is added as a third access technique and if each access technique has an availability of 99%, then the system as a whole has an availability of 99.9999%.

Dual ISPs and PBR can be used in conjunction with WOCs to further alleviate the shortcomings of Internet VPNs, bringing the service quality more in line with MPLS at a much lower cost point. For example, a WOC can classify the full range of enterprise applications, apply application acceleration and protocol optimization techniques, and shape available bandwidth in order to manage application performance in accordance with enterprise policies. As a result,

In many situations, a dual ISP-based Internet VPN with PBR can deliver a level of CoS and reliability that is comparable to that of MPLS at a significantly reduced price.

Part of the cultural challenge that IT organizations have relative to migrating traffic away from their MPLS network and onto an Internet based network is that Internet based networks don't provide a performance based SLA. However, as previously described, the majority of IT organizations don't place much value in the SLAs that they receive from their network service providers.

Hybrid WANs with Policy Based Routing

As noted, some IT organizations are reluctant to abandon traditional enterprise services such as MPLS. An alternative design that overcomes their concerns is a hybrid WAN that leverages multiple WAN services, such as traditional enterprise WAN services and the Internet, and which uses PBR for load sharing. The advantage of a hybrid WAN is that the CoS of MPLS can be leveraged for delay sensitive, business critical traffic with the Internet VPN used both for other traffic and as a backup for the MPLS network. As in the case of the dual ISP based Internet VPN, the major disadvantage of this approach is the static nature of the PBR forwarding policies. Since PBR cannot respond in real time to changing network conditions, it will consume more costly bandwidth than would a dynamic approach to traffic allocation. A second drawback of hybrid WANs based on PBR is that they can prove to be overly complex for some IT departments. As with many other types of WAN services, hybrid WANs can also be used in conjunction with WOCs and ADCs.

Aggregated Virtual WANs

A relatively new class of device has emerged to address the shortcomings of PBR-based hybrid WANs. WAN path controller (WPC) is one phrase that is often used to describe devices that work in conjunction with WAN routers to simplify PBR and to make the selection of the best WAN access link or the best end-to-end WAN path from a number of WAN service options.

Some members of this emerging class of products are single-ended solutions whereby a device at a site focuses on distributing traffic across the site's access links on a per-flow basis. Typical capabilities in single-ended solutions include traffic prioritization and bandwidth reservation for specific applications. These products, however, lack an end-to-end view of the available paths and are hence limited to relatively static path selections.

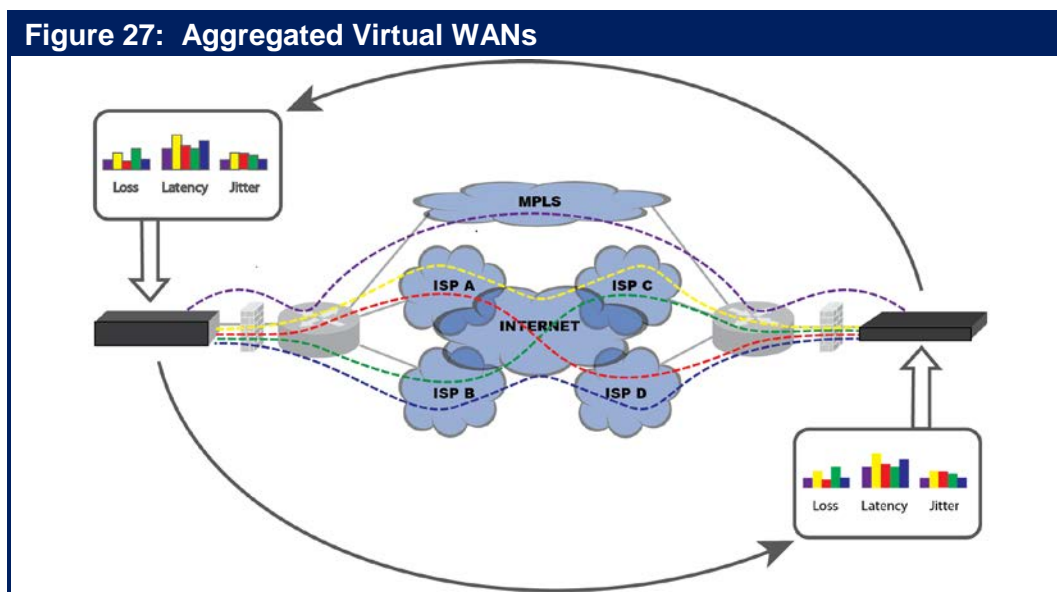
In contrast, symmetrical or dual-ended solutions are capable of establishing an end-to-end view of all paths throughout the network between originating and terminating devices and these solutions can distribute traffic across access links and specific network paths based on either a packet-by-packet basis or a flow basis. These capabilities make the multiple physical WAN services that comprise a hybrid WAN appear to be a single *aggregated virtual WAN*.

Aggregated virtual WANs (avWANs) represent another technique for implementing WANs based on multiple WAN services (e.g., MPLS, Frame Relay and the Internet) and/or WANs based on just multiple Internet VPN connections. An aggregated virtual WAN transcends simple PBR by dynamically recognizing application traffic and allocating traffic across multiple paths through the WAN based on real-time traffic analytics, including:

- The instantaneous end-to-end performance of each available network: This allows the solution to choose the optimal network path for differing traffic types. One differentiator among virtual WAN solutions is whether the optimal path is chosen on a per packet basis or on a per flow basis. Per packet optimization has the advantage of being able to respond instantaneously to short term changes in network conditions.
- 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 previously noted, one of the primary reasons why IT organizations backhaul their Internet traffic to a central site over an enterprise WAN service is because of security concerns. In order to mitigate those concerns when using an avWAN for direct Internet access, the avWAN should support security functionality such as encryption.

Like other hybrid WANs, an avWAN (**Figure 27**) allows IT organizations to add significant amounts of additional bandwidth to an existing MPLS-based WAN at a relatively low incremental cost. In addition to enabling the augmentation of an MPLS WAN with inexpensive Internet connectivity, aggregated virtual WANs also give 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 network design provided.



As shown in **Figure 27** because the two avWAN appliances work together to continuously measure loss, latency, jitter and bandwidth utilization across all of the various paths between any two locations, an aggregated virtual WAN can rapidly switch traffic away from a path that is exhibiting an unacceptable level of performance. This capability, combined with the availability advantages of parallel systems as depicted in **Figure 26**, 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 virtually all of the time. This combination of capabilities also underscores the ability of aggregated virtual WANs to deliver performance predictability that equals, and in many cases exceeds, that of a single MPLS network.

Because of the high availability and performance predictability of aggregated virtual WANs, 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 there is an ongoing 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.

Network-as-a-Service

As shown in **Table 25**, the two biggest concerns that IT organizations have with the use of MPLS are its cost and the amount of time it takes to implement new circuits. An emerging WAN service, referred to as Network-as-a-Service (NaaS), is intended to avoid those concerns. As shown in **Figure 27**, NaaS is built using a core network that interconnects a distributed set of Points of Presence (POPs). The phrase *NaaS* implies that unlike MPLS, the service can be deployed rapidly – typically within a day by leveraging Internet links for the first and last mile connections while providing a reliable private core network and additional network intelligence. The service also allows IT organizations to add capacity on demand, rather than provisioning and paying for bandwidth to support future requirements.

In order to meet enterprise requirements, the NaaS must deliver extremely high quality, predictable performance. It must also have enough POPs so that it is close to customers' sites. Some of the other specific characteristics of a NaaS that IT organizations should expect include:

- Centralized visibility across the WAN
- Low latency
- Diversity and redundancy
- Low packet loss
- Instant access to cloud based services or applications
- Support for multiple access methods

- Enterprise class security based on IPSec
- Low total cost of ownership


To mitigate the impact of packet loss on the first and last mile, the service should support a multi-segment TCP optimization architecture on those links as well as on the links that connect the POPs in order to ensure a rapid response to packet loss. The service should also honor industry standard QoS markings.

