The 2011 Cloud Networking Report

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

Produced by:

Webtorials

Part 2: The Emerging Data Center LAN

Featured Sponsors of Part 2:

Alcatel·Lucent Enterprise
AVAYA The Power of We™
CISCO
DELL Force10
extreme networks™ Make Your Network Mobile
hp
IBM
LSI

Additional Sponsors:

A10 Networks
Blue Coat
ca
certeon Accelerate Your Business
CITRIX
exinda
Infoblox
ipanema Technologies
NETSCOUT
Packet Design
radware
riverbed
Silver Peak
TALARI NETWORKS
VISUAL NETWORK SYSTEMS
VYATTA Open Networking

Table of Contents

EXECUTIVE SUMMARY 1

THE EMERGING DATA CENTER LAN 2

FIRST AND SECOND GENERATION DATA CENTER LANS 2

Drivers of Change 3

THIRD GENERATION DATA CENTER LAN ARCHITECTURE AND TECHNOLOGY OPTIONS 7

Two Tier Data Center LAN Design 8

Controlling and Managing Inter-VM Traffic 15

Software Defined Networks and Network Virtualization 17

Network Convergence and Fabric Unification 18

Network Support for the Dynamic Creation and Movement of VMs 22

Summary of Third Generation Data Center LAN Technologies 25

Executive Summary

The **2011 Cloud Networking Report** will be published both in its entirety and in a serial fashion. This is the second of the serial publications. One goal of this publication is to provide a very brief overview of how data center LAN technology and design has evolved and to identify the factors that are currently driving the vast majority of IT organizations to rethink how they design their data center LANs. Another goal of this publication is to provide insight into the technologies and design choices that IT organizations are making. The third and primary goal of this publication is to describe the data center LAN architecture and technology options that either are currently available in the market or are likely to be available within two years.

Given the breadth of fundamental technology changes that are impacting the data center LAN, this section is very technical.

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

The Webtorials Respondents were asked "Has your IT organization already redesigned, or within the next year will it redesign, its data center LAN in order to support cloud computing in general, and virtualized servers in particular?" Their responses are shown in Table 1.

Table 1: Redesign of the Data Center LAN

	Already Have	Will Within the Next Year	No Plans
Cloud Computing in General	21.8%	51.1%	27.1%
Virtualized Servers in Particular	53.7%	34.0%	12.2%

One conclusion that can be drawn from the data in Table 1 is that the majority of IT organizations have already begun the process of redesigning their data center LANs. Another conclusion is that:

> ***One of the key factors driving IT organizations to redesign their data center LANs is the deployment of virtual servers.***

In order to quantify the interest that IT organizations have in implementing server virtualization, The Webtorials Respondents were asked to indicate the percentage of their company's data center servers that have either already been virtualized or that they expected would be virtualized within the next year. Their responses are shown in Table 2.

Table 2: Deployment of Virtualized Servers

	None	1% to 25%	26% to 50%	51% to 75%	76% to 100%
Have already been virtualized	15%	33%	21%	18%	14%
Expect to be virtualized within a year	6%	25%	28%	20%	20%

In early 2010, the Webtorials survey base was asked to indicate the percentage of their data center servers that had already been virtualized. Their responses are shown in Table 3.

Table 3: Deployment of Virtualized Servers as of Early 2010

	None	1% to 25%	26% to 50%	51% to 75%	76% to 100%
Have already been virtualized	30%	34%	17%	11%	9%

The data in Table 2 and Table 3 show the strength of the ongoing movement to virtualize data center servers. For example, in early 2010 20% of IT organizations had virtualized the majority of their data center servers. Today, 32% of IT organizations have virtualized the majority of their data centers servers. In addition, The Webtorials Respondents predict that within a year, that 40% of IT organizations will have virtualized the majority of their data center servers. Another way to look at the data in Table 2 and Table 3 is that in early 2010 30% of IT organizations had not virtualized any data center servers. Today, only 15% of IT organizations have not virtualized any data center servers and The Webtorials Respondents predict that within a year, that only 6% of IT organizations will not have virtualized any of their data center servers.

As pointed out in Virtualization: Benefits, Challenges and Solutions[1], 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 Webtorials Respondents were asked to indicate whether they agreed or disagreed with the statements in the left hand column of Table 4.

Table 4: Movement of VMs

	Agree	Disagree
We currently manually migrate VMs between servers in the same data center	67.4%	32.6%
We currently automatically migrate VMs between servers in the same data center	55.9%	44.1%
We currently manually migrate VMs between servers in disparate data centers	42.6%	57.4%
We currently automatically migrate VMs between servers in disparate data centers	26.6%	73.4%

The data in Table 4 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.

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**

[1] http://www.webtorials.com/content/2010/06/virtualization.html

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 5 contains a list of the factors that are driving data center redesign. The center column shows the percentage of The Interop Respondents who in the fall of 2010 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 Webtorials Respondents who recently indicated that the corresponding factor was the primary factor that is driving their organization to redesign their data center LAN.

Table 5: Factors Driving Data Center LAN Redesign

Factor	Percentage of The Interop Respondents in 2010	Percentage of The Webtorials Respondents in 2011
To reduce the overall cost	22.4%	24.6%
To support more scalability	11.6%	20.8%
To create a more dynamic data center	11.6%	12.6%
To support server virtualization	11.2%	12.1%
To reduce complexity	9.9%	5.3%
To make it easier to manage and orchestrate the data center	9.2%	13.0%
To support our storage strategy	7.5%	3.4%
To reduce the energy requirements	6.5%	1.0%
Other (please specify)	6.1%	3.4%
To make the data center more secure	4.1%	3.9%

The data in Table 5 indicates that a broad range of factors are driving IT organizations to re-design their data center LANs. For example, making it easier to manage and orchestrate the data center is becoming a key driver in how IT organizations design their data center LANs. However, as was the case with the adoption of the second generation of data center LANs:

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

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 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 Webtorials 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 1.

