Issue 3

Enabling the Enterprise Network for Cloud Computing

Reducing Cost and Complexity, Enabling the Agile Enterprise

Featuring research from Gartner

Introduction

Avaya leads the industry with **'Fit for Purpose' data networking** that will redefine the IT value proposition and offer the best return on investment for business communications. Furthering our industry leadership, we recently launched a **revolutionary end-to-end virtualization strategy** and architecture that will simplify data center and campus networking.

Avaya Virtual Enterprise Network Architecture, or Avaya VENA, allows organizations to more easily optimize business applications and service deployments in and between data centers and campuses. Avaya VENA reduces costs and significantly improves time to service by more efficiently optimizing the network connections between application servers and end users – essentially providing an end-to-end connection from the desktop all the way to the data center. It can also reduce risks for CIOs by simplifying provisioning and policy configurations, reducing network re-designs, and providing new, streamlined tools for network management.

Avaya also announced a series of relationships in support of Avaya VENA with companies including VMware, an industry leader in virtualization and cloud infrastructure; QLogic, an industry leader in converged networking; Coraid, an industry innovator in converged Ethernet SAN storage; and Silver Peak Systems, the leader in data center WAN optimization - in addition to many other companies that have joined as members of Avaya's DevConnect program, an initiative to develop, market and sell innovative third-party products that interoperate with Avaya technology. Avaya VENA is based on open industry standards, and Avaya is aggressively pursuing and building out a broader ecosystem of technology collaborators to support the architecture.

No one else can offer this level of simplicity around virtualization and collaboration. Together with our technology partners, we are delivering a truly revolutionary and unique solution that maximizes the business model for virtualization and data center networking. The new Avaya VENA architecture, coupled with our cutting-edge innovation in business collaboration, will redefine the way enterprises communicate.

You will be hearing a lot about Avaya VENA in the coming days and months. But on one thing you can rest assured… With Avaya's data solutions, your network can adapt to your business needs quickly and more efficiently by delivering immediate and long term value.

Today's CIOs are focused on saving money, driving new revenue streams, increasing the loyalty of their customers and creating competitive differentiation.

Our customers are telling us they want unique solutions that are tailored and best-designed to support their specific communications requirements.

And our strategy is focused on addressing these needs with some of the most competitive products in the industry.

So when you are scrutinizing your next data networking investment, be sure to look at Avaya Data Solutions.

Steve Bandrowczak, Vice President and General Manager, Avaya Data Solutions

Fit-for-Purpose Data Center Networking

Simplified and optimized service orchestration maximizes the return from a virtualized computing environment

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 – scalable, however many existing deployment methodologies are no longer an option.

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. Many IT organizations are burdened with the additional challenge of supporting a mosaic of servers, operating systems, and applications which have been collected during numerous purchasing cycles. Tales abound of 'specials', too awkward to port but too important to lose. A closed, single-vendor approach is simply not an option for many; versatility and the ability to seamlessly transition the underlying network are key in delivering transformation without introducing risk.

The progressive evolution of the network also creates a new opportunity: the convergence of storage and real-time traffic. While researchers caution against immediate adoption, they do agree that the writing is probably on the wall for stand-alone storage; the network of the future will be required to provide low-latency, shortest path connectivity, and purpose-built products will deliver the end-to-end lossless architecture mandatory for the successful convergence of storage.

Section 1: Creating a plug and play Data Center

Simply put, the goal is to make the Core of the network transparent. By evolving the Data Centre and Campus Core infrastructures into one that seamlessly integrates services – and by this we typically mean applications delivered by virtualised computing systems – availability, performance, and efficiency can be dramatically improved. Simplified service orchestration is possible when the network facilitates an easy and effective, yet highly granular, mapping capability.

The limitations of legacy networks are many. They include unreliability and unpredictability, difficulty in achieving effective scale, unavoidable complexity, inefficient resource utilization, delays in provisioning new services and slow time-to-service, and difficulty in maintaining separation of traffic. These limitations manifest themselves in an inability to decouple networking considerations from service delivery which introduces a genuine element of constraint.

AVAYA VIRTUAL ENTERPRISE NETWORK ARCHITECTURE

To simplify communication and understanding a solution name is used: the 'Avaya Virtual Enterprise Network Architecture'. Components of this architecture include:

- Products and Tools – existing and new hardware and software elements that combine to deliver an integrated solution
- The 'Virtual Services Fabric' – the foundational network fabric delivered by the combined functionality of standards-based Shortest Path Bridging and Avaya's value-adding features, overlaid with various management tools
- 'Virtual Service Networks' – implementations of virtual networks spanning and being transported by the Virtual Service Fabric

Potential customers will be able to clearly appreciate how individual portfolio products and components fit into the wider strategy. Initially solutions empowered by Avaya's Virtual Enterprise Network Architecture will be available using the Ethernet Routing Switch 8800, and the upcoming Virtual Services Platform 9000. The architecture will also see this capability extended to future Top-of-Rack products.

Network transparency is achieved by taking a new approach to design, including doing away with unnecessary complexity, and leveraging new capabilities.

The technology of choice for creating the foundations of network transparency will be that of Shortest Path Bridging, or SPB, as defined by the Institute of Electrical and Electronics Engineers (IEEE) in the 802.1aq standard. SPB is based on Carrier-grade technology developed by Avaya engineers and field-proven in Service Provider networks today. SPB delivers new capabilities in the crucial areas of simplicity, scalability, performance, reliability, and service orchestration and abstraction.

Built around a fault-tolerant, powerful, and self-aware Core, the transparent network features a design where service provisioning occurs only at the Edge. The advantage is immediate and pronounced; administrative effort is reduced, errors are avoided, and time-to-service is vastly enhanced.

