# Provider Backbone Bridging and MPLS: Complementary Technologies for Next-Generation Carrier Ethernet Transport

*Samer Salam and Ali Sajassi, Cisco Systems*

## ABSTRACT

Although provider backbone bridging is sometimes cast as an alternative technology to MPLS for Ethernet transport, the fact is that both technologies can be leveraged to complement one another in a service provider network. This provides the network operator with the best of what each technology has to offer in terms of scalability, manageability, cost, and flexible support for services. An example of the two technologies working in unison is provider backbone bridging interoperating with H-VPLS over networks with an MPLS core. The combined infrastructure leverages the strengths of Ethernet and MPLS with the added advantage of addressing the major shortcomings of standard H-VPLS.

## INTRODUCTION

Ethernet is rapidly becoming the access technology of choice for service provider next-generation networks. This is market driven and can be attributed to the advantages of the technology. Driving the market is the ongoing surge in end customer bandwidth requirements fueled by new high-bandwidth applications such as video, network-based storage, and online gaming, to name a few. Second, there is the pressure on service providers to reduce capital and operating expenditures, combined with a drive to introduce high-margin services to enable those providers to sustain profitability and increase revenues. Ethernet technology provides the cost and flexibility characteristics that service providers require. It has a lower cost-per-bit compared to other carrier access technologies. Furthermore, it is an enabler for both Layer 2 and Layer 3 services (doubling as a service, as well as a transport). To that, one can add the fact that recent and ongoing standardization efforts in IEEE and the Metro Ethernet Forum (MEF) are addressing the scalability, manageability, and other carrier-grade aspects of Ethernet; for example, [1] and [2] define the operation of Ethernet bridges in service provider access and backbone networks, respectively. All this positions Ethernet as an access technology for next generation service provider and carrier networks.

Service provider core networks are converging toward employing Internet Protocol (IP) and multi-protocol label switching (MPLS) technologies. This trend towards IP/MPLS core has existed for some time, with service providers heavily investing in those technologies. It is hard to envision this infrastructure being replaced by Ethernet bridging (e.g., provider backbone bridging). This is especially true because IP/MPLS core networks benefit from the full suite of IP and MPLS control protocols, which were developed and enhanced over a period of more than ten years, and accommodate both Layer 3 and Layer 2 (including Ethernet) transport services. These control protocols cover a broad spectrum of functions, including: flexible and scalable unicast routing, multicast, resource reservation, virtual private networks, and traffic engineering [3]. This is in addition to accommodating inter-autonomous system operations and inter/intra area scenarios. The counterpart for some of these protocols does not even exist in Ethernet and would take years to develop and to reach the same level of maturity as in IP/MPLS. Hence, provider backbone bridging will not replace IP/MPLS technologies but rather complement them in the context of carrier networks — Ethernet for efficient access and IP/MPLS for large-scale aggregation.

One embodiment of Ethernet and IP/MPLS technologies working in unison is provider backbone bridging (PBB) interoperating with hierarchical virtual private LAN service (H-VPLS). This embodiment is the focus of this article. First, a brief overview of the PBB and the H-VPLS technologies is presented, along with the shortcomings of existing H-VPLS. Then, PBB and H-VPLS interoperability is discussed, including the advantages of the solution. This is followed by a summary of related technologies proposed in this area. Finally, a snapshot of the current status of the evolving standards relating to this area is provided.


0163-6804/08/$25.00 © 2008 IEEE

*The customer VLANs remain invisible to the service provider. The service VLAN serves the purpose of uniquely identifying the customer service instance in the service provider's network and allows for the transparency of customer VLANs within the carrier's network.*