Figure 1: Planning Horizon for Data Center LANs

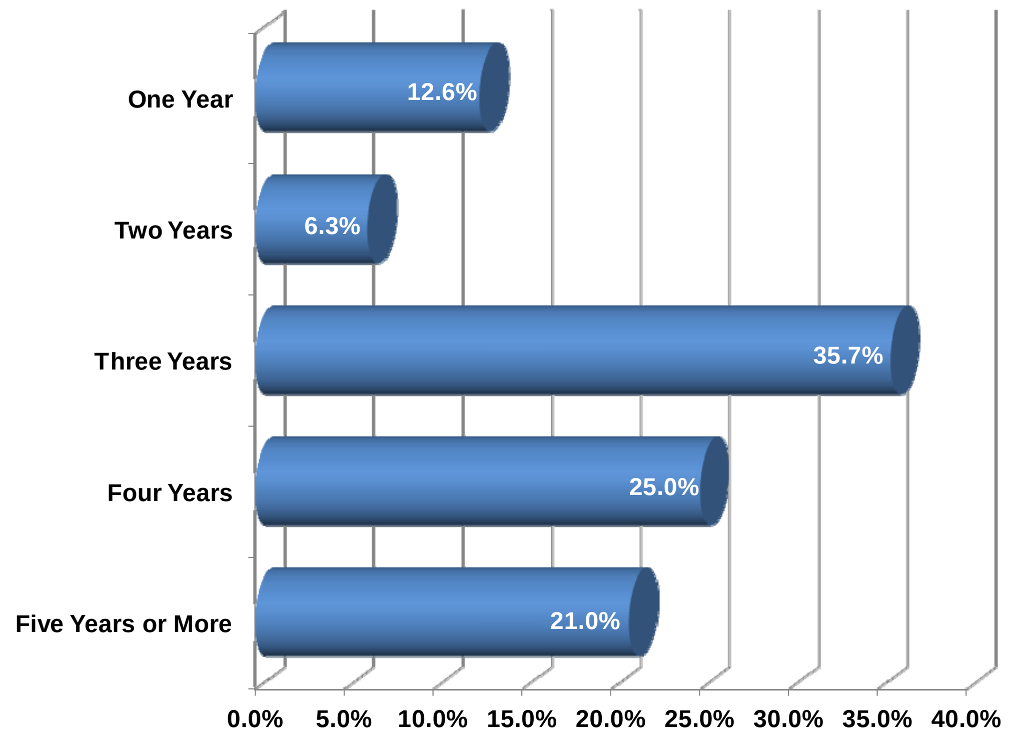

The data in Figure 1 indicates that almost 75% of IT organizations have a planning horizon of three years or longer. Since most of the technologies discussed in this chapter will be standardized and ready for production use in three years, that means that the vast majority of IT organizations can incorporate 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, 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.

The Interop Respondents were asked, “Two years from now, what is the fewest number of layers that you expect will be in any of your company’s data center LANs.” The answers of one hundred and ninety respondents are summarized in Table 6.

Table 6: Anticipated Number of Layers in a Data Center LAN

Number of Layers	Percentage of Respondents
4	8%
3	37%
2	38%
1	17%

The data in Table 6 indicates that while just over a third of IT organizations expect to still be running traditional, three-tier data center LANs in two years, the majority of the IT professionals who answered the question expect to be running a flatter data center LAN in that time frame. However, what is even more interesting is that two hundred and sixty five members of the pool of survey respondents answered the question with “don’t know”. That means that the number of survey respondents who don’t know how many layers will be in their data center LANs in two years is notably greater than the number of survey respondents that do know.

There is significant desire on the part of IT organizations to flatten their data center LANs, but there is also significant uncertainty relative to how flat they will become in the next two years.

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 improvement[2] 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. 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. Today, high-speed uplinks are often comprised of multiple 10 GbE links that leverage Link Aggregation (LAG)[3]. 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 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 2012 or 2013.

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

[2] 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.

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

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. Implementing one of the two emerging shortest path first bridging protocols, TRILL and SPB, that support equal cost multi-path bridging can also eliminate loops. 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) technology[4], as shown in Figure 2. 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-

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

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 2, 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 2. The traffic is load balanced via the LAG hashing algorithms rather than being based on VLAN membership, as is the case with more traditional redundant LAN designs. The virtual switch not only improves resource utilization but also enhances availability because the relatively long convergence times of STP topology calculations are circumvented. Virtual switch technology also simplifies management because multiple physical switches can be managed as a single entity.