The beauty of SPB is that it is masking devices, links, and protocols and delivering what is logically an extended Ethernet LAN that provides connectivity for multiple end-points. That's the simple concept and SPB achieves this is an interesting and quite unique way. It leverages an established dynamic link-state routing protocol called Intermediate System-to-Intermediate System (IS-IS) and extends it to share topology, reachability, and device information between every node in the SPB domain. With nodes holding their own self-determined view of the network, including the optimal path to any destination, a fully distributed and dynamically maintained solution is created.

Around the perimeter of the domain each SPB node builds and distributes a database of all attached services; the servers and networks that each directly supports. Encapsulating the Customer MAC Addresses (C-MACs) within a header based on the 802.1ah format, SBP effectively hides the Edge from the Core; nodes within the SBP cloud are logically decoupled from customer networks, ensuring scalability and stability.

Importantly this solution also delivers simplicity, particularly in the area of provisioning. Most will be familiar with the concept of Virtual LANs (VLANs) and the fact that each VLAN is normally associated with a unique IP Address range; servers are assigned to an appropriate VLAN and issued with an address. Application and server virtualization complicate the network by making it very difficult to support dynamic virtualization and potentially messy when coupled with managing the distribution of VLANs across extended Data Centers. SPB elegantly solves this dilemma by mapping each VLAN instance to a service identifier (known as an I-SID) which advertises these services to the network; those SPB nodes that require connectivity automatically calculate their shortest path or paths and populate their forwarding database. There is no requirement to configure the Core to pass service VLANs, to administer path selection or redundancy, and no risk of configuration errors.

An added advantage of the SPB model is that many stability problems associated with traditional networks can be totally avoided. Due to factors such as topology being managed by IS-IS (and not by Spanning Tree) and MAC Address from within the customer networks being hidden from the Core, loops or broadcast storms in the Edge cannot impact the Core or any other customer environment. Loop prevention is a fundamental aspect of traditional Ethernet networks however SPB introduces loop mitigation to improve forwarding resolution during topology changes.

Performance will be fully optimized because SPB is fundamentally all about 'shortest path' connectivity. In creating symmetrical shortest paths between common services SPB delivers the best possible utilization of all available resources, indeed multiple equal cost routes are supported so that load-sharing is automatic. The limitations within old-fashioned Spanning Tree-based networks of blocked and sub-optimal paths based on a remote root disappear. Due to the symmetrical nature of SPB paths, SPB provides predictable path selection which aids in operational diagnosis and management.

It should be noted that because SPB implements the exchange of a single digest of link state information covering the entire network view, negotiation of each individual path-to-root is not required. This means that the volume of messaging exchanged in order for the network to converge is proportional to the change in topology instead of the number of multicast trees in the network. A simple link event that may change many trees is communicated by signalling the link event only; similarly, the removal of a bridge, which might involve the rebuilding of hundreds of trees, is signalled only with a few link state updates.

Finally, SPB can claim to be truly Data Center fit-for-purpose; being totally compatible with various server/application virtualization technologies such as VMware and Hyper-V. It fully supports Microsoft's L2 network load-balancing solution (NLBS), and it is transparent to established L3 redundancy techniques (VRRP). In short SPB extends the capabilities of L2 within the Data Center, increasing the efficiency of what is widely recognized as the most efficient transport layer, and certainly making it more scalable.

Not unusual for our industry, there is a competing standardization effort. This is being led by the Internet Engineering Task Force (IETF) and is known as 'TRILL'; Transparent Interconnect of Lots of Links. Ironically TRILL also leverages IS-IS for topology; however, it doesn't offer SPB's service abstraction of customer networks from the Core. This necessitates a complex solution and ultimately limits scalability. In short, whereas SPB is based on a proven Carrier-grade L2 technology, TRILL introduces a totally new, and as yet unproven, version of IS-IS and packet encapsulation that attempts to make an L3 network look and feel like one operating at L2.

Section 2: Avaya's fit-for-purpose solution for the data network

Leveraging SPB as the foundational technology, Avaya augments this with extensions to add further usability and functionality – SPB is after all a Layer 2 technology (as is TRILL) – and, unimproved, SPB probably lacks the flexibility to be a real-world option. Avaya's strategy is to integrate our field-proven techniques with new capabilities that facilitate a truly fit-for-purpose solution for the data-centric enterprise.

Building on the numerous advantages that SPB provides for L2 connectivity, Avaya leverages the versatility of IS-IS to also extend and distribute L3 Virtual Routing and Forwarding (VRF) instances. No additional protocol support is required to implement this functionality, simply map the VRF instances to the appropriate service ID. The network now operates as a fully functioning L2 and L3 domain.

Another example of fit-for-purpose networking is the ability to implement selective routing between virtual networks. Inter-network routing provides for effective and highly granular control of traffic flows between separate networks. Typical examples might be to facilitate north-south traffic between user and business applications, and the placement of firewalls or other Application Delivery Controllers.