The next subsection of The Report discusses cloud-based network and application optimization that is based on a network similar to the one described above. Another key feature of a NaaS is that it should allow a customer to use the basic service if that is their choice, but it should also enable the customer to quickly upgrade to add the optimization capabilities discussed in the following subsection of The Report.

Cloud-Based Network and Application Optimization

As mentioned in the section of this report entitled *The Emergence of Cloud Computing and Cloud Networking*, network and application optimization has become available from CCSPs as a Cloud Networking Service (CNS). In this situation, instead of a physical or virtual WOC at each site, the WOC functionality is provided at the CCSP's cloud data centers or POPs, which ideally are in close proximity to the enterprise users, the data centers and the providers of other cloud services. As shown in **Figure 28**, the PoPs are interconnected by the CCSP's core network with customer access to each PoP provided via the Internet or via an enterprise WAN service. The CNS core network could be an Internet overlay, a private IP network or possibly a multi-carrier MPLS/IP network that uses intelligent routing capabilities similar to an aggregated virtual WAN or ADS in order to provide high levels of performance and reliability.

Figure 28: Network and Application Optimization CNS



In **Figure 28** a variety of types of users (e.g., mobile users, branch office users) access WAN optimization functionality at the service provider's points of presence (POPs). These POPs are inter-connected by a dedicated, secure and highly available network. To be effective, the

solution must have enough POPs so that there is a POP in close proximity to the users. In addition, the solution should support a wide variety of WAN access services.

There are at least three distinct use cases for the type of solution shown in **Figure 28**. One such use case is that this type of solution can be leveraged to solve the type of optimization challenges that an IT organization would normally solve by deploying WOCs; e.g., optimizing communications between branch office users and applications in a corporate data center or optimizing data center to data center communications. In this case, the factors that would cause an IT organization to use such a solution are the same factors that drive the use of any public cloud based services; e.g., cost savings, reduce the time it takes to deploy new functionality and provide functionality that the IT organization could not provide itself

The second use case is the ongoing requirement that IT organizations have to support mobile workers. Some IT organizations will resolve the performance challenges associated with supporting mobile users by loading optimization software onto all of the relevant mobile devices. There are two primary limitations of that approach. One limitation is that it can be very cumbersome. Consider the case in which a company has 10,000 mobile employees and each one uses a laptop, a smartphone and a tablet. Implementing and managing optimization software onto those 30,000 devices is very complex from an operational perspective. In addition, as previously discussed the typical smartphone and tablet doesn't support a very powerful processor. Hence, another limitation is that it is highly likely that network and application optimization software running on these devices would not be very effective.

The third use case for utilizing a solution such as the one shown in **Figure 28** is the expanding requirement that IT organizations have to support access to public cloud services. As previously mentioned, in some instances it is possible for an IT organization to host a soft WOC at an IaaS provider's site. However, that is generally not possible at a SaaS provider's site, and in any case a solution with a WOC at either end of a long distance Internet connection cannot address the congestion-based loss that occurs on the Internet. A Cloud-based optimization solution can improve users' access to cloud services by providing to the users the type of functionality typically provided in a WOC: reducing the amount of loss and high latency experienced between the end user's location and the location of the cloud service as well as and minimizing the impact of packet loss when it does occur.


VPLS

As previously mentioned:

VPLS represents the combination of Ethernet and MPLS.

VPLS is a class of VPN that supports the connection of customer edge (CE) Layer 2 switches at multiple sites into a single bridged, multipoint-to-multipoint domain over a service provider's IP/MPLS network, as shown in **Figure 29**. 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 29: 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.

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 assigned inner label, and associates these labels with the VPLS instance.

Table 29 provides a high level comparison of the different types of Ethernet WAN 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.

Table 29: Ethernet WAN Service Types				
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

Software Defined Networking (SDN)

As mentioned in the section of The Report entitled *Software Defined Networking* the most common discussion about implementing SDN focuses on the data center. However, as was also previously mentioned, Google has implemented SDN in their WAN, referred to as the G-Scale WAN, which interconnects their data centers. While SDN will not be a mainstream WAN technology for at least a couple of years, it does potentially represent a new approach to wide area networking.

As previously discussed, the G-Scale control plane is based on BGP and IS-to-IS and the OpenFlow-only switches are very simple 128 port 10 GbE switches that were built by Google using merchant silicon. It is important to note that when Google built these switches, 128 port 10 GbE switches had not yet been introduced in the commercial market. Google also built their own traffic engineering (TE) service. Their TE service collects both real-time utilization metrics and topology data from the underlying network as well as bandwidth demands from applications and services. The Google TE service uses this data to compute the best path for traffic flows and then programs those paths into their switches. The way that Google implemented the G-Scale WAN each site is comprised of multiple switch chassis to provide both scalability and fault tolerance. The sites are connected together and multiple controllers communicate with the switches using OpenFlow.

Google started this project in January 2010 and by early 2012 all of their data center backbone traffic was being carried on the G-Scale WAN. According to Google, some of the benefits of the G-Scale WAN include:

- Unified view of the network fabric: This simplifies configuration, management and provisioning.
- High Utilization: The centralized traffic engineering allows Google to achieve network utilization of up to 95%.
- Faster failure handling: In addition to handling failures faster, the systems converge more rapidly to target optimum and the behavior is predictable.
- Faster time to market/deployment: This comes in part from the fact that only features that are needed are developed.
- Hitless upgrades: The separation of the control plane from the forwarding plane enables hitless software upgrades without packet loss or capacity degradation.
- Elastic compute: The compute capacity of network devices is no longer a limiting factor. Large scale computation is done using the latest generation of servers.

According to Google, some of the challenges they faced were:

- OpenFlow: At the time they started the project, the OpenFlow protocol was just being developed and hence was not feature rich.
- Fault tolerant OpenFlow controllers: To provide high availability and scalability, multiple OpenFlow controllers must be provisioned which at the time that Google deployed the G-Scale WAN, required extra work on their part.
- Partitioning functionality: There is an ongoing lack of clarity in the industry as to what functionality should reside in network devices and what should reside in controllers.
- Flow programming: For large networks, programming of individual flows can take a long time.

Emerging Cloud Networking Specific Solutions

The preceding discussion of WAN services provided some insight into the interplay between the general requirements of cloud computing and the capabilities of WAN services to meet those requirements. One of the goals of this subsection of The Report is to describe the functionality that is required to support a particular form of hybrid cloud computing – cloud balancing. Another goal of this subsection of The Report is to describe some of the optimization functionality that is being developed specifically to support cloud computing.

Cloud Balancing

The phrase *hybrid cloud computing* refers to an IT organization providing IT services in such a way that each of the services is based in part on the private cloud that the IT organization operates and in part on the applications or services provided by one or more CCSPs. A hybrid cloud relies on a WAN to provide the connectivity between the enterprise's locations, including the enterprise's data center(s) and its remote sites, and the CCSP's data center. One of the goals of cloud balancing is to have the collection of individual data centers appear to both users and administrators as a single cloud data center, with the physical location of application resources as transparent as possible. The goal of having the location of application resources be transparent creates a number of requirements. This includes:

- **VLAN Extension**
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. The VLANs within which VMs are migrated must be extended over the WAN between and amongst the private and public data centers. This involves the creation of an overlay network that allows the Layer 2 VLAN traffic to be bridged or tunneled through the WAN.
- **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 enhances security as well as transparency by allowing these application services to be provisioned from the private enterprise data center and by eliminating 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 the location of the users or the IT resources that the users are accessing. This means that the public cloud data centers need to offer the same WAN optimization and application acceleration capabilities that are deployed within the enterprise. In addition, WOCs may well be needed between the enterprise's private cloud data center(s) and the public cloud data

center(s) in order to accelerate VM migration, system backups, and other bulk data transfers between these data centers.