Figure 2: Switch Virtualization and Multi-Chassis LAG

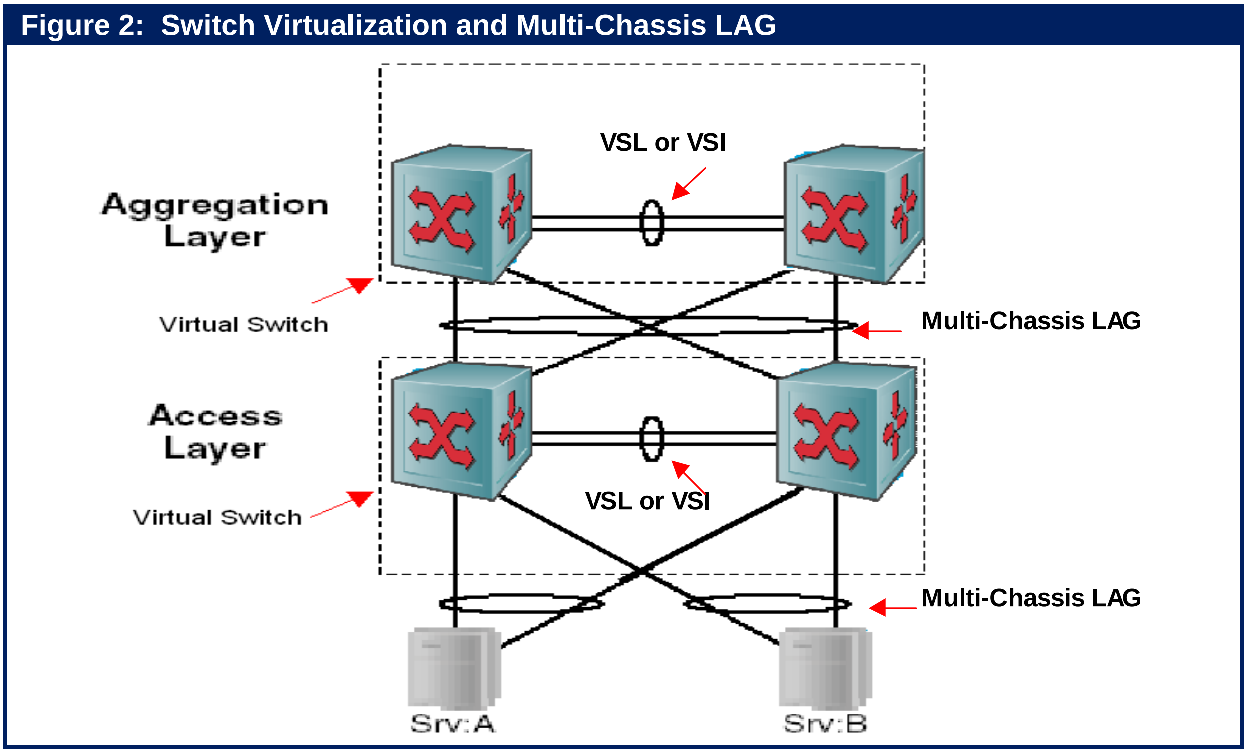

Most vendors of data center switches support switch virtualization and MC LAG in their ToR and modular switches, and these technologies are fully utilized in the two-tier LAN designs that they are currently recommending to enterprise customers. As a result, most two tier LAN designs being proposed by 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. 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 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.

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) project to develop a Layer 2 shortest-path first (SPF) routing protocol for Ethernet. The TRILL RFC (RFC 6325) is currently on the standards track and is being used as the basis for some pre-standard implementations. A similar competing effort is being pursued by the IEEE 802.1aq working group which is defining a standard for shortest path bridging (SPB) of unicast and multicast frames and which supports multiple active topologies. The SPB standard is expected to be ratified by the IEEE by early 2012.

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.

- With TRILL, network virtualization is limited to 4K VLANs, while SPB supports a 16 million service instances via Q-in-Q.

- 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.

SPF bridging should have major implications for data center LAN designs and most of the larger switch vendors are well along in developing switches that can support either TRILL or SPB and network designs based on these technologies. A number of vendors are already shipping pre-standard versions of these protocols, in some cases with proprietary enhancements. 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 2, to a highly meshed or even fully meshed array of switches that appears to the attached devices as a single switch. SPF bridging can support a variety of other topologies, including the fat tree switch topologies[5] that are popular in cluster computing approaches to HPC. Fat trees are also used by Ethernet switch vendors to build high density, non-blocking 10 GbE switches using merchant silicon switch chips. This trend may eventually lead to the commoditization of the data plane aspect of Ethernet switch design. Figure 3 shows how a 48 port 10 GbE TOR switch can be constructed using six 24-port 10 GbE switch chips. By increasing the number of leaf and spine switches, larger switches can be constructed[6]. A number of high density 10 GbE switches currently on the market use this design approach.

[5] http://www.mellanox.com/pdf/whitepapers/IB_vs_Ethernet_Clustering_WP_100.pdf

[6] 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 3: TOR Switch Fat Tree Internal Architecture

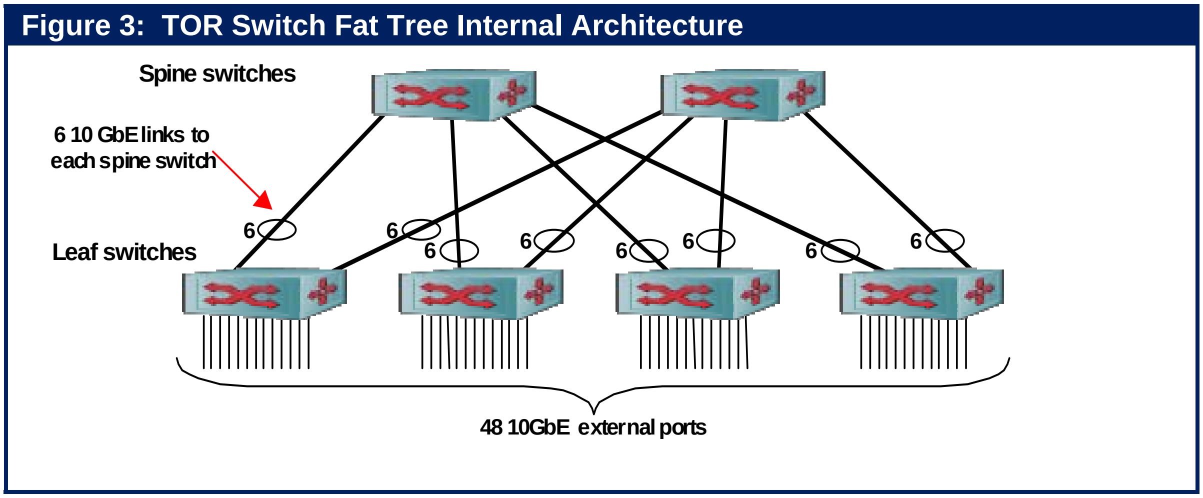

The Interop Respondents were asked, “Two years from now, which of the following is likely to be the most commonly used L2 Ethernet protocol in your company’s data center LANs?” The answers of one hundred and eighty one respondents are summarized in Table 7.