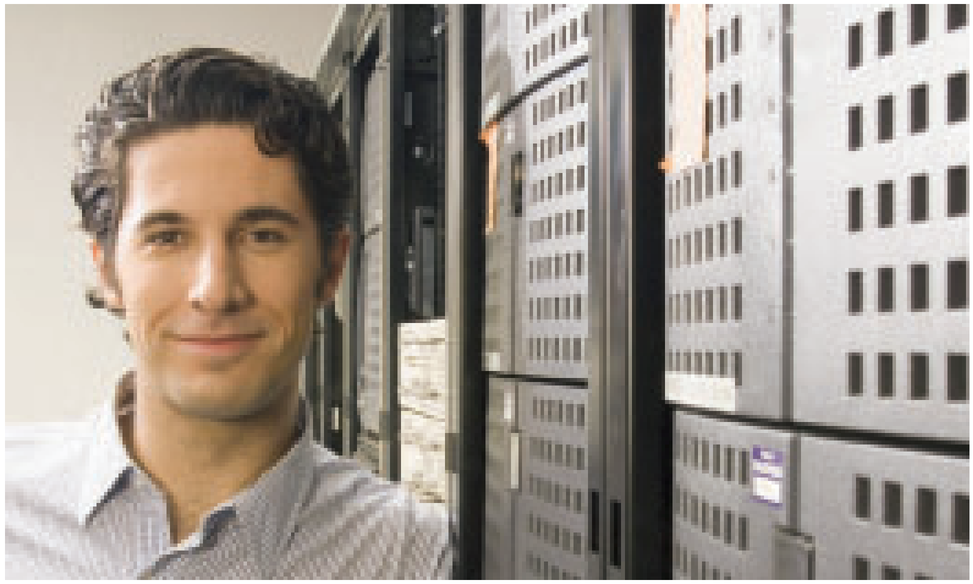

Avaya, uniquely positioned based on decades of networking experience, helps ensure that the transition to an SPB network is both low-risk and seamless. Avaya's ground-breaking Switch Cluster technology is deployed on the perimeter of the SPB domain providing resilient active-active connections for attached devices: servers or access switches. These 'domain gateway' nodes map services into SPB, in addition to performing the traditional functions of VLAN separation and IP Default Gateway; Avaya's Routed SMLT and VRRP Back-Up Master options again deliver crucial resiliency and efficiency differentiation.

Top-of-Rack switches are provided with high-availability and high-performance connectivity into the SPB domain; and the Horizontal Stacking feature that Avaya delivers is leveraged to optimize the availability and quality-of-experience for hosted applications.

In addition to the standards-based SPB architecture Avaya offers a number of value-add features:

- Avaya's Switch Cluster pairs delivering always-on fault-tolerant access to the SPB Core; a model that scales seamlessly
- Avaya's Horizontal Stacking within the Data Center provides always-on high- performance connectivity within and between racks
- VRF-Lite and Avaya's IP VPN-Lite offer multiple virtualization options and enable extremely versatile solution creation
- Avaya's Stackable Chassis delivers always-on high-performance user connectivity in the wiring closet

Offered together, these products and technologies form the Avaya Virtual Enterprise Network Architecture, a truly fit-for-purpose solution for the Data Center and Enterprise Campus; delivering simplified yet optimized end-to-end connectivity between users and their content.

Section 3: Learn More

To learn more about fit-for-purpose Data Center networking or the Avaya Virtual Enterprise Network Architecture, contact your Avaya Account Manager, Avaya Channel Partner, or visit us at www.avaya.com.

STORAGE NETWORKING AND THE NETWORK

The desire to converge storage onto the network is driven by the potential for significant cost savings: reducing adaptors, cabling, and networking components such as Top-of-Rack switches.

iSCSI (Internet Small Computer System Interface) currently delivers an IP-based solution but is perceived by some as a low-end network-attached storage option because of its "PC Server" heritage. Fibre Channel comes with supercomputer credibility, however it also comes with a complex protocol stack and assumes a dedicated, lossless environment.

Fibre Channel over Ethernet (FCoE) is intended to leverage the pervasiveness and scalability of Ethernet to deliver a cost-effective yet high-performance solution. However the need to have a lossless – or virtually lossless – environment remains and more standards work, under the auspices of the IEEE's Data Center Bridging Task Group, is required in order to deliver the necessary framework ahead of interoperable products. Importantly, FCoE needs to find a way to accommodate Fibre Channel in a multi-hop architecture; it doesn't have the inherent "routability" of iSCSI.

This is another opportunity for SPB to prove its worth. By transparently connecting nodes and automatically forwarding traffic on the most direct path over a simplified Core, SPB emulates the kind of network that is native to Fibre Channel.

TRILL, POSSIBLY NOT AS TRANSPARENT AS THE NAME IMPLIES

Touted by advocates as the 'next big thing', displacing VPLS (Virtual Private LAN Service), for providing Ethernet-based many-to-many connectivity, TRILL leverages IS-IS as a topology management protocol but introduces the need for new IS-IS packet formats.

TRILL will build one or more rooted spanning trees to support flooding of packets. Unicast traffic takes the shortest path with forwarding decisions being made on a hop-by-hop basis; unknown and Broadcast traffic travels to the root of the tree. This implies that different types of traffic may take different paths through the network, even though all traffic may be going between the same source and destination; creating the potential for out-of-sequence packets.

A new header format has also been created for TRILL. This new header sits behind the standard MAC header and establishes communications between TRILL nodes. There is also a new TTL (Time-to-Live) field that is needed to minimize the impact of loops within the TRILL network. This TTL is needed mainly because the formation of the non-congruent trees for Unicast, Multicast, and Broadcast traffic.

TRILL is able to support up to 4,000 VLANs and introduces a new protocol to advertise End Station Address Information (ESADI). TRILL supports Customer MAC Addresses and 802.1Q VLAN IDs. There is no abstraction and no simple way to map VLANs into different services; without this abstraction, TRILL lacks the ability to deliver granular control of traffic.

Source: Avaya

Kutztown University Case Study

CASE STUDY

Challenge

With its data center reaching capacity, Kutztown University needed to strategize for future growth to support state-of-the-art educational and administrative applications, with 24/7 uptime for students and faculty. Their goal was to develop a scalable, cost-efficient, energy and space-saving solution that could provide the required density and a clear pathway to network virtualization.

Solution

One Avaya VSP 9000 has already been installed for the residential network, and the University is strategizing toward an Avaya VSP 9000 installation for the administrative/academic network. (Please see the Systems and Applications chart on page 4 for additional details.)