- **Interoperability Between Local and Global ADC Functions**

Cloud balancing is based on making routing decisions based on a combination of local and global variables. This requires interoperability between local and global ADC functions.

- **Synchronizing Data between Cloud Sites**

In order for an application to be executed at the data center that is selected by the cloud balancing system, the target server instance must have access to the relevant data. In some cases, the data can be accessed from a single central repository. In other cases, the data needs to co-located with the application. The co-location of data can be achieved by migrating the data to the appropriate data center, a task that typically requires highly effective optimization techniques. In addition, if the data is replicated for simultaneous use at multiple cloud locations, the data needs to be synchronized via active-active storage replication, which is highly sensitive to WAN latency.

WAN Optimization and Application Delivery for Cloud Sites

One of the most significant trends in the WAN optimization market is the development of new products and new product features that are designed to enable IT organizations to leverage public and hybrid clouds as extensions of their enterprise data centers. Some recent and anticipated developments include:

- **Cloud Optimized WOCs**

These are purpose-built virtual WOC appliances for deployment in public cloud environments. Cloud optimized features include compatibility with cloud virtualization environments, SSL encryption and acceleration, and automated migration or reconfiguration of virtual WOCs in conjunction with VM provisioning or migration. As previously mentioned, WOCs can either be deployed in a symmetric fashion, with a WOC on each end of the WAN link; or in an asymmetric fashion, with a WOC deployed just in a branch office.

- **Cloud-based WAN Optimization Service**

As mentioned in the Cloud-based Network and Application Optimization section above, this solution both leverages the Internet ecosystem and is a solution that provides accelerated, reliable access to public cloud services. It combines cloud-based WAN Optimization technology with a reliable core network, using globally distributed POPs and centralized WAN and application-layer visibility. Optionally an appliance can be deployed on premise for last mile bandwidth scaling. The service is intended to deliver the performance of WOC solutions without the high cost of MPLS or the cost and management overhead of traditional WAN Optimization appliance solutions, in a single combined, fully-managed service with no capital expenditures.

- **Cloud Storage Optimized WOCs**

These are purpose-built virtual or physical WOC appliances for deployment in the enterprise's data center(s) and also at public cloud Storage as a Service environments that are used for backup and archival storage. Cloud optimized features can include support for major backup and archiving tools, de-duplication to minimize the required

data transfer bandwidth and the storage capacity that is required, and support for SSL and AES encryption.

- **Cloud Optimized Application Delivery Controllers**

One trend in the evolution of ADCs is increasing functional integration with more data center service delivery functions. As organizations embrace cloud computing models, service levels need to be assured irrespective of where the applications are hosted. As is the situation with WOCs, ADC vendors are in the process of adding enhancements that support the various forms of cloud computing, including:

- **Hypervisor-based Multi-tenant ADC Appliances**

Partitioned ADC hardware appliances have for some time allowed service providers to support a multi-tenant server infrastructure by dedicating a single partition to each tenant. Enhanced tenant isolation in cloud environments can be achieved by adding hypervisor functionality to the ADC appliance and by dedicating an ADC instance to each tenant. Each ADC instance is then afforded the same type of isolation as a virtualized server instance, with protected system resources and address space. A combination of hardware appliances, virtualized hardware appliances and virtual appliances provides the flexibility for a cloud service provider to offer highly customized ADC services that are a seamless extension of an enterprise customer's IT environment.

- **Cloud Bursting and Cloud Balancing ADCs**

Cloud bursting refers to directing user requests to an external cloud when the enterprise private cloud is at or near capacity. Cloud balancing refers to routing user requests to application instances deployed in the various different clouds within a hybrid cloud. Cloud balancing requires a context-aware load balancing decision based on a wide range of business metrics and technical metrics characterizing the state of the extended infrastructure. By comparison, cloud bursting can involve a smaller set of variables and may be configured with a pre-determined routing decision. However, cloud bursting may require rapid activation of instances at the remote cloud site or possibly the transfer of instances among cloud sites. Cloud bursting and balancing can work well where there is consistent application delivery architecture that spans all of the clouds in question. This basically means that the enterprise's application delivery solution is replicated in the public cloud. One way to achieve this is with virtual appliance implementations of GSLBs and ADCs that support the range of variables needed for cloud balancing or bursting. If these virtual appliances support the IaaS cloud hypervisors, they can be deployed as VMs at each cloud site. The architectural consistency insures that each cloud site will be able to provide the information needed to make global cloud balancing routing decisions. When architectural consistency extends to the hypervisors across the cloud, integration of cloud balancing/bursting ADCs with the hypervisors management systems can help the routing of application traffic synchronized with private and public cloud resource availability and performance. Access control systems integrated within the GSLB and ADC make it possible to maintain control of applications wherever they reside in the hybrid cloud.

Planning for WAN Evolution

The **Survey Respondents** were asked “As your organization evolves its WAN over the next two years, which of the following describes the expectations that your organization will have for the functionality that the WAN will provide. The question had seven classes of WAN functionality and the **Survey Respondents** were asked to indicate all of the classes that applied in their environment. The responses of all of the **Survey Respondents** as well as just the **Survey Respondents** who work in large companies⁴¹ are shown in **Table 30**.

Table 30: WAN Expectations		
	All of The Survey Respondents	The Survey Respondents who work for Large Companies
Provide high level functionality such as security or optimization	50%	62%
Provide basic connectivity between users and business critical IT resources	57%	59%
Utilize basic QoS functionality to support voice, video and telepresence	49%	51%
Be aware of the applications and end points that It supports and adjust accordingly	41%	48%
Utilize not only basic QoS functionality, but also media-aware controls to support enhance voice	37%	44%
Provide integrated security	36%	43%

The data in **Table 30** indicates that the majority of IT organizations continue to see that one role of their WAN is to provide basic connectivity. However, the data also indicates that the majority of all IT organizations and an even bigger majority of large IT organizations also see that on a going forward basis, that their WAN must provide a range of higher value services that correspond closely to the functionality previously discussed in this section of The Report.

The Survey Respondents were also asked two additional questions. Those questions were:

1. *Does your organization have an architecture or strategy document that outlines the current state and likely evolution of your WAN?*
2. *Does the document have a significant influence on decision making around issues such as the choice of technologies, services and vendors (a.k.a., is it effective)?*

⁴¹ Throughout this section of The Report, the phrase *large companies* refers to companies with 10,000 or more employees.

Their responses are shown in **Table 31** and **Table 32**.

Table 31: Does Your Organization have a Documented WAN Strategy?		
	Yes	No
All Companies	50%	50%
Large Companies Only	76%	24%

Table 32: Is Your WAN Strategy Effective?		
	Yes	No
All Companies	76%	24%
Large Companies Only	77%	23%

One conclusion that can be drawn from the data in **Table 31** and **Table 32** is that:

Slightly over a third of all companies, and slightly over a half of large companies have an effective WAN strategy.