Table 7: Most Common L2 Ethernet Protocol

L2 Ethernet Protocol	Percentage of Respondents
Spanning Tree Protocol (STP)	43%
A Vendor’s Proprietary Protocol	18%
Multi-Switch Link Aggregation (M-LAG)	16%
Transparent Interconnect of Lots of Links (TRILL)	12%
Shortest Path Bridging (SPB)	10%
Other	1%

The data in Table 7 indicates that just under a half of IT organizations expect to still be running STP in their data center LANs in two years. For those IT professionals who indicated that STP would not be the most commonly used L2 Ethernet protocol in two years, there was no consensus as to what protocol would be the most common. However, similar to the situation with flattening the data center LAN, what is even more interesting is that two hundred and seventy four members of the pool of survey respondents answered the question with “don’t know”. That means that the number of survey respondents that don’t know which L2 Ethernet protocol will be the most commonly used in their data center LANs in two years is notably greater than the number that do know.

There is significant desire on the part of IT organizations to move away from using STP in their data center LANs, but there isn’t a consensus as to what the most common replacement technology will be.

A discussion of the alternatives to STP amongst six of the primary data center LAN switch vendors can be found at Webtorials[7].

[7] http://www.webtorials.com/content/tls.html

Controlling and Managing Inter-VM Traffic

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

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

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

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

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

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

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

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

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

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

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

Software Defined Networks and Network Virtualization

With DVS, the switch control plane is decoupled from the data plane and placed in a separate server or controller. This concept can also be applied to the entire data center or campus LAN by removing the control plane from every physical and virtual switch and centralizing it in a control plane server. This centralization would make it relatively easy to programmatically control the entire network. Programmatic control is a key aspect of the concept of a Software Defined Network (SDN) that uses an abstraction layer or *network hypervsior* between the network operating system (NOS) control software and the packet forwarding data plane hardware.

OpenFlow[8] is an open API/protocol that is used between a network controller and a controlled physical switch that provides the forwarding hardware. The protocol is used to set flow table entries within the physical switch. The abstraction layer allows OpenFlow-enabled switches from different vendors to be mixed and matched without impacting the NOS. The Open Networking Foundation (ONF)[9] established in 2012 is now responsible for maintaining the OpenFlow specification, which is currently at Version 1.1.

Building an SDN with OpenFlow requires two components:

- A NOS supporting OpenFlow that is also capable of presenting a logical map of the entire network to the network administrators. This NOS could a modification of an existing proprietary NOS or possibly an open source NOS. The NOS should also be extensible by providing a northbound API to allow new functions to be added.

- Packet forwarding hardware that also supports OpenFlow. In principle, the SDN could be based on a physical network built with OpenFlow switches from a number of different vendors.

Some the potential benefits of an SDN with OpenFlow include:

- Network virtualization where multiple independent virtual networks can share a common physical infrastructure. Virtual networks are based on segmenting flows. Within OpenFlow, flows are defined using a ten-tuple of header fields including Ethernet SA/DA, IP SA/DA, TCP/UDP ports, and VLAN ID. This also provides enhanced security via firewall-style granular control of traffic flows within virtual networks. Network virtualization beyond VLANs is of particular interest in public cloud data centers that provide services to multiple tenants. This concept is elaborated upon in the subsection below that discusses the network support that is required to support the dynamic creation and migration of VMs.

- Network operations are streamlined via the global nature of the network-wide NOS, which results in lower OPEX. In the public cloud, OpenFlow allows the network to be programmatically controlled in conjunction with server and storage resources in order to provision and modify services to tenants. An OpenFlow-enabled NOS with an open API to server virtualization and cloud management systems can potentially be exploited to achieve higher levels of management integration across the data center or the cloud.

[8] http://www.openflow.org/
[9] http://www.opennetworking.org

- SDNs are well suited for highly meshed data center switching fabrics based on the fat tree topologies common in HPC and in web data centers that are dealing with the challenges that are associated with Big Data. Because the control plane has a global view of the network topology, loops can be avoided without resorting to bridging protocols such as STP, TRILL, or SPB.

- Where both the NOS and the packet forwarding hardware support open APIs, network designers would be free to independently optimize each level of the network; e.g., NOS, switches, and applications that extend the functionality of the network. OpenFlow proponents believe this would make the networking industry more innovative and competitive, lowering the overall CAPEX and OPEX cost of network infrastructure.

At Interop 2011, twelve vendors demonstrated prototype switches supporting OpenFlow. Three of these vendors are already shipping switches that support OpenFlow. There are also a number of vendors working on open source OpenFlow-enabled NOS packages and applications that extend NOS functionality. The Open Networking Summit in October 2011 had several demonstrations of OpenFlow implementations on physical hardware combined with network controllers from various vendors. A discussion of OpenFlow amongst six of the primary data center LAN switch vendors can be found at Webtorials[10].

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.

[10] http://www.webtorials.com/content/tls.html

- **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 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 4.

