Fit-for-Purpose Data Center Networking

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

Table of Contents

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.

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.

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[1] 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 data centre and the 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.

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

[1] Gartner. Myth: A Single FCoE Data Center Network = Fewer Ports, Less Complexity and Lower Costs [11 March 2010].

The technology of choice for creating the foundations of network transparency will be that of Shortest Path Bridging[2], 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.[3] 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.

---

[2]SPB has two variants – MAC-on-MAC (SPBM) and VLAN ID (SPBV) – and this paper focuses on SPBM.

[3]Provider Link State Bridging (PLSB) is the fore-runner to SPBM, and was presented to the IEEE by the-then Nortel, now part of Avaya.

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.

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

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.

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 Authorized Partner, or visit us at **www.avaya.com**.

About Avaya

Avaya is a global provider of business collaboration and communications solutions, providing unified communications, contact centers, data solutions and related services to companies of all sizes around the world. For more information please visit www.avaya.com.

The Power of We™

06/11 • DN4621-02

avaya.com