In order to successfully respond to the challenges described in this report, IT organizations must create an effective strategy for how they will evolve their WAN. As described in this report, a key component of the WAN strategy that IT organizations must develop is to identify how the organization will continue to provide the same functionality as it does today, as companies make increasing use of public cloud computing services, independent of whether or not traffic is backhauled to a corporate data center prior to being handed off to the Internet. This functionality includes the ability to:

- Optimize application performance
- Provide intelligent QoS that reflects business priorities, not network priorities
- Provide end-to-end visibility of application performance over all segments of the network
- Dynamically route network traffic according to changing conditions
- Enable the growing adoption of all forms of Cloud Computing; e.g., Public, Private, Hybrid.
- Support a variety of end user devices and mobile workers
- Provide integrated network security regardless of the end user device and whether or not they are mobile
- Provide the ability to manage network performance and security policies centrally no matter where and who owns the hardware of the IT infrastructure

Management & Security

Management

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

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

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

A New Set of Management Challenges

Management Challenges Associated with Server Virtualization

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

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

Additional management challenges that are associated with server virtualization include:

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

- **Limited VM-to-VM Traffic Visibility**
The first generation of vSwitches doesn't have the same traffic monitoring features as does physical access switches. This limits the IT organization's ability to do security filtering, performance monitoring and troubleshooting within virtualized server domains.
- **Poor Management Scalability**
Many IT organizations have experienced VM proliferation sometimes called VM sprawl. In addition, the normal best practices for virtual server configuration call for creating separate VLANs for the different types of traffic to and from the VMs. The combined proliferation of VMs and VLANs places a significant strain on the manual processes that are traditionally used to manage servers and the supporting infrastructure.
- **Contentious Management of the vSwitch**
Each virtualized server includes at least one software-based vSwitch. This adds yet another layer to the existing data center LAN architecture. It also creates organizational stress and leads to inconsistent policy implementation.
- **Inconsistent Network Policy Enforcement**
Traditional vSwitches lack some of the advanced features that are required to provide a high degree of traffic control and isolation. Even when vSwitches support some of these features, they may not be fully compatible with similar features that are offered by physical access switches. This situation leads to the implementation of inconsistent end-to-end network policies.
- **Multiple Hypervisors**
It is becoming common to find IT organizations using multiple hypervisors, each of which comes with their own management system and their own management interface. In addition, the management functionality provided by each hypervisor varies as does the degree to which each hypervisor management system is integrated with other management systems.
- **Management on a per-VM Basis**
IT organizations typically perform management tasks such as discovery, capacity planning and troubleshooting on a per server basis. While that is still required, IT organizations must also perform those tasks on a per-VM basis.

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

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

The responses of the **Survey Respondents** are summarized in **Table 33**.

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

One conclusion that can be drawn from the data in **Table 33** is that:

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

Management Challenges Associated with Cloud Computing

Even in the traditional IT environment⁴² when the performance of an application is degrading the degradation is typically noticed first by the end user and not by the IT organization. In addition, when IT is made aware of the fact that application performance has degraded, the process to identify the source of the degradation can be lengthy.

Unfortunately:

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

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

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

⁴² This refers to an IT environment prior to the current wave of virtualization and cloud computing.

resources needed by the application server and that as a result, the application server occasionally performs poorly.

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

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

1. **Public Cloud Computing**

SmartCompany acquires BusApp functionality from a SaaS provider. The employees of SmartCompany that work in branch and regional offices use an MPLS service from a network service provider (NSP) to access the application, while home office workers and mobile workers use the Internet.

2. **Hybrid Cloud Computing**

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

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

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

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

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

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

Importance of Managing Cloud Computing

Three of the twenty tasks that were included in The Question were managing private, hybrid and public cloud computing solutions in an end-to-end manner. The responses of the **Survey Respondents** are summarized in **Table 34**.

Table 34: Importance of Managing Cloud Solutions			
	Private Cloud	Hybrid Cloud	Public Cloud
Extremely	16%	11%	9%
Very	25%	25%	19%
Moderately	25%	28%	23%
Slightly	25%	24%	29%
Not at All	10%	13%	19%

One observation that can be drawn from the data in **Table 34** is that

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

Another observation that can be drawn from the data in **Table 34** is that managing a private cloud is more important than managing a hybrid cloud which is itself more important than managing a public cloud. One of the reasons for this phenomenon is that enterprise IT organizations are making more use of private cloud solutions than they are of either public or hybrid cloud solutions. Another reason for this phenomenon is that as complicated as it is to manage a private cloud, it is notably more doable than is managing either a hybrid or public cloud and IT organizations are placing more emphasis on activities that have a higher chance of success.

The Traditional Management Environment

Network Performance Management Systems

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

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

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

Application Performance Management

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

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

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

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

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

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

Synthetic Transactions

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

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

Internal SLAs

As recently as two or three years ago, few IT organizations offered an SLA to the company's business and functional managers; a.k.a., an internal SLA. However, that situation has changed and now it is common for IT organizations to offer internal SLAs. To understand the prevalence and effectiveness of internal SLAs, The **Survey Respondents** were asked to indicate their agreement or disagreement with three statements. The three statements and the percentage of the **Survey Respondents** that agreed with the statement are shown in **Table 35**.

The data in **Table 35** highlights the fact that:

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


Table 35: Status of Internal SLAs

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

One of the answers to The Question was “managing internal SLAs for one or more business-critical applications”. The responses of the **Survey Respondents** are summarized in **Figure 30**.

The data in **Figure 30** leads to the conclusion that:

Figure 30: The Importance of Getting Better at Managing Internal SLAs



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

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

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

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

Delay Sensitive Traffic

Over the last few years the majority of IT organizations have adopted VoIP and video, which are examples of applications that have high visibility and which are very sensitive to transmission impairments. To identify the emphasis that IT organizations place on managing this type of traffic, the **Survey Respondents** were asked to indicate how important it was over the next year for their IT organization to get better at ensuring acceptable VoIP quality. Their answers are shown in **Table 36**.

Table 36: Importance of Getting Better at Managing VoIP <i>n</i> = 127	
	Percentage
Extremely Important	14%
Very Important	32%
Moderately Important	32%
Slightly Important	15%
Not at all Important	8%

The data in **Table 36** shows that almost 50% of the **Survey Respondents** indicated that getting better at managing VoIP quality is either very or extremely important to their IT organization.

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

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

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

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

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


The Emerging Management Environment

The Evolving Focus on Services

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

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

To quantify this shift in thinking on the part of IT organizations, the **Survey Respondents** were asked to indicate how important it was over the next year for their organization to get better at managing a business service, such as ERP, that is supported by multiple, interrelated applications. Their responses are shown in **Figure 31**.




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

applications that comprise a service. As was the case with SLAs, that follows because the adoption of cloud computing in general and of virtualization in particular, creates significant management challenges.

Another way to measure this transformation is to identify how IT organizations currently focus their management efforts. To that end, the **Survey Respondents** were asked to indicate the approach their organization takes to service or performance management. They were given the following choices and allowed to choose all that applied to their environment.

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

Their responses are summarized in **Figure 32**.



The data in **Figure 32** indicates that the most frequent approach that IT organizations take to management is to focus on individual technology domains. However:

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

Service Delivery Management

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

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

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

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


One of the reasons why it is important to get better at managing the end-user's experience was highlighted in The 2012 Application and Service Delivery Handbook⁴³. That handbook presented market research that highlighted the fact that in spite of all of the effort that has gone into implementing IT management to date, that it is the end user, and not the IT organization who typically is the first to notice when the performance of an application begins to degrade.

The data in **Figure 33** demonstrates the growing importance that IT organizations place on managing end-user experience. That figure shows the results of a question in which the **Survey Respondents** were asked how important it was over the next year for their organization to get better at monitoring the end-user's experience and behavior. As shown in **Figure 33**,

⁴³ <http://www.webtorials.com/content/2012/08/2012-application-service-delivery-handbook-2.html>

getting better at managing end-user's experience is either very or extremely important to roughly half of all IT organizations.

Figure 33: The Importance of Getting Better at Managing the End-User's Experience



Dynamic Infrastructure Management

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

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

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

Virtualized Performance and Fault Management

In a traditional IT environment it is common to implement adaptive performance thresholding solutions that can identify systemic deviations from normal patterns of behaviour as well as time over threshold violations and can also automatically update thresholds based on changes to

historic levels of utilization. As previously discussed, that same capability is needed in a virtualized environment so that IT organizations can monitor the performance of individual VMs.