Value Created

- **A fully scalable architecture** that provides immediate support for very high-density 1 and 10 Gigabit Ethernet, and future-readiness for 40 and 100 Gigabit Ethernet standards.
- **Investment protection:** infra-structure will accommodate future technological advances for many years to come.
- **Cost savings and operational efficiencies** are achieved through a simplified network infrastructure that speeds deployment of new systems and applications and allows for expansion without costly network re-engineering.
- **High reliability, resiliency, and availability**—helping to ensure 24/7 uptime for day and night usage by students and faculty.

Kutztown, Pennsylvania – A quiet transformation is taking place in the cool, well-orchestrated data center on the Kutztown University campus.

What's new on campus?

A next-generation data network places Kutztown University on the leading edge of educational technologies— now, and for many years to come

With the installation of an Avaya Virtual Services Platform 9000 (Avaya VSP 9000), the IT team has launched the University on a journey into next-generation data networking. It's a journey that is expected to reduce overall operating expenses, ensure the highest levels of flexibility and reliability, and enable the latest advances in data, voice, and video technology for years to come.

Building on a legacy of success

Since 2001, Kutztown University has successfully developed its ResNet (residential) and administrative/academic data networks with two Avaya Ethernet Routing Switches 8600 at the core of each network. The infrastructure has successfully accommodated a 25% increase in student body size and the construction of several new buildings on campus.

However, demands on the network continue to grow. With over 10,000 students and approximately 1,000 full- and part-time employees working at all hours of the day and night, research and advanced scientific projects and applications requiring high bandwidth and throughput, and administrative processes constantly advancing, Kutztown University needs to be both proactive in its deployment of network technology and careful in its allocation of resources.

During the latest refresh cycle for the residential network, the University's IT team had an opportunity to carefully consider the current and future network needs of the University and to establish a strong foundation for future growth.

According to Kevin Schukraft, Manager of Network Technology, many factors defined Kutztown University's networking needs going forward. "Our data center was at capacity, and instead of adding more and more hardware, we wanted to have a solution that would give us greater capacity without increasing energy usage.

> "We also wanted a network that could enable us to stay current with technology. For this we would require much greater density and throughput, with a clear pathway to network virtualization."

Senior Network Engineer Josh Heller commented, "There are many applications that already require enormous throughput, and the number of these are certainly going to expand in the future. For example, the science department often transfers massive amounts of data from various atmospheric research institutions, and these are used to teach geological modeling. Applications such as these make IT work on a campus very exciting but at the same time very data-intensive and IT-critical."

"Thirdly, our campus environment requires maximum uptime, because our students, faculty, and staff use the network round the clock, and that usage is stretched by the growth of online education and an increase in the number of non-traditional students who study with us. We knew that Avaya solutions offer high resiliency and reliability, and through virtualizing our two networks, we could take those capacities to the next level, with a live back-up for any type of failure on either side."

Phasing in a powerful solution to accommodate growth

Working closely with Avaya and an Avaya Connect channel partner, the Kutztown IT team examined a scenario that would involve enhancing their existing Avaya ERS 8600 infrastructure and another that would involve replacing the two ERS 8600 switches used for the residential network with the new Avaya VSP 9000.

Today, the University's data center is in production with the Avaya VSP 9000 for ResNet, while retaining the two Avaya ERS 8600s for the administrative network.

"When we started comparing the options, it seemed that from the perspective of both finance and performance, we would gain by moving ahead with the Avaya VSP 9000," Schukraft commented. "We felt that this technology would continue to develop and that it will be the best way to pursue our plans for virtualization."

Why Avaya? Schukraft stated, "Our Avaya equipment has performed very well for many years, and the reliability it gives us is extremely important. The Avaya team has a track record of providing excellent service in both the planning and execution of new installations. Their future roadmaps are very clear and progressive, so we can have confidence in where we are moving for the future. As we consider all of our options, Avaya is by far the best fit for us."

Heller added, "We also have to consider time of deployment, and the importance of having a good support team in place. Typically we operate on very, very short timeframes to do large-scale deployments. The relationships we have with Avaya and our Avaya channel partner make a big difference in a successful deployment." According to Heller, moving into production with the Avaya VSP 9000 was done on an extremely accelerated timeframe in order to be ready for the return of Kutztown University students for the fall semester 2010.

"Instead of being in a user-vendor relationship, I think we function more as a partnership," Schukraft said. "Avaya and our Avaya channel partner have accommodated us very quickly and really taken care of us as a customer. Being a partnership helps. No matter what the issue may be, they take care of it with no questions asked."

Future-proofing the campus networks with an advanced Avaya platform that is always-on, efficient, and scalable

By moving to the Avaya VSP 9000 platform, Kutztown University has taken an important step to future-proofing their networks. The platform delivers a wide range of benefits, including:

- **High-density, fully scalable architecture.** The Avaya VSP 9000 provides immediate support for very high-density 1 and 10 Gigabit Ethernet, and future-readiness for the emerging 40 and 100 Gigabit Ethernet standards. The fully scalable architecture can scale seamlessly in line with performance requirements, without complex or expensive re-engineering.

Schukraft commented, "10-Gig will set us up to accommodate needs in the academic arena, where applications are starting to drive more demand for high-speed performance. Currently, some of the greatest needs for enhanced performance and throughput come from our business college, the computer science program, and our research programs where there is heavy exchange of data. We anticipate that these needs will continue to increase, and that we will also need to accommodate new high-performance applications in the administrative area as well."

Schukraft mentioned that the University is also acutely aware of the students' residential needs. "This is their home for many months of the year, and they expect to get the same service or better than they would get at home. Enabling their work and satisfying their needs is a top priority. We realize that the students today are brought up with the technology, and they're more advanced than a lot of people because they're exposed to the technology at such a young age. Everything they do is based around technology. This new Avaya network will help us keep pace with their technology expectations."