Figure 4: FCoE Converged LAN

Source: Cisco Systems

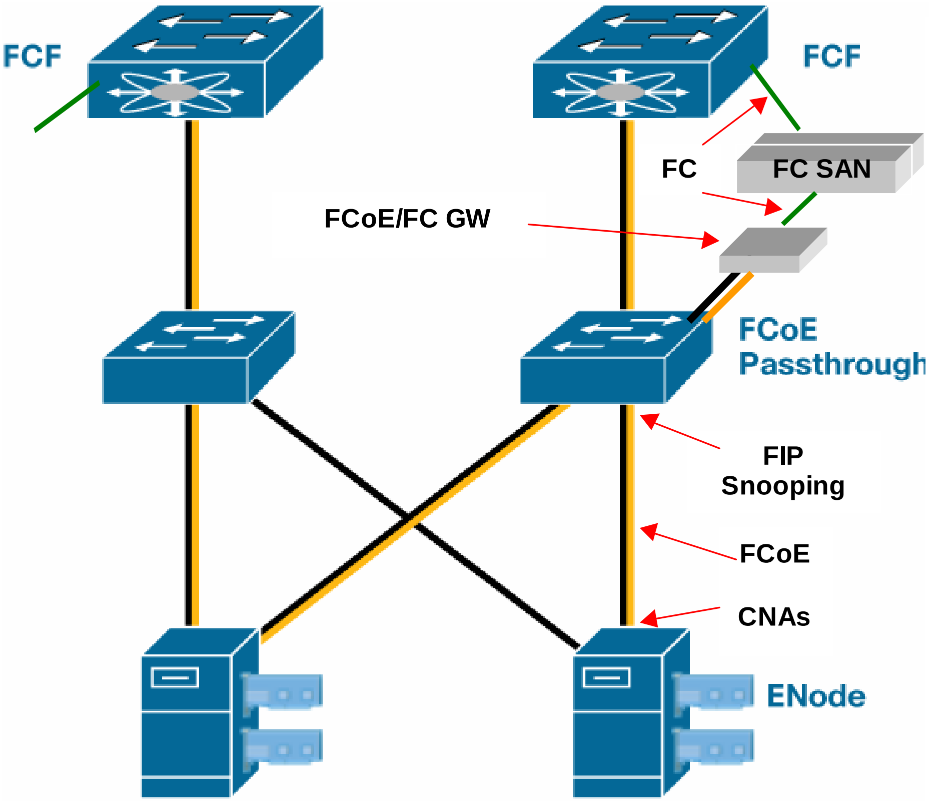

As shown in Figure 4, End Nodes (servers) don't need to connect directly to a FCF capable switch. Instead the FCoE traffic can pass through one or more intermediate FCoE passthrough switches. The minimal requirements for a simple FCoE passthrough switch is support for lossless Ethernet or DCB. The FCoE 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 5.

Figure 5: FCF Support Options

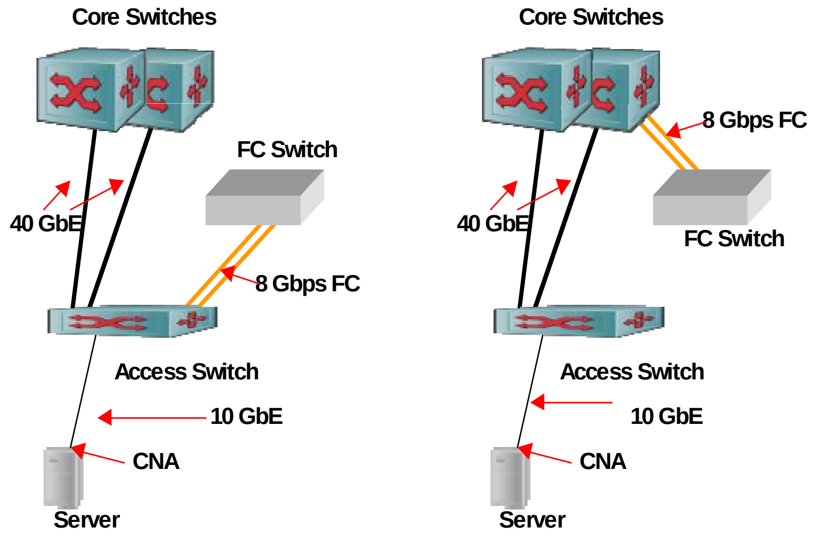

The left side of the figure shows single hop FCoE with the FCF function integrated into the access switch. It would also be possible to use intervening FCoE/FCF gateways, either standalone or incorporated in the FC switch, which would be connected to the access switch via 10 GbE, making the access switch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

The Interop Respondents were asked about their company's current approach to converging the LAN and SAN in their data centers as well as what they thought their company's approach will be two years from now. Fifty nine percent of the Interop Respondents indicated that their company has not currently made any implementation of a converged LAN and SAN. Almost half of the Interop Respondents didn't indicate what they thought their company's approach would be two years from now. Of The Interop Respondents who did indicate what they thought their company's approach would be two years from now, thirty percent indicated that over the next two years that they would either make a significant effort to converge their data center LANs and SANs or they would converge all of their data center LANs and SANs. A discussion of converging the data center LAN and SAN amongst six of the primary data center LAN switch vendors can be found at Webtorials[11].

The ambiguity expressed by The Interop Respondents about their company's direction relative to converging their LANs and SAN, combined with their previously discussed ambiguity about how they will replace the spanning tree protocol and how many layers of switches they will have in their data center LANs leads to the observation that:

> ***The majority of IT organizations have not developed concrete, broad-based plans for the evolution of their data center LANs.***

Network Support for the 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 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 LAN[12].

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.

MPLS/VPLS offers one approach to bridging remote data center LANs together. Another alternative is to tunnel Layer 2 traffic through a public or private IP network using Generic Router Encapsulation (GRE). A more general approach that addresses some of the major limitations of live migration of VMs across a data center network is the Virtual eXtensible LAN

[11] http://www.webtorials.com/content/tls.html
[12] www.vce.com/pdf/solutions/vce-application-mobility-whitepaper.pdf

(VXLAN)[13]. VXLAN is the subject of a recently submitted IETF draft supported by VMware, Cisco, Arista Networks, Broadcom, Red Hat and Citrix. In addition to allowing VMs to migrate transparently across Layer 3 boundaries, VXLAN provides support for virtual networking at Layer 3, circumventing the 802.1Q limitation of 4,094 VLANs, which is proving to be inadequate for VM-intensive enterprise data centers and for multi-tenant cloud data centers. VXLAN also addresses the requirement for multi-tenancy where multiple tenants within a cloud data center could have overlapping MAC and IP addresses.