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

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

Converged Infrastructure Management

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

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

- **Integrated and Automated Infrastructure and Service Management:** Integrated management reduces the number of management interfaces that are involved in implementing administrative workflows. Automation allows services to be dynamically provisioned, modified or scaled without requiring time-consuming manual configuration across the various technology domains of the data center; e.g., compute, network, storage and security. The management suite should also include the application and service level management capabilities that will support end-to-end SLAs. From an operational management perspective, the management system should provide additional capabilities, such as cross-domain root cause analysis and service impact analysis, to support the highest levels of service reliability.
- **Secure Multi-tenancy:** A robust multi-layer security architecture is required to ensure confidentiality and integrity of the services and the subscriber's data, particularly in a multi-tenant environment.

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

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

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

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

Culture eats strategy for breakfast.

That cliché refers to the fact that in many cases the culture of an IT organization resists any changes that involve changing the roles of the members of the organization. Exacerbating the challenge of the IT organization's resistance to change is the fact that, as was pointed out in the section of this report entitled *The Emergence of Cloud Computing and Cloud Networking*, the culture of an IT organization typically changes very slowly.

Orchestration and Provisioning

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

virtualized services. Orchestration engines are available as standalone management products or as part of complete suites of management tools that are focused on the data center. In addition, the management systems that are integrated with converged infrastructure solutions typically include some orchestration capabilities.

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

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

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

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

Orchestration solutions would benefit greatly from the emergence of an open standard for the exchange of information among the full range of devices that may be used to construct a dynamic virtual data center. In the Cloud Computing arena there are a number of standards under development, including the Open Cloud Computing Interface (OCCI) from the Open Grid Forum⁴⁴. 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 Group⁴⁵ and implemented by a number of companies in the security and network industries. It is a publish/subscribe protocol that allows hosts to lookup meta-data and to subscribe to service or host-specific event notifications. IF-MAP can enable auto-discovery and self-assembly (or re-assembly) of the network architecture. As such, IF-MAP has the potential to support the automation and dynamic orchestration of not only security systems, but also other elements of the virtual data center. For example, IF-MAP could facilitate the automation of the processes associated with virtual machine provisioning and deployment by publishing all of the necessary

⁴⁴ <http://www.gridforum.org/>

⁴⁵ <http://www.trustedcomputinggroup.org/>

policy and state information to an IF-MAP database that is accessible by all other elements of the extended data center.

Application Performance Management

Impediments

Application performance management has been deployed for several years and yet only a small percentage of the **Survey Respondents** indicated that their organization did a good job of managing application performance. To understand why IT organizations are not more successful with application performance management, the **Survey Respondents** were asked to indicate the two primary impediments to their organization being more successful with application performance management. The impediments and the percentage of the **Survey Respondents** who indicated that the impediment was one of the two primary impediments to successful application performance management are shown in **Table 37**.

Table 37: Impediments to Successful Application Performance Management	
Impediment	Percentage of the Survey Respondents
Our organization tends to be more reactive than proactive	33%
We focus too much on managing technology domains and not enough on managing business transactions	32%
The tools we use don't give us an end-to-end view of the user's experience	29%
The various sub-groups within the IT organization don't work effectively to identify and resolve problems	26%
We don't have the ability to manage the performance of applications and services acquired from cloud service providers	17%
The tools we use don't allow us to perform rapid root cause analysis	14%
The tools we use don't give us the ability to link the performance of a transaction as seen by the user with all of the various applications that comprise that application	13%
The tools we use don't give us the ability to link the performance of a transaction as seen by the user with the components of the infrastructure that support those transactions	13%
We don't have the ability to gather management data across both the physical and the virtual components of the infrastructure	12%
Other	10%

One observation that can be drawn from the data in **Table 37** is that there isn't a single impediment that is the primary reason why IT organizations aren't successful with application performance management. Rather, there is a wide range of impediments that limit the ability of IT organizations to be successful with application performance management. Another observation is that

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

A Top Down Approach

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

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

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

Root Cause Analysis

As previously mentioned, one of the questions (The Question) that was administered to the **Survey Respondents** was “Please indicate how important it is to your organization to get better at each of the following tasks over the next year.” The question included twenty wide-ranging management tasks. **Table 38** lists the three management tasks that were the most important to the **Survey Respondents** and the percentage of the **Survey Respondents** that indicated that getting better at those tasks was either very or extremely important.

Table 38: Primary Management Challenges

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

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


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

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

Designing for Application Performance

One of the traditional challenges to effective application performance is that most IT organizations don't place much emphasis on application performance during application development. To quantify that phenomenon, the **Survey Respondents** were asked "When your IT organization is in the process of either developing or acquiring an application, how much attention does it pay to how well that application will perform over the WAN?" Their answers are shown in **Figure 34**.

Figure 34: The Emphasis on Performance over the WAN



The data in **Figure 34** shows that almost three quarters of all IT organizations place at most moderate emphasis on performance while either developing or acquiring an application.

The lack of emphasis on an application's performance over the WAN during application development often results in the development and implementation of applications that run poorly once they are placed into production. One of the reasons for that phenomenon is that due to factors such as chatty protocols (**Figure 35**), an application can run well over a high-speed, low latency LAN in a development environment but run poorly over a relatively low-speed, high latency WAN in a production environment.

Figure 35: Chatty Protocol



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

Application Performance Engineering

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

Application Performance Engineering (APE) is the practice of first designing for acceptable application performance and then testing, measuring and tuning performance throughout the application lifecycle.

During the operational, or production phase of the lifecycle, application performance management is used to monitor, diagnose, and report on application performance. Application performance management and APE are therefore highly complementary disciplines. For example, once an application performance management solution has identified that an application in production is experiencing systemic performance problems, an APE solution can be used to identify the root cause of the problem and to evaluate alternative solutions. Possible solutions include modifying the application code or improving application performance by making changes in the supporting infrastructure, such as implementing more highly performing servers or deploying WAN Optimization Controllers (WOCs). Throughout this section of The Report, implementing products such as WOCs will be referred to as a Network and Application Optimization (NAO) solution. Independent of which remedial option the IT organization takes, the goal of APE can be realized – performance bottlenecks are identified, root causes are determined, alternative remedies are analyzed and bottlenecks are eliminated.

An IT organization could decide to ignore APE and just implement NAO in a reactive fashion in an attempt to eliminate the sources of the degraded application performance. Since this

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

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

- **Setting Performance Objectives**

This involves establishing metrics for objectives such as user response time, transaction completion time and throughput. A complex application or service, such as unified communications, is comprised of several modules and typically different objectives need to be established for each module.

- **Discovery**

Performance modeling and testing should be based on discovering and gaining a full understanding of the topology and other characteristics of the production network.

- **Performance Modeling**

APE modeling focuses on creating the specific usage scenarios to be tested as well as on identifying the performance objectives for each scenario. A secondary focus is to identify the maximum utilization of IT resources (e.g., CPU, memory, disk I/O) and the metrics that need to be collected when running the tests.

- **Performance Testing and Analysis**

Test tools can be configured to mimic the production network and supporting infrastructure, as well as to simulate user demand. Using this test environment, the current design of the application can be tested in each of the usage scenarios against the various performance objectives. The ultimate test, however, is measured performance in the actual production network or in a test environment that very closely mimics the actual production environment.

- **Optimization**

Optimization is achieved by identifying design alternatives that could improve the performance of the application and by redoing the performance testing and analysis to quantify the impact of the design alternatives. In conjunction with the testing, an ROI analysis can be performed to facilitate cross-discipline discussion of the tradeoffs between business objectives, performance objectives, and cost. This component of APE is one of the key ways that APE enables an IT organization to build better relationships with the company's business unit managers.

Application Performance Management Tools

Enterprise IT organizations can choose among several types of tools for monitoring and managing application performance over a private enterprise network. These include:

application agents, monitoring of real and synthetic transactions, network flow and packet capture, analytics, and dashboard portals for the visualization of results.