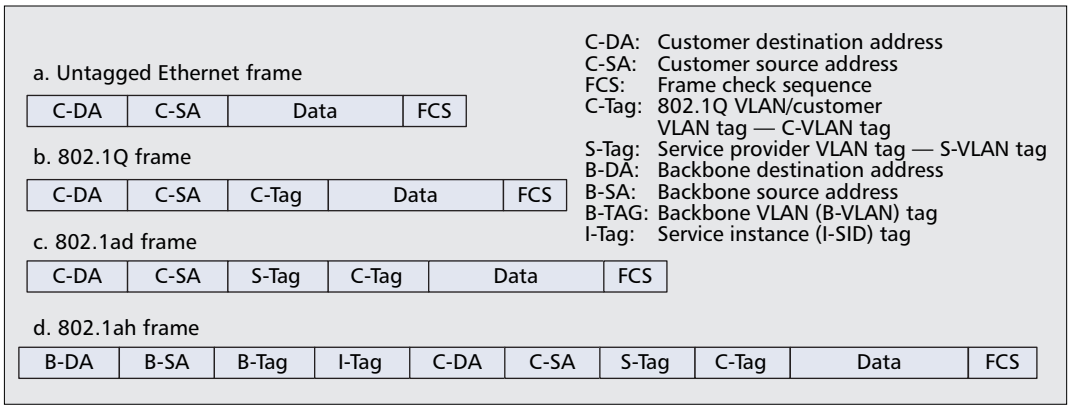


**Figure 1.** *Various Ethernet frame formats.*

## Provider Backbone Bridges

Given its roots as a local area network (LAN) technology, Ethernet initially lacked some of the flexibility, security, and scaling characteristics of carrier technologies. In the beginning, Ethernet offered a shared transport medium where all end devices had connectivity to each other and were placed on a common broadcast domain. Later, IEEE standard 802.1Q defined the concept of a virtual LAN (VLAN) as a means of creating independent logical LANs, with disjoint broadcast domains, over a single physical network. The 802.1Q standard extended the Ethernet frame format to include a VLAN tag (Fig. 1b), which was added by switches, to identify the frame's virtual LAN over inter-switch links. A switch would admit and forward frames only on ports that were configured with the same VLAN identifier. The motivation for the VLAN was two-fold: providing security over a shared physical network through logical traffic segregation and more efficient bandwidth utilization by limiting the scope of flooded broadcast or multicast traffic to a VLAN [4] rather than the entire physical network. This was developed in the context of modest sized LANs, where a few hundred VLANs were sufficient to address most, if not all, deployments. Given those constraints, IEEE 802.1Q converged on assigning 12 bits for the VLAN identifier. This limited the number of identifiers to 4094 (the identifiers, 0 and 4095, were reserved by the standard).

As service providers started venturing into Metro Ethernet services, they leveraged 802.1Q to offer services to their end customers over bridged networks. A service provider would assign each customer service instance its own VLAN identifier [5] to guarantee traffic security and optimal bandwidth utilization. This limited the maximum number of service instances that a provider could offer over a network to 4094 — the number of allowed VLANs. The latter was clearly a scalability limitation for service providers that prevented them from adding services and increasing revenue on their installed networks.

Another issue that faced service providers was that 802.1Q did not provide a mechanism to distinguish the customer VLAN identifier space from the service provider identifier space. For example, if a service provider offers an intercity transparent LAN service to an enterprise and that enterprise required some number *N* of VLANs, then the service provider must assign the enterprise *N* VLAN identifiers for the single service. This implied that the service provider had to coordinate the identifier space amongst its customers to prevent two customers from using the same VLAN, a matter which introduced operational complexity.

IEEE standard 802.1ad (now part of IEEE standard 802.1Q-2005) solves this problem by introducing a service VLAN tag that is pre-pended to the 802.1Q VLAN — hereafter referred to as the customer VLAN — in the Ethernet frame (Fig. 1c). The service VLAN tags are added by provider edge bridges to customer frames as they ingress into the service provider's network. The tags are removed from those frames, again by provider edge bridges, as the frames egress the provider's network. The frames are bridged within the provider's network based solely on the service VLAN. As such, the customer VLANs remain invisible to the service provider. The service VLAN serves the purpose of uniquely identifying the customer service instance in the service provider's network and allows for the transparency of customer VLANs within the carrier's network. For example, in a network employing 802.1ad Ethernet technology, a service provider may use service VLAN 100 for a transparent LAN service offered to *enterprise Y* and service VLAN 200 for another service to *enterprise Z*. *Enterprise Y* has five customer VLANs (11 through 15), whereas *enterprise Z* has 10 customer VLANs (11 through 20). Note how the two customers have overlapping customer VLANs (11 through 15). This would not have been possible without 802.1ad technology.

It is worth noting that IEEE 802.1ad does not solve the problem of service instance scalability: The service VLAN field remains 12 bits in width. Thus, service providers are still limited to a maximum of 4094 service instances per bridged network.

Furthermore, with IEEE 802.1ad and the original IEEE 802.1Q, service provider equipment operating at the Ethernet service layer must learn the MAC addresses of the customer devices, a matter that poses yet another scalability challenge. Because the service provider equipment is bridging customer frames, it must learn the addresses of either the customer edge devices (when those devices are routers) or the addresses of the customer end stations (when the customer