VXLAN creates a Layer 2 overlay on a Layer 3 network via encapsulation. The VXLAN segment is a Layer 3 construct that replaces the VLAN as the mechanism that segments the network for VMs. Therefore, a VM can only communicate or migrate within a VXLAN segment. The VXLAN segment has a 24 bit VXLAN Network identifier, which supports up to 16 million VXLAN segments within an administrative domain. 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), which is typically present in a hypervisor or a physical 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 draft was submitted to the IETF in August 2011, so ratification of a standard is not imminent. However, VMware and Cisco are likely to include pre-standard implementations in their hypervisor switches in the relatively near future. The IETF draft also discusses VXLAN gateways that connect VXLAN environments to the current VLAN based environments. These gateways are likely to be implemented in hardware switches within the data center.

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 with VLAN boundaries has provided a major impetus for flattening the LAN with two-tier designs featuring Layer 2 connectivity end-to-end.

Many of the benefits of cloud computing depend on the ability to dynamically provision VMs and to migrate them at will among physical servers located in the same data center or in geographically separated data centers. The task of creating or moving a VM is a relatively simple function of the virtual server'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

[13] http://searchservervirtualization.techtarget.com/news/2240074318/VMware-Cisco-propose-VXLAN-for-VM-mobility

implemented logging and auditing functionality. In addition to the VM's network configuration state, this logging and auditing capability also has to be transferred to the new physical server.

The most common approach to automating the manual processes involved in VM provisioning and migration is based on communication between the Hypervisor Management system and the switch element management system (EMS) via APIs supported by both vendors[14].

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.

An interesting benefit of the VXLAN overlay over Layer 3 networks is that IT organizations no longer need to plumb the VLAN for a VM on the link connecting to the ToR switch when the VM migrates. This requirement has been addressed to date through static configuration at the destination hypervisor/ToR switch. As indicated above, dynamic configuration is available via protocols being defined for Edge Virtual Bridging within the IEEE 802.1Qbg working group. With VXLAN, all VM traffic from the hypervisor to the ToR switch is encapsulated within an IP packet so there is no need to plumb the VM's VLAN information on the link between the hypervisor and the ToR switch.

Most data center switch vendors have already implemented some 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.

A somewhat different approach to automating data center configuration, including the provisioning and migration of VMs is based on orchestration engines, which are discussed in more detail in the management section of this report. Service orchestration is a centralized server function that can automate many of the manual tasks involved in provisioning and controlling the capacity of dynamic virtualized services across myriad technology domains; e.g., networking, servers, storage and security. In the case of VM provisioning and migration, the orchestration engine would function as the point of integration between the network device EMS and the hypervisor's management system. This capability requires that third generation data center LANs provide APIs that enable integration with third party orchestrations solutions. Orchestration solutions are available from a number of network management vendors and hypervisor vendors.

[14] While this approach is the most common, some vendors have alternative approaches.

Summary of Third Generation Data Center LAN Technologies

The data center LAN is on the cusp of a number of quite dramatic technology developments, as summarized in Table 8. As shown in the table, most the items on this list are still in flux and require additional development, and/or additional work from the standards bodies[15].

Table 8: 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
Reduced role for blade switches to eliminate switch tier proliferation.	On-going
Changing role for the hypervisor vSwitch as a port aggregator (VEPA) for EVB, essentially eliminating the vSwitch tier.	A standard is in progress and pre-standard implementations are available.
STP issues are being addressed by switch virtualization and multi-chassis LAG technology, as well as by newer protocols such as TRILL/SPB.	On-going deployment
Multi-core servers with notably more VMs per server and 10 GbE connectivity to the LAN.	Early adoption stage.
40 GbE and 100 GbE uplinks and core switches.	A standard is in place: 40 GbE is available 100 GbE due in 2012
DCB delivering lossless Ethernet for 10 GbE and higher speed Ethernet	Standards are in place. Implementations are being announced.
SDN and OpenFlow	Specifications have bees released There is some prototype switch support and some NOS support from startups.
VXLAN extended virtual networks address VLAN scalability, multi-tenancy and switch hardware Layer 2 table capacity issues that are caused by the proliferation of virtualization.	A draft was recently submitted to the IETF. Pre-standard implementations are expected.
FCoE approach to fabric unification	FCoE standard is in place. Early implementations are based on pre-standard DCB.
10 GbE iSCSI approach to fabric unification	Early implementations were over pre-standard DCB.
TRILL/SPB enabling new data center LAN topologies; e.g., fully meshed, fat tree with equal cost multi-path forwarding	Standards are in progress. Pre-standard implementations of both SPB and TRILL are available. SPB is expected to be finalized by early 2012
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.

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

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.

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.