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

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

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

Management as a Cloud Provided Service

As pointed out in the section of The Report entitled *The Emergence of Cloud Computing and Cloud Networking*, a new class of solutions has begun to be offered by CCSPs. These are solutions that have historically been provided by the IT infrastructure group itself and include VoIP, network management, security, network and application optimization, application performance management, Unified Communications (UC) and virtualized desktops. This new class of solutions is referred to as [Cloud Networking Services](#) (CNS). That section of The Report also presented the results of a survey in which The **Survey Respondents** were asked to indicate the CNSs that their organization currently acquires from a CCSP and the CNSs that they would like acquire from a CCSP in the next year. Their responses are shown in **Table 39**.

Table 39: Current and Planned Adoption of CNSs			N = 142
	Currently Acquire	Will Likely Acquire	
VoIP	20.4%	17.6%	
Network Management	19.7%	8.5%	
Security	18.3%	9.9%	
Unified Communications	15.5%	23.2%	
Application Performance Management	10.6%	10.6%	
Network and Application Optimization	8.5%	9.2%	
Virtual Desktops	7.0%	19.0%	

The data in **Table 39** shows that

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

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

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

Security

The Current Environment for Security Breaches

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

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

The shift in the security landscape has been documented in a number of reports. For example, IBM's X-Force 2011 Trend and Risk Report⁴⁶ made a number of observations relative to the current environment for security breaches. Some of the key observations made in that report are:

- **Mobile Devices**

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

- **Social Media**

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

- **Cloud Computing**

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

⁴⁶ [X-Force 2011 Trend and Risk Report](#)

Blue Coat Systems' 2012 Web Security Report⁴⁷ also made a number of observations relative to the current environment for security breaches. According to the Blue Coat report, "In 2011, malnets emerged as the next evolution in the threat landscape. These infrastructures last beyond any one attack, allowing cybercriminals to quickly adapt to new vulnerabilities and repeatedly launch malware attacks. By exploiting popular places on the Internet, such as search engines, social networking and email, malnets have become very adept at infecting many users with little added investment." That report also noted the increasing importance of social networking and stated that, "Since 2009, social networking has increasingly eclipsed web-based email as a method of communications" and that, "Now, social networking is moving into a new phase in which an individual site is a self-contained web environment for many users – effectively an Internet within an Internet."

The Current Environment for Implementing Security

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

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

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

In order to determine how IT organizations are responding to the traditional and emerging security challenges, the **Survey Respondents** were asked a series of questions. For example, to get a high level view of how IT organizations are providing security, the **Survey Respondents** were asked to indicate which of a number of network security systems their organization supports. The **Survey Respondents** were asked to check all of the alternatives that applied in their environment. Their responses are shown in **Table 40**.

⁴⁷ http://www.bluecoat.com/sites/default/files/documents/files/BC_2012_Security_Report-v1i-optimized.pdf

Table 40: The Network Security Systems in Use	
Network Security Systems	Percentage
Remote Access VPN	86.30%
Network Access Control	73.50%
Intrusion Detection/Protection Systems (IDS/IPS)	65.70%
Next Generation Firewalls (Firewall+IPS+Application Control)	56.90%
Secure Web Gateways	46.10%
Web Application and/or XML Firewalls	36.30%
Mobile Device Security/Protection	36.30%
Security Information Event Management	31.40%
Data Loss Prevention	24.50%
Password Vault Systems (either local or portal based)	12.70%
SAML or WS-Federation Federated Access Control	8.80%

One obvious conclusion that can be drawn from **Table 40** is that IT organizations use a wide variety of network security systems. A slightly less obvious conclusion is that

On average, IT organizations use 4.8 network security systems.

The **Survey Respondents** were asked to indicate the approach that best describes how their company uses data classification to create a comprehensive IT security environment. Their responses are shown in **Table 41**.

Table 41: Approach to Comprehensive IT Security	
Approach	Percentage
We have a data classification policy and it is used to determine application access/authentication, network and end user device security requirements.	42.90%
We do not have a data classification policy.	33.00%
We have a data classification policy and it is used to determine application security requirements.	13.20%
We have a data classification policy, but it is not used nor enforced.	11.00%

The data in **Table 41** represents a classic good news/bad news situation. The good news is that the majority of IT organizations have a data classification policy that they use to determine requirements. The bad news is that

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

In order to understand how IT organizations are responding to the BYOD movement, the **Survey Respondents** were asked, "If your organization does allow employee owned devices to connect to your network, please indicate which of the following alternatives are used to register employee owned devices and load authentication (e.g. certificate/private key) data onto those devices before they are allowed to connect to your company's network." The **Survey Respondents** were asked to check all of the alternatives that applied in their environment. Their responses are shown in **Table 42**.

Table 42: Alternatives to Support Employee Owned Devices	
Alternative	Percentage
Employees must install a VPN client on their devices for network access	53.90%
IT Administrator and/or Service Desk must register employee owned device for network access	47.40%
Employees can self-register their devices for network access	28.90%
Employees must generate and/or load X.509 certificates & private keys network access	13.20%
Employees must install a token authentication app on their devices for network access	10.50%

The data in **Table 42** indicates that while using a VPN is the most common technique that a wide range of techniques are used. VPN's popularity comes in part from the fact that remote access VPN solutions implemented on new generation mobile devices have various capabilities to enforce security policies when connecting to the corporate network. Popular security checks include ensuring that a screen password is present, that anti-virus software is present and is up to date, that there is not rogue software on the device and that the operating system has not been modified.

Two different approaches have emerged to protect against lost devices. For the traditional PC, full disk encryption is typically used to protect data if the PC is lost or stolen. However, on new generation mobile devices, remote erase solutions are typically used to protect data. New generation mobile devices with smaller displays are often used more for content reading rather than content creation. As screen sizes and resolution improves, this situation may change. In order to understand how IT organizations have implemented full disk encryption, the **Survey Respondents** were asked to indicate which alternatives their organization implements relative to using full disk encryption on laptops and desktop PCs. Their responses are shown in **Table 43**.

Table 43: Techniques for Implementing Full Disk Encryption	
Alternative	Percentage
We do not use full disk encryption on PCs.	52.5%
We use software based disk encryption on PCs.	49.5%
We use hardware based self-encrypting rotating drives on PCs.	6.1%
We use hardware based self-encrypting Solid State Drives on PCs.	6.1%

The data in **Table 43** indicates that

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

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

The **Survey Respondents** were asked to indicate the approach that best describes their company's approach to Identity and Access Management (IAM). Their responses are shown in **Table 44**.

Table 44: How IAM is Implemented	
Approach	Percentage
We do not have a formal IAM program.	36.6%
We have an IAM program, but it only partially manages identities, entitlements and policies/rules for internal users.	25.8%
We have an IAM program and it manages identities, entitlements and policies/rules for all internal users.	20.4%
We have an IAM program and it manages identities, entitlements and policies/rules for end users for internal, supplier, business partner and customers.	17.2%

The data in **Table 44** indicates that only a minority of IT organizations has a IAM program that has broad applicability.

The **Survey Respondents** were asked to indicate how their company approaches the governance of network and application security. Their responses are shown in **Table 45**.

Table 45: Governance Models in Use	
Approach	Percentage
Network Security and Application Security are funded, architected, designed and operated together.	46.9%
Network Security and Application Security are funded, architected, designed and operated separately.	30.2%
Network Security and Application Security are funded jointly, but architected, designed and operated separately.	22.9%

The data in **Table 45** indicates that

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

Security as a Cloud Provided Service

As previously mentioned, IT organizations have shown a great interest in acquiring from CCSPs a wide range of functionality that historically has been provided by the IT infrastructure group; a.k.a., cloud networking services (CNS). This includes security. In particular, as was also previously discussed (**Table 39**), over a quarter of the **Survey Respondents** indicated that their company either currently acquires security functionality from a CCSP or they expect that their company will within the next year.