- **Cost savings, investment protection, and green benefits.** A less complex, more agile virtual network infrastructure simplifies the network and helps reduce the cost of deploying new services. Upgrades to meet evolving network and application requirements can be done efficiently, reducing operating expenses and protecting network investments.

The Avaya VSP 9000 equipment runs on either 110 or 220 volts. Kutztown University plans to convert to using 220 volts to provide higher performance without increasing energy usage. The overall footprint is reduced, with about 30% less rack space required to house the equipment. (Note: as indicated below, power consumption and cooling needs can be reduced, with industry estimates showing possible reductions of 50% or more when multi-servers are consolidated and virtualized.)

- **Reliability/Resiliency/Availability.** The Avaya VSP 9000 platform provides density of the high-speed ports, which the IT team has identified as the main driving factor in where they need to go. They plan to have 4 virtual routers on two Avaya VSP 9000s, to achieve the levels of performance and reliability that are critical to their future growth.

The Avaya VSP 9000 was designed and developed on the basis of leading-edge hardware resiliency with numerous mechanisms in place to ensure system availability. It can provide unmatched resiliency for Kutztown University because it is powered by Avaya's unique switch clustering capability, which uses both split multi-link trunking and routed split multi-link trunking technology. The high-availability mode engages all links when forwarding traffic. Instantaneous all-port re-routing results in the elimination of packet loss, which is an important feature when many students and faculty are using next-generation applications. Redundant and hot-swappable control processor and switch fabric modules, plus redundant cooling fans and power supplies, contribute to reliability of the Kutztown University network. As the University continues on its path to virtualization, the IT team can expect to experience further enhanced performance features, to the point where utilization rates can run to 60-70% or more (a typical current utilization rate in the IT industry is roughly 10%). Management benefits can be expected to include application provisioning, streamlined maintenance, high-availability, disaster recovery, and the ability to run multiple operating systems.

- **Enabling future advances in IP telephony.** The university currently utilizes Avaya Communication Server 1000 (Avaya CS1000) with Contact Center Manager to support agents at the various help desks for the IT department, the registrar's office, the financial aid office, the admissions office, the bursar's office, and the health and wellness center. The University is also phasing in Avaya CallPilot™, which provides voice/fax messaging and integrated unified messaging capabilities through the user's familiar desktop e-mail environment, plus Web-based unified messaging and personal mailbox management through My CallPilot.

"We've had very positive results with the Avaya telephony products, and we plan to continue to develop their potential," Schukraft mentioned. "We've seen the roadmap demonstrating how these systems will be supported by the new Avaya VSP 9000 solution. It appears that we will have many opportunities to enhance the services that we offer for staff, students, parents, and even grandparents who rely heavily on the phones for information and services."

Schukraft concluded, "Keeping our network always on and available for every user, with the throughput that is needed, is our first priority. Our data center strategy is now centered around the capabilities we will have with the Avaya VSP 9000. We believe it will be very effective in enabling our vision of supporting the university mission of educating students."

Learn More

For more information on how Avaya Intelligent Communications can take your enterprise from where it is to where it needs to be, contact your Avaya Account Manager or a member of the Avaya Connect channel partner program, or access other collaterals by clicking on Resource Library at www.avaya.com.

All statements in this case study were made by Kevin Schukraft, Manager of Network Technology; and Josh Heller, Senior Network Engineer.

"Keeping our network always on and available for every user, with the throughput that is needed, is our first priority. Our data center strategy is now centered around the capabilities we will have with the Avaya VSP 9000. We believe it will be very effective in enabling us to support the university's mission of educating students."

— Kevin Schukraft, Manager of Network Technology

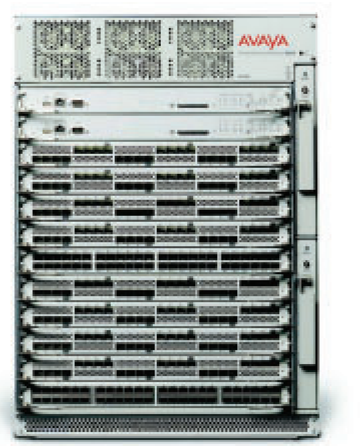

SYSTEMS AND APPLICATIONS

Data Center

- Avaya Virtual Services Platform 9000
- Avaya Ethernet Routing Switch 5650-48T
- Avaya Ethernet Routing Switch 8600
- Avaya VPN 2600
- Avaya VPN 3050
- Avaya Ethernet Routing Switch5510-48T

Telephone Room

- Avaya Communication Server 1000e
- Avaya Contact Center Manager
- Avaya CallPilot™
- Avaya Fiber Remote Shelves

Telecommunication Rooms: ResNet (student residential) Network

- Avaya Ethernet Routing Switch 470-48T
- Avaya Ethernet Routing Switch BPS-2000
- Avaya Ethernet Routing Switch 5510-48T
- Avaya Ethernet Routing Switch 5520-48T
- Avaya Ethernet Routing Switch 2550T-PWR
- Avaya Ethernet Routing Switch 5698TFD-PWR

Administrative/Academic Network

- Avaya Ethernet Routing Switch 470-48T
- Avaya Ethernet Routing Switch BPS-2000
- Avaya Ethernet Routing Switch 5510-48T
- Avaya Ethernet Routing Switch 5520-48T
- Avaya Ethernet Routing Switch 470-24T
- Avaya Ethernet Routing Switch 5698TFD-PWR

Source: Avaya

From the Gartner Files:

Myth: A Single FCoE Data Center Network = Fewer Ports, Less Complexity and Lower Costs

Joe Skorupa:

Joe Skorupa is a research vice president in Gartner Research, where he is part of the Data Center Transformation and Security Research unit. His research focuses on maximizing application performance via advanced networking technologies, including Layer 4-7 switching, and WAN optimization. He also covers data center network convergence, fabric-based infrastructure and delivery of rich media, including video within the enterprise. In addition, Mr. Skorupa advises vendors on product, market and partnership strategies. Through extensive collaboration, Mr. Skorupa's research extends to intersections between networking and adjacent research areas, including storage, servers and data center optimization. The research covers topics including Fibre Channel over Ethernet (FCoE), I/O virtualization, and emerging data center architectures, such as cloud computing and fabric-based computers.