A10 Networks

The Fastest Growing Application Networking Company

64-bit AX Series

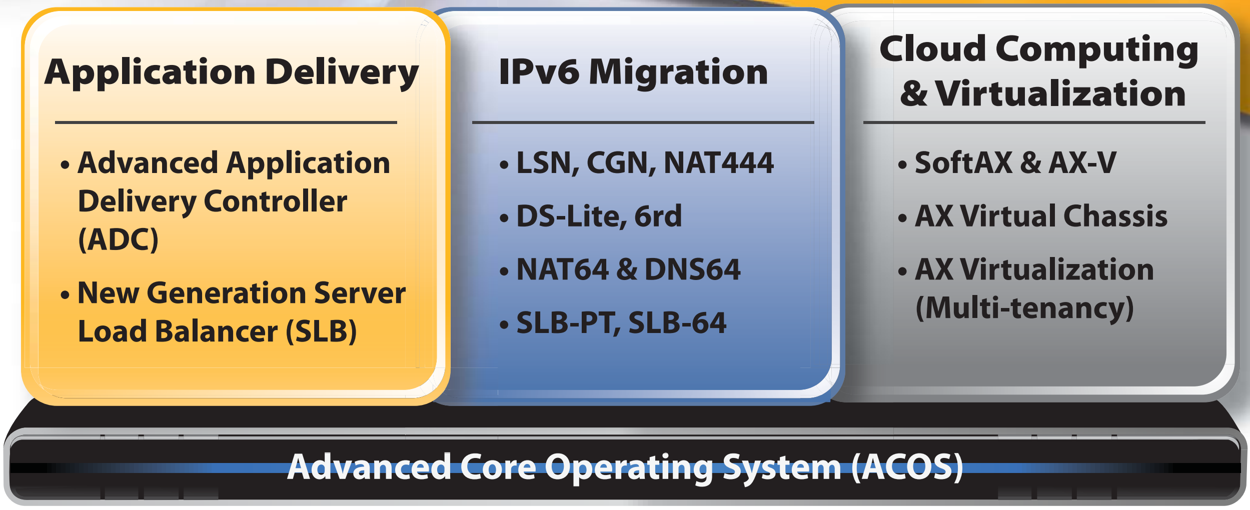

Application Delivery

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

IPv6 Migration

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

Cloud Computing & Virtualization

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

Advanced Core Operating System (ACOS)

AX Series Advantage

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

Application Solutions

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

Bb Blackboard   Microsoft   Microsoft Exchange   Microsoft Lync

IMPERVA   Infoblox   JUNIPER NETWORKS   ORACLE   vmware READY

AVAYA

The Power of We™

Cloud Networking – the not-so-quiet revolution

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

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

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

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

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

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

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

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

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

Cisco Unified Fabric

Converged. Scalable. Intelligent.

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

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

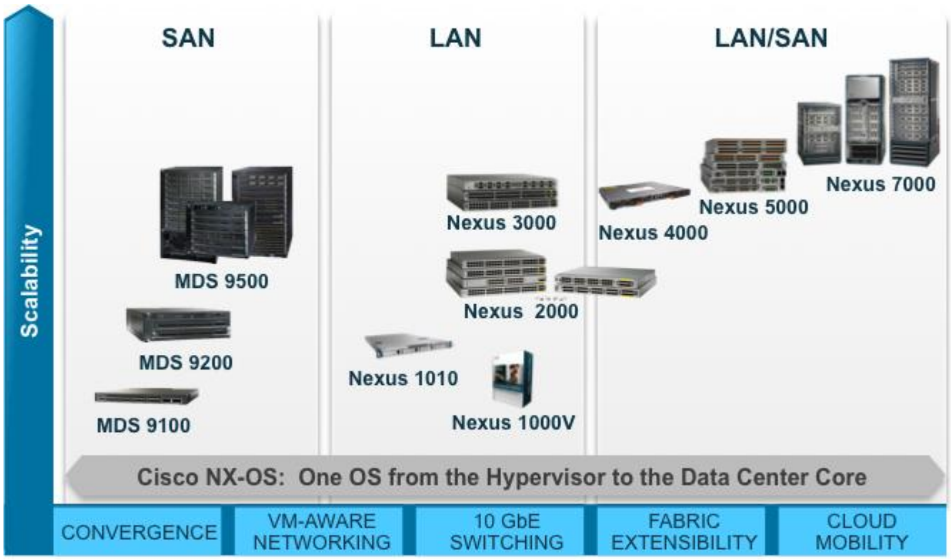

- Revolutionary fabric scale with over twelve thousand 10 GbE server connectivity with Cisco Nexus
- Highest 10Gb Ethernet density in the industry with Cisco Nexus 7000
- High performance and ultra-low latency networking at scale with Cisco Nexus
- Network services delivered in virtual and physical form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