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

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

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

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

One part of the value proposition of a CBSS is the value proposition of any cloud based service. For example, a CBSS reduces the capital investment in security that an organization would have to make. In addition, a CBSS reduces the amount of time it takes to deploy new functionality. The speed at which changes can be made to a CBSS adds value in a variety of situations, including providing better protection against zero-day attacks⁴⁸. Another part of the value proposition of a CBSS is that unlike a traditional security solution that relies on the implementation of a hardware based proxy, a CBSS can also protect mobile workers. The CBSS does this by leveraging functionality that it provides at its POPs as well as functionality in a software agent that is deployed on each mobile device. The use of a Cloud-based solution to provide mobile device management and security was discussed previously in this section.

In many instances, the best security solution is a hybrid solution that combines traditional on-premise functionality with one or more Cloud-based solutions. For example, in many cases IT organizations already have functionality such as web filtering or malware protection deployed in CPE at some of their sites. In this case, the IT organization may choose to implement a CBSS

⁴⁸ http://en.wikipedia.org/wiki/Zero-day_attack

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

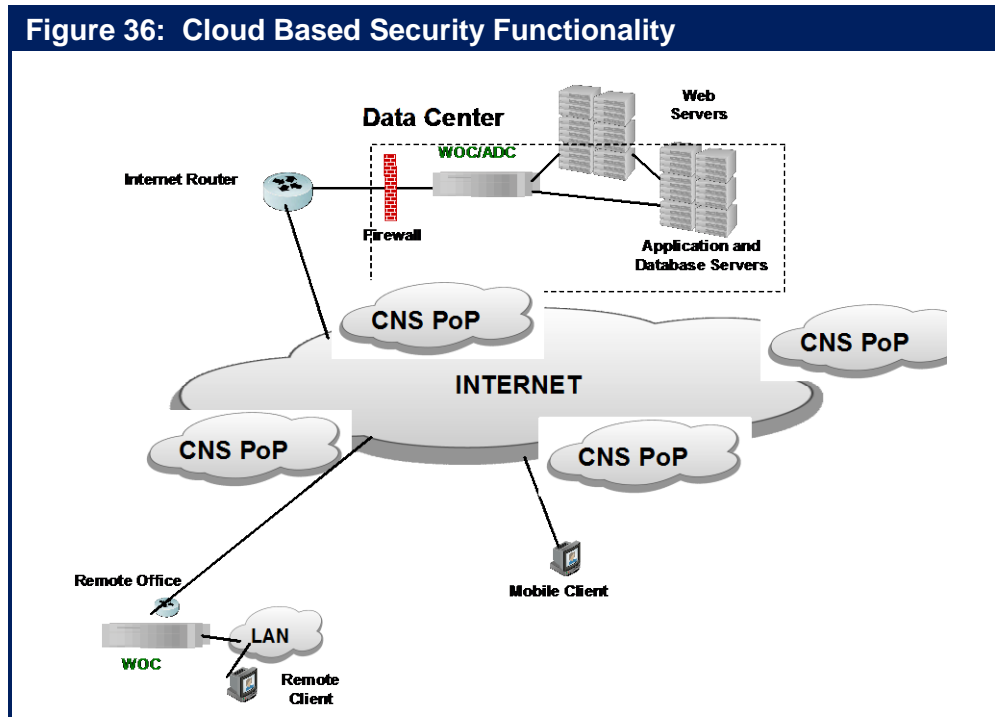
Other situations in which a CBSS can serve to either be the only source of security functionality, or to compliment CPE based implementations include cloud-based firewall and cloud-based IPS services. Such a service should support equipment from the leading vendors. Given the previously mentioned importance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

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

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

Web Application Firewall Services

The section of this report entitled *Wide Area Networking*, discussed how a Cloud-based service, such as the one shown in **Figure 36**, can be used to optimize the performance of the Internet. As will be discussed in this sub-section of the handbook, that same type of service can also provide some CCSBs.



The Role of a Traditional Firewall

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

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

Whereas network firewalls are focused on parameters such as IP address and port numbers, a more recent class of firewall, referred to as a Web application firewall, analyzes messages at layer 7 of the OSI model. Web application firewalls are typically deployed as a hardware appliance and they sit behind the network firewall and in front of the Web servers. They look for violations in the organization's established security policy. For example, the firewall may look for abnormal behavior, or signs of a known attack. It may also be configured to block specified

content, such as certain websites or attempts to exploit known security vulnerabilities. Because of their ability to perform deep packet inspection at layer 7 of the OSI model, a Web application firewall provides a level of security that cannot be provided by a network firewall.

The Role of a Web Application Firewall Service

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

In the current environment, high-end DDoS attacks can generate 100 Gbps of traffic or more⁴⁹. Attacks of this magnitude cannot be prevented by onsite solutions. They can, however, be prevented by utilizing a CBSS that includes security functionality analogous to what is provided by a Web application firewall and that can identify and mitigate the DDoS-related traffic close to attack traffic origin.

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

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

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

- **Minimizing the points of vulnerability**
If an organization has most or all of its important assets in a small number of locations, this makes the organization more vulnerable to successfully being attacked as the attacker has fewer sites on which to concentrate their attack.
- **Protecting DNS**
Many IT organizations implement just two or three DNS servers. As such, DNS is an example of what was discussed in the preceding bullet – how IT organization are vulnerable because their key assets are located in a small number of locations.

⁴⁹ [DDoS-attacks-growing-in-size](#)

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

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

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

- Block or redirect requests based on characteristics such as the originating geographic location and whether or not the originating IP addresses are on either a whitelist or a blacklist.
- Direct traffic away from specific servers or regions under attack.
- Issue slow responses to the machines conducting the attack. The goal of this technique, known as tarpits⁵⁰, is to shut down the attacking machines while minimizing the impact on legitimate users.
- Direct the attack traffic back to the requesting machine at the DNS or HTTP level.

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

⁵⁰ [Wikipedia Tarpit\(networking\)](#)

Conclusions and Observations

Throughout the 2012 Cloud Networking Report the following conclusions were drawn and observations were made.

- The phrase cloud networking refers to the LAN, WAN and management functionality that must be in place to enable cloud computing.
- In order to support cloud computing, a cloud network must be dramatically more agile and cost effective than a traditional network.
- 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.
- 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.
- SLAs from both traditional network service providers as well as public cloud computing providers are a work in progress.
- Roughly half of all IT organizations are currently in the process of developing a strategy for how they will use public and private IaaS solutions.
- Concern about the security and confidentiality of data is the primary impediment to the broader adoption of private IaaS solutions.
- The SaaS marketplace is comprised of a small number of large players such as Salesforce.com, WebEx and Google Docs as well as thousands of smaller players.
- The primary factors that are driving the adoption of SaaS are the same factors that drive the adoption of any form of out-tasking.
- There is strong interest on the part of IT organizations in acquiring both virtual private data center and disaster recovery services from IaaS providers.
- By a wide margin, agility is the most important factor driving the adoption of Cloud-based IaaS solutions.
- Concern about the security and confidentiality of data is by a wide margin the number one factor inhibiting the adoption of Cloud-based IaaS solutions.
- There is a strong desire on the part of IT organizations to manage the security related network services that are part of an IaaS service.
- The evaluation of the supporting network services is a key component of the overall process of evaluating IaaS solutions.