Mr. Skorupa has more than 25 years of experience in enterprise, service provider and storage networking. He has held a variety of engineering, marketing and analyst roles at companies such as FORE Systems, Bytex, Motorola and RHK.

This document evaluates the potential benefits and drawbacks of converging data center networks, and debunks a myth about lower costs and complexity. It includes quantitative analysis, along with insights gained from discussions with Gartner clients about organizational issues that affect convergence decisions.

Key Findings

- Don't assume that a single converged Fibre Channel over Ethernet (FCoE) network is desirable, or even feasible.
- Standards for building large, scalable, Layer 2, converged Ethernet backbones are at least a year away. Products proven to be interoperable are much further off.
- Combining storage area network (SAN) and local-area network (LAN) traffic on a single backbone network increases costs and complexity.
- Organizational issues often dwarf the technical issues surrounding network convergence.
- Staff reductions are unlikely to be feasible even if physical networks are converged.
- Maintaining two separate data center networks doesn't mean you can't use the same technology for both.

Recommendations

- Plan to maintain separate data center SANs and LANs for at least the next three years.
- For virtualized servers, employ input/output (I/O) virtualization and network convergence to top of rack (ToR) switches to simplify cabling and reduce costs and complexity.
- To ensure competitive pricing, insist that vendors offering new technologies – such as Transparent Interconnection of Lots of Links (TRILL) – demonstrate interoperability with at least one credible alternative vendor before you buy.
- Consider Data Center Bridging (DCB) as a possible long-term (more than three-year) alternative to Fibre Channel (FC).
- Be wary of claims that you can build a standards-compliant, end-to-end converged data center network either today or within the next 24 months.
- Start integrating your server, storage and network teams under a single operations structure to prepare for longer-term synergies as their associated disciplines become more closely linked. For most organizations this represents a political challenge that will take years to complete, but until it is done operational efficiencies from staff integration are unlikely.

ANALYSIS

The Network Backbone Convergence "Buzz"

Once again the networking industry is abuzz with the promise of a single converged backbone infrastructure, this time in the data center core. Variously described as FCoE, Data Center Ethernet (DCE) and, more precisely, DCB, this latest development is intended to succeed where InfiniBand failed to unify computing, networking and storage networks. The argument goes like this: "It must be less expensive to build and manage one larger backbone network than two smaller ones." But this is a case where the facts don't support a seemingly obvious assumption.

The barriers to building a single network range from a dearth of available products, and the price premiums charged for those products, to the requirement to "forklift-upgrade" your entire data center backbone network in order to overcome long-standing organizational barriers. These barriers are expected to remain for at least the next three years. Over time, they may be lowered by product improvements and organizational integration, supporting a more convincing argument for backbone network convergence.

Note that Gartner believes that in-rack server network convergence can be both achievable and advisable. However, at current prices, a converged FCoE in-rack network can be more expensive than maintaining separate Ethernet and FC networks.

Convergence: Why 2<1

If convergence within the server and rack is a good idea, why shouldn't you extend it across the data center network backbone? There are several reasons.

Financial Reasons

The promise that a single converged data center network will require fewer switches and ports, be simpler, consume less power and require less cooling doesn't stand up to scrutiny. This is because as networks grow beyond the capacity of a single switch, ports must be dedicated to interconnecting switches. In large mesh networks entire switches do nothing but connect switches to each other. This results in a non-linear relationship between usable ports at the edge of the backbone and ports used for inter-switch links. In smaller networks fewer ports are required to perform this interconnect function.

A network topology known as a "fat tree" or "folded Benes" is popular when building large-port-count non-blocking networks using smaller-port-count switches. Figure 1 illustrates how a 32-port non-blocking network can be constructed using 8-port switches.

FIGURE 1

A 32-Port Non-Blocking Network in a Two-Tier "Fat Tree" Topology

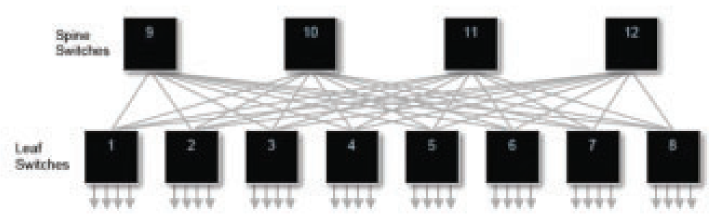

Source: Fulcrum Microsystems

FIGURE 2

Three-Tier Fat-Tree Network Topology

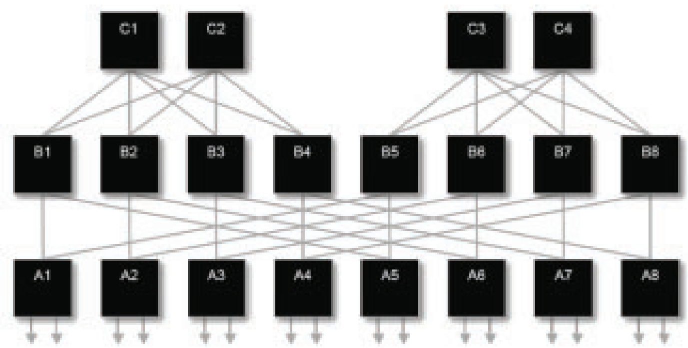

Source: Fulcrum Microsystems

Tables 1 and 2 show that, at best, the converged network requires the same number of ports as two separate networks. And if the same number of ports and switches is required, there are no savings on power or cooling. Assumptions for these configurations are detailed in Note 1.