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

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

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

Consolidation and Cloud Computing Without Compromise

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

Starting Point

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

Next Step

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

No Compromises

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

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

Unique NetScaler Strengths

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

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

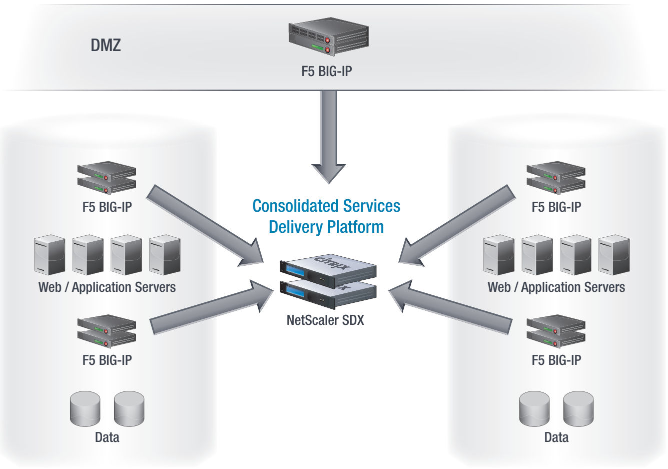

Network Services Consolidation

OPEN

HP FLEXNETWORK ARCHITECTURE

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

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

HP FlexNetwork—an architectural blueprint for cloud-optimized networking

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

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

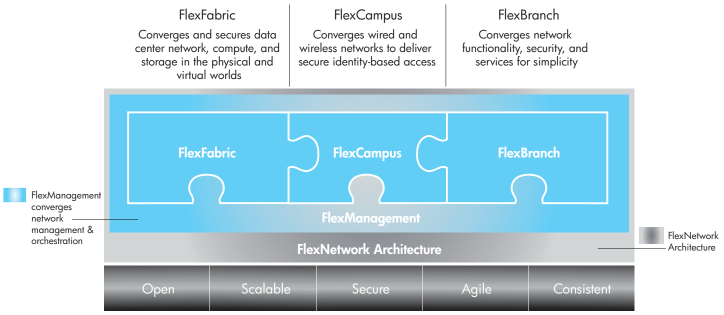
FlexFabric
Converges and secures data center network, compute, and storage in the physical and virtual worlds
FlexCampus
Converges wired and wireless networks to deliver secure identity-based access
FlexBranch
Converges network functionality, security, and services for simplicity
FlexFabric
FlexCampus
FlexBranch
FlexManagement
FlexManagement converges network management & orchestration
FlexNetwork Architecture
FlexNetwork Architecture
Open
Scalable
Secure
Agile
Consistent

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

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

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

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

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

HP CloudSystem and HP FlexNetwork networking solutions deliver:

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

For more information

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

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

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

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

Created November 2011

EXECUTIVE VIEWPOINT

Maximizing Your IT Resources

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

Steve Nye

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

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

Infoblox®

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

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

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

WHAT IS THE IMPACT OF VIRTUALIZATION ON NETWORK STAFF?

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

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

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

HOW DOES THE MOVEMENT TO IPV6 AFFECT NETWORK STAFF?

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

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

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

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

nano|engine

Full application control at 10% of the cost

www.ipanematech.com

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

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

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

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

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

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

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

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

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

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

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

Packet Design Solutions:

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

Route Explorer: Industry-Leading Route Analytics Solution

Optimize IP Networks with Route Explorer

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

Deployed in the world's largest IP networks

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

Overview of Route Explorer

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

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

Optimize IP Networks with Traffic Explorer

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

Traffic Explorer Architecture:

Traffic Explorer consists of three components:

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

Traffic Explorer Applications

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

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

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

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

Overview of Traffic Explorer

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

For more information, contact Packet Design at:

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

IT Organizations Find Key Enabling Technologies for Adopting Cloud Architectures

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

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

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

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

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

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

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

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

riverbed.com

riverbed®

www.vx-xpress.com

VISUAL NETWORK SYSTEMS™

Application Performance Management in the Cloud

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

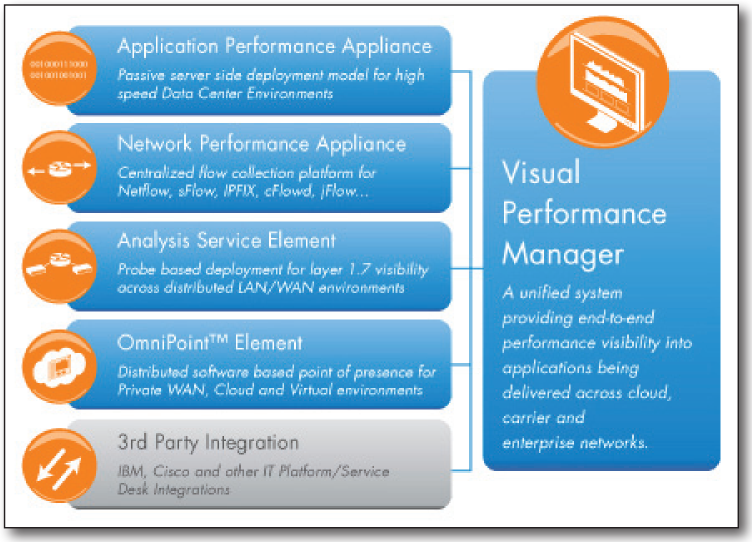

Application Performance Appliance
Passive server side deployment model for high speed Data Center Environments

Network Performance Appliance
Centralized flow collection platform for Netflow, sFlow, IPFIX, cFlowd, jFlow...

Analysis Service Element
Probe based deployment for layer 1-7 visibility across distributed LAN/WAN environments

OmniPoint™ Element
Distributed software based point of presence for Private WAN, Cloud and Virtual environments

3rd Party Integration
IBM, Cisco and other IT Platform/Service Desk Integrations

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

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

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

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

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

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

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

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

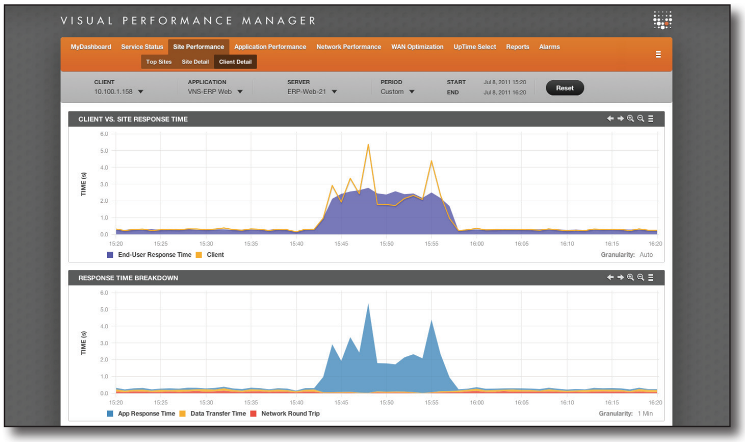

VISUAL PERFORMANCE MANAGER

CLIENT 10.100.1.158 ▼ | APPLICATION VNS-ERP Web ▼ | SERVER ERP-Web-21 ▼ | PERIOD Custom ▼ | START Jul 8, 2011 15:20 | END Jul 8, 2011 16:20

CLIENT VS. SITE RESPONSE TIME

Granularity: Auto

RESPONSE TIME BREAKDOWN

Granularity: 1 Min

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

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

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