- Roughly 20% of the times that a company is evaluating public IaaS solutions, the company's IT organization is either not involved at all or plays a minor role.
- Cloud balancing can be thought of as the logical extension of global server load balancing (GSLB).
- Cloud Networking Services represents the beginning of what could be a fundamental shift in terms of how IT services are provided.
- One way for an IT organization to evaluate the agility of a CCSP is to identify the degree to which the CCSP has virtualized their infrastructure.
- IT organizations provide considerable value by being the broker between the company's business unit managers and cloud computing service providers.
- The culture of an IT organization changes very slowly.
- The primary factors driving IT organizations to re-design their data center LAN is the desire to reduce cost, support server virtualization and reduce complexity.
- One approach for improving server-to-server communications is to flatten the network from three tiers to two tiers consisting of access layer and aggregation/core layer switches.
- The 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.
- 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.
- 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.
- The combination of switch virtualization and multi-chassis LAG can be used to create a logically loop-free topology
- In many cases, the best technology doesn't end up being the dominant technology in the marketplace.
- With technologies like TRILL and SPB, the difference between access switches and core switches may shrink significantly.
- There is very strong interest on the part of IT organizations to implement network virtualization.
- A possible characteristic of Third Generation Data Center LANs is the convergence of block-level storage and data traffic over a common high-speed Ethernet data center switching fabric.

- There are several levels of support that data center switch vendors can provide for Fibre Channel over Ethernet (FCoE).
- The primary drivers of FCoE are the vendors that offer both Ethernet and Fibre Channel products.
- Most enterprise IT organizations have little if any knowledge of SDN.
- The vast majority of IT organizations that understand SDN believe that OpenFlow is an important component of an SDN.
- There is not a consensus amongst IT organizations about whether or not SDN will relegate switches and routers to be just dumb forwarding engines.
- IT organizations believe the primary value that SDN offers in the data center is that it can help IT organizations to reduce costs, automate management, and enforce security policies.
- One of the key promises of SDN is that developer communities will be created and that these communities will offer a wide range of applications.
- The majority of IT organizations believe that implementing SDN will make networks more secure.
- The primary inhibitor to SDN adoption is the overall confusion in the market and the immaturity of products and vendor strategies.
- Today there is not a fundamentally new generation of technology under development that is focused on the WAN.
- The WAN doesn't follow Moore's Law.
- WAN budgets are notably more constrained than they were a year ago.
- IT organizations need to make changes relative to how they use WAN services in order to support a significant increase in WAN traffic while experiencing a highly constrained WAN budget.
- Over the next year, the percentage of IT organizations that have not implemented any desktop virtualization will be cut roughly in half.
- IT organizations are required to support a wide range of end user devices.
- The primary WAN services used by IT organizations are MPLS and the Internet.
- The primary concerns that IT organizations have with the use of MPLS are cost, the lead time to implement new circuits and uptime. The primary concerns that IT organizations have with the use of the Internet are uptime, latency and cost.

- One viable approach to WAN design is to use both the Internet and MPLS in ways that maximize the benefits of each while minimizing their deficiencies.
- In a growing number of instances, Internet-based VPNs that use DSL for access are 'good enough' to be a cloud network.
- The key concept behind an aggregated virtual WAN is that it simultaneously utilizes multiple enterprise WAN services and/or Internet connections in order to optimize reliability and minimize packet loss, latency and jitter.
- Some of the concerns that IT organizations have about the use of the Internet are exacerbated by backhauling Internet traffic to a central site.
- Over the next year, IT organizations will make an increased use of distributed access to the Internet from their branch offices.
- In roughly forty percent of the instances that business users are accessing public cloud computing services, the primary WAN service that they intend to use is not the traditional Internet.
- In almost two thirds of the instances that business users are accessing private cloud computing services, the primary WAN service that they intend to use is not the traditional Internet.
- The majority of IT organizations don't regard the SLAs that they receive from their network service providers as being effective.
- The majority of IT organizations believe that factors such as the growth in the number of mobile workers and the increase in the use of virtualization and cloud computing will make ensuring acceptable service and application delivery either harder or notably harder.
- An ADC provides more sophisticated functionality than an SLB does.
- 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.
- Virtual appliances make is easier to conduct a proof of concept trial.
- In many instances the benefits of the dynamic movement of a VM from one server to another are maximized if the supporting infrastructure is virtualized and can also be dynamically moved.
- The majority of IT organizations are either undecided about how they will optimize the performance of IaaS services or they intend to do nothing.
- In many situations, a dual ISP-based Internet VPN with PBR can deliver a level of CoS and reliability that is comparable to that of MPLS at a significantly reduced price.

- Slightly over a third of all companies, and slightly over a half of large companies have an effective WAN strategy.
- IT organizations need to adopt an approach to management that is based on the assumption that the components of a service, and the location of those components, can and will change frequently.
- Almost half of the IT organizations consider it to be either very or extremely important over the next year for them to get better performing management tasks such as troubleshooting on a per-VM basis.
- The adoption of cloud computing makes troubleshooting application performance an order of magnitude more difficult than it is in a traditional environment.
- A fundamental issue relative to managing either a public or hybrid cloud computing service is that the service has at least three separate management domains: the enterprise, the WAN service provider(s) and the various cloud computing service providers.
- IT organizations need to adopt an approach to IT management that is based on gathering management data across myriad data centers, including ones that are owned and operated by a third party.
- A majority of IT organizations believe that getting better at managing all forms of cloud computing solutions is at least moderately important.
- The vast majority of IT organizations provide an internal SLA for at least some applications, but that only half of all IT organizations are successful managing those SLAs.
- Two thirds of IT organizations believe that it is either very or extremely important to get better at effectively managing internal SLAs.
- 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.
- IT organizations need to adopt an approach to management in which one set of tools is used to manage enterprise data applications as well as voice, video and complex interrelated applications.
- A significant percentage of IT organizations focus their management activities on the performance of applications and/or services.
- IT organizations should adopt an approach to service delivery management that is unified across the various IT domains so that IT organizations have visibility across all of the applications, services, locations, end users and devices.
- 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.

- Culture eats strategy for breakfast.
- 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.
- Organizational impediments are more likely to limit an IT organization's success with application performance management than are technical impediments.
- Only a small minority of IT organizations has a top down, tightly coordinated approach to application performance management.
- Application Performance Engineering (APE) is the practice of first designing for acceptable application performance and then testing, measuring and tuning performance throughout the application lifecycle.
- IT organizations have a significant interest in acquiring network management functionality for a cloud service provider.
- Over the last few years, the sophistication of hackers has increased by an order of magnitude.
- On average, IT organizations use 4.8 network security systems.
- Almost half of all IT organizations either don't have a data classification policy or they have one that isn't used or enforced.
- Just over half of all IT organizations don't use full disk encryption on PCs.
- In the majority of instances, network security and application security are architected, designed and operated separately.
- A cloud-based security service needs to be able to allow access to a social media site such as Facebook, but block specific activities within the site, such as gaming or posting.

About the Webtorials® Editorial/Analyst Division

The Webtorials® Editorial/Analyst Division, a joint venture of industry veterans Steven Taylor and Jim Metzler, is devoted to performing in-depth analysis and research in focused areas such as Metro Ethernet and MPLS, as well as in areas that cross the traditional functional boundaries of IT, such as Unified Communications and Application Delivery. The Editorial/Analyst Division's focus is on providing actionable insight through custom research with a forward looking viewpoint. Through reports that examine industry dynamics from both a demand and a supply perspective, the firm educates the marketplace both on emerging trends and the role that IT products, services and processes play in responding to those trends.

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

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

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

Professional Opinions Disclaimer

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

Copyright © 2012, Webtorials

For editorial and sponsorship information, contact Jim Metzler or Steven Taylor. The Webtorials Editorial/Analyst Division is an analyst and consulting joint venture of Steven Taylor and Jim Metzler.



Application Performance for Business Efficiency

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

82% *

of organizations suffer application performance problems.

63% *

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

72% *

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

Business and IT performance are tightly coupled...

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

**Ipanema Killer Apps survey 2012*

IT departments are witnessing change at a pace never seen before

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

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

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

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

Business efficiency requires guaranteed application performance

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

With Ipanema, control all your IT transformations!



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

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

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

What our customers say about us:

Do more with less

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

Guarantee Unified Communications and increase network capacity

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

Reduce costs in a cloud environment

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