In the top section of Table 1 all servers are connected to a DCB backbone through ToR switches only or embedded blade switches. Connections to existing FC storage arrays would require DCB-FC bridges, further increasing the cost of the converged backbone configuration.

The lower section of Table 1 shows the results of building separated DCB-based SAN and LAN. In both cases the backbone topology is a two-tier fat-tree network.

Table 1 assumes a fat tree of non-blocking 256-port core switches. The requirement for a non-blocking core is detailed in "Use Top-of-Rack Switching for I/O Virtualization and Convergence; the 80/20 Benefits Rule Applies" (see Note 2 for an update on this topic).

An additional attribute of fat trees is that they enable very large, fault-tolerant non-blocking networks to be constructed from a large number of relatively inexpensive low-port-count switching elements. The availability of low-cost 48-port ToR DCB switches suggests that, if TRILL support is added, a lower-cost converged core might be possible. In this particular case, using 48-port switches requires a three-tier – spine, aggregation and leaf – fat tree to construct the converged network, whereas only two-tier fat trees are required to construct the two separate networks (Figure 2 shows a three-tier fat-tree network topology).

This is because, in general, switches with N-ports enable two-tier fat-tree topologies of N2/2 external ports, and that two-tier topology requires 3N/2 switches to build a maximum configuration. So 48-port switches are limited to (482)/2 ports or 1,152 ports. To build the required 1,490-port backbone, a three-tier fat-tree topology is required.

Even when the converged backbone can be built with a two-tier network topology, the best-case scenario is that the converged network requires the same number of ports and costs the same as two separate networks. Although a network built with low-cost 48-port switches is less expensive than one using more expensive 256-port switches, the saving comes at the expense of complexity. In either case, networks constructed using smaller switches require significantly more interconnect cables deployed in a very complex topology, which makes installation and troubleshooting very difficult.

In much smaller configurations with only a few hundred servers the entire network could be constructed using a single very large data center backbone switch. While this topology is very simple, 256- to 512-port 10-Gbps Ethernet non-blocking switches are very expensive, often costing over $1,000,000. Additionally, configuration and management complexities remain.

Since convergence brings no reduction in the amount of equipment needing to be acquired, maintenance and support costs are unlikely to be lower. It should also be noted that committing to a single vendor to increase purchasing volumes may actually increase costs.

Reasons of Design and Management Complexity

With today's switches, the large Layer 2 topologies required by VMware's VMotion virtualization technology make Ethernet network design a very complex multi-overlapping VLAN problem. Designing a large SAN is no simpler. And when the two networks are overlaid on a single infrastructure the complexity increases significantly. As traffic shares ports, line cards and inter-switch links, avoiding congestion ("hot spots") becomes extremely difficult. Over time, emerging standards such as TRILL may make it easier to avoid these hot spots, but mature, standards-compliant implementations are at least two years away. Avoiding hot spots and single points of failure in the event of switch or link failure is a very large design challenge.

Debugging problems is also more difficult in the converged network, since interactions between LAN and SAN traffic can make root-cause analysis harder. Since many problems are transient in nature, events must be correlated across the two virtual networks, which increases complexity. Should an outage be required to solve a problem or simply to perform maintenance, a downtime window that is acceptable for both environments may be required. This increases complexity and may increase costs as well.

Table 1. Price Comparison: Single Converged DCB Backbone Network Versus Separate DCB-Based LAN and SAN, All Using 256-Port Switches

Configuration	Spine Ports	Leaf Ports	Ports per Switch	Leaf Switches	Spine Switches	Total Switches	Total Ports	Price per Port	Price
5,000 servers, DCB end to end, not counting FCoE bridges at the SAN edge	1,490	2,980	256	11.75	6.00	17.75	4,544	$1,000	$4,544,000
5,000 servers, FCoE ToR, two backbones, partial fill on core chassis									
Ethernet backbone	544	1,088	256	4.25	2.25	6.50	1,664	$1,000	$1,664,000
FCoE SAN backbone	946	1,892	256	7.50	3.75	11.25	2,880	$1,000	$2,880,000
Total	1,490	2,980	256	11.75	6.00	17.75	4,544	$1,000	$4,544,000
DCB= Data Center Bridging; FCoE = Fibre Channel over Ethernet; SAN = storage area network; ToR = top of rack For the assumptions underlying this table, see Note 1									
Source: Gartner (March 2010)									

Table 2. Price Comparison: Equivalent Backbones to Those of Table 1, Constructed Using 48-Port Switches

Configuration	Spine Ports	Aggregation Ports	Leaf Ports	Ports per Switch	Leaf Switches	Aggregation Switches	Spine Switches	Total Switches	Total Ports	Price per Port	Price
5,000 servers, DCB end to end, not counting FCOE bridges at the SAN edge	1,490	2,980	2,980	48	63	63	32	158	7,584	$500	$3,792,000
5,000 servers, FCOE ToR, two backbones, partial fill on core chassis											
Ethernet backbone	544		1,088	48	23		27	50	2,400	$500	$1,200,000
FCOE SAN backbone	946		1,892	48	40		22	62	2,976	$500	$1,488,000
Total	1,490		2,980	48	63		49	112	5,376	$500	$2,688,000
DCB= Data Center Bridging; FCoE = Fibre Channel over Ethernet; SAN = storage area network; ToR = top of rack For the assumptions underlying this table, see Note 1											
Source: Gartner (March 2010)											

The simple solution is to segment the switches to isolate LAN and SAN traffic from each other. Alternatively, of course, you could simply maintain two separate networks and avoid the problem altogether.

Organizational Reasons

Perhaps the greatest impediment to backbone network convergence is organizational. Simply put, in most large organizations, the SAN and LAN administration teams report to different managers, have very different cultures, and don't get along. Most LAN staff see storage staff as Luddites stuck in the previous century, while most storage staff view LAN staff as people who don't know how to run a production network. Additionally, the SAN staff remember what happened to the voice network engineers when Internet Protocol (IP) telephony was introduced – they were absorbed into the LAN team as "second-class citizens."

Although Gartner recommends that server, storage and network teams be integrated, for most organizations this will prove a political challenge that takes years to resolve. But until this happens, improved operational efficiencies from staff integration are unlikely.

Gartner's IT Key Metrics data for voice networks and data networks shows that convergence of these networks did not result in significant staff reductions. Some external contractors may have been eliminated, but the unique skills required prevented significant reductions in staffing levels. Although moves, adds and changes are simplified, more highly skilled personnel are required to design, implement and operate the converged infrastructure. We see no reason to think that the results of LAN-SAN convergence would differ from those for time division multiplexing (TDM) voice-VoIP/LAN convergence.

Converge All Data Center Traffic on a Single Technology, but Not a Single Network

There are benefits to standardizing on a single technology for all data center networking if that technology adequately supports the needs of applications. Doing this simplifies acquisition, training and "sparing." However, settling on a single technology does not require that separate networks be combined. Design, operations and troubleshooting are much easier with separate networks and, as this document demonstrates, they may also cost less to build.

Note 1. Assumptions for Tables 1 and 2

Assumptions for Table 1

- 5,000 servers, in a mix of 1RU, 2RU and blade server chassis, supporting a mix of dedicated and virtualized workloads.
- 256-port non-blocking switches at $1,000 per port.
- Two-tier fat-tree (folded Benes) network topology.
- All devices are dual-pathed to provide redundancy and some margin of headroom.
- Server traffic is aggressively aggregated in ToR switches to minimize required backbone ports.
- Less aggressive aggregation would reduce the need for a non-blocking core, but would increase the number of leaf ports and negate any significant savings in core ports.

Assumptions for Table 2

- 5,000 servers, in a mix of 1RU, 2RU and blade server chassis, supporting a mix of dedicated and virtualized workloads.
- 48-port non-blocking switches at $500 per port.
- Three-tier fat tree (folded Benes) network topology due to smaller switch size.
- All devices are dual-pathed to provide redundancy and some margin of headroom.
- Server traffic is aggressively aggregated in ToR switches to minimize required backbone ports.
- Less aggressive aggregation would reduce the need for a non-blocking core, but would increase the number of leaf ports and negate any significant savings in core ports.

Traditionally, the technology that adapts to assimilate others has been Ethernet. If DCB proves a suitable substitute for Fibre Channel and 40/100-Gbps Ethernet is delivered in a timely manner, it will make sense over time to move all traffic to Ethernet. In our model we assume that 10-Gbps DCB switch ports will cost approximately the same as normal 10-Gbps Ethernet ports and 8-Gbps FC switch ports.

We expect DCB to become standard on most data-center-oriented switches. Since 10-Gbps Ethernet already provides a much higher data rate that 8-Gbps FC (10 Gbps versus 6.4 Gbps), DCB promises better price/performance. However, due to the size of the installed base of FC, the promise of the all-Ethernet data center will take at least five years for most enterprises to realize.

The Payoff

Although the promise that a unified fabric will require fewer switches and ports, and result in a simpler network that consumes less power and needs less cooling, may go unfulfilled, this doesn't mean that enterprises should forgo the benefits of adopting a unified network technology. In fact, this approach may prove that in some cases, 2<1.

Gartner RAS Core Research Note G00174456,
Joe Skorupa, 11 March 2010

Note 2. The Need for a Non-Blocking Converged Core

Recently there has been some suggestion that heavily over-subscribed, and therefore lower-cost, backbones will be suitable in heavily virtualized environments if virtual machine (VM) affinity is employed. VM affinity keeps all the VMs associated with a particular application co-located on the same blade chassis or in the same rack. While this may remove some server-server TCP/IP traffic from the backbone, it fails to account for the following:

- Approximately two-thirds of all server network input/output is SAN traffic, and this traffic will still have to cross the backbone to reach storage arrays. Distributing ("Balkanizing") storage to localize that traffic reduces efficiencies and increases complexity. And as VMs move across the data center to balance loads or to recover from failures, the storage will remain in place, causing the traffic to cross the backbone.
- Some portion of the LAN traffic comprises end-user interactions, which will cross the backbone.
- As VMs move around the data center to balance loads and recover from failures, traffic patterns are at best difficult to predict. The resulting network congestion will impair application performance, negating any savings from reduce capital costs.
- VM mobility is only practical if very-high-performance links are available between source and destination physical servers. Unanticipated congestion during VM migrations can cause application degradation or disruption.
- Although it is not a storage best practice, many data centers still use server-based backups. These backups can generate an order of magnitude more LAN traffic than they do application traffic. An oversubscribed converged backbone can result in failed backups and interruptions to running applications.

About Avaya

[Avaya Data Solutions]

[Creating a Fit for Purpose Network Architecture]

[Gartner Newsletter Issue 1: Creating a Fit for Purpose Network Architecture]

[Avaya Fit For Purpose Data Networking Solutions Podcast]

[Avaya Data Solutions - A Heritage of Technology and Innovation Podcast]

[Avaya Virtual Enterprise Network Architecture Podcast]