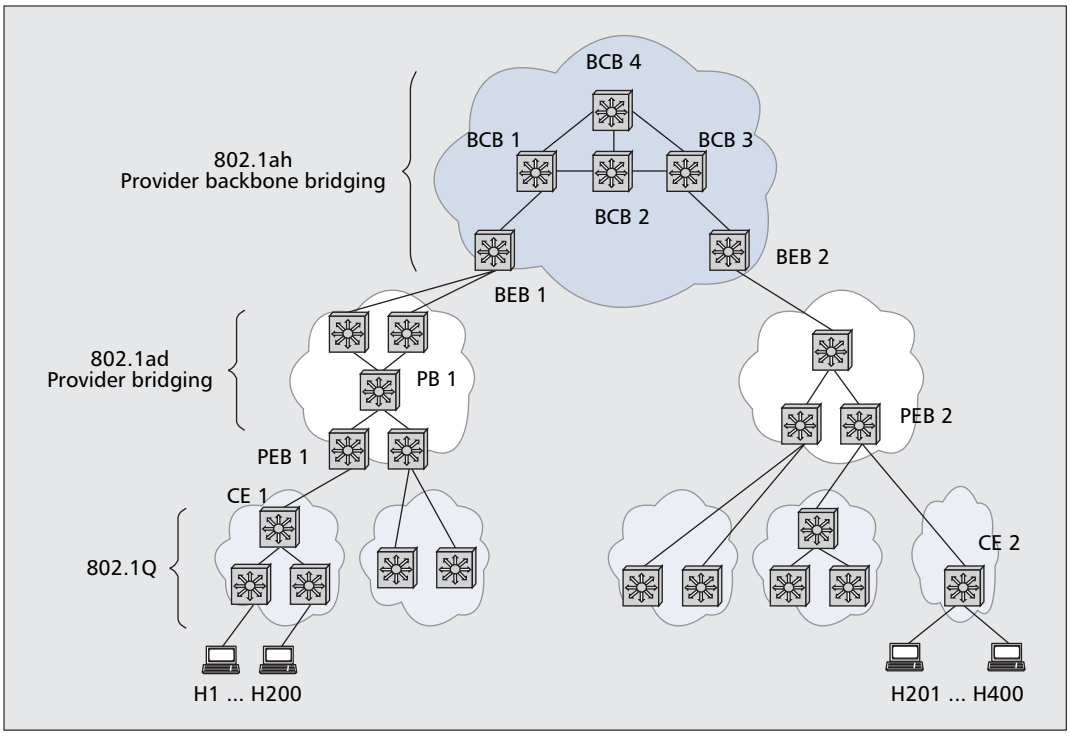


**Figure 2.** *Provider backbone bridging network hierarchy.*

*IEEE draft standard 802.1ah, or PBB, aims to solve the service instance and MAC address scalability problems. PBB extends IEEE 802.1ad and introduces a hierarchical network architecture model that enables service providers to build large bridged networks.*

edge devices are bridges). Disabling address learning is not an option, especially in the context of multipoint services, because it leads to excessive flooding of unknown unicast frames. This, in turn, wastes the bandwidth of the network. As such, service providers are bound by the size of their equipment hardware address tables when attempting to scale the number of customers or customers' sites over their existing networks. For example, consider a service provider offering a transparent LAN service for an enterprise with 50 sites and 200 end stations per site. For this service, the service provider equipment must learn 10,000 addresses (50 × 200). If the provider were to accommodate only 10 such customers, the provider's devices would each have to learn 100,000 address entries. This exceeds the capability of most Ethernet switches on the market, because typical state-of-the-art Ethernet switches have MAC address table sizes ranging from 4,000 to 64,000 entries [6].

IEEE draft standard 802.1ah, or PBB, aims to solve the service instance and MAC address scalability problems highlighted previously. PBB extends IEEE 802.1ad and introduces a hierarchical network architecture model that enables service providers to build large bridged networks. In this model, networks of IEEE 802.1Q bridges are aggregated into networks of IEEE 802.1ad provider bridges, which in turn, are aggregated into a network of provider backbone bridges. Figure 2 depicts this hierarchy in an example network.

To solve the MAC address scalability problem, PBB introduces a new frame format (Fig. 1d) that provides a MAC tunneling encapsulation scheme: customer Ethernet frames are encapsulated within provider Ethernet frames as they ingress the PBB network, thereby hiding the customer addresses from the PBB core. Devices in the core of a PBB network forward traffic based on backbone MAC (B-MAC) addresses. This effectively confines the requirement to learn customer addresses to the edge devices of the PBB network — these edge devices are called backbone edge bridges. A given backbone edge bridge is required to learn the addresses of only those customers that it supports, and a given core device is required to learn the addresses of only the backbone edge bridges (as opposed to having to learn addresses of all of the end customer devices). This greatly enhances the scalability of the solution. For example, in Fig. 2, assume that a service provider employing 802.1ad access and PBB core networks is offering a transparent LAN service to a customer with two sites corresponding to *CE1* and *CE2*. Furthermore, assume that 200 customer end stations (*H1* through *H200*) are sitting behind customer edge bridge *CE1* , and another 200 end stations (*H201* through *H400*) are behind customer edge bridge *CE2*. The provider devices that transport the service frames between the two customer sites that employ 802.1ad technology (e.g., devices *PEB1*, *PB1*, and *PEB2*) will learn 400 customer addresses, corresponding to all end stations. Shifting our focus to the PBB section of the network: backbone edge bridge *BEB1* will learn customer addresses for the directly connected site, namely, the addresses corresponding to hosts *H1* through *H200*, in addition to the customer addresses for the remote site of the service, corresponding to hosts *H201* through *H400* (a total of 400 addresses). Similarly, *BEB2* will learn a total of 400 addresses. However, the bridges in the core of the PBB network will have only two MAC address entries in their forwarding tables for that service: these correspond to the backbone

A pseudowire is an emulated point-to-point link that consists of two unidirectional label switched paths (LSPs) that enable the extension of an Ethernet (and other technologies) physical wire over a packet-based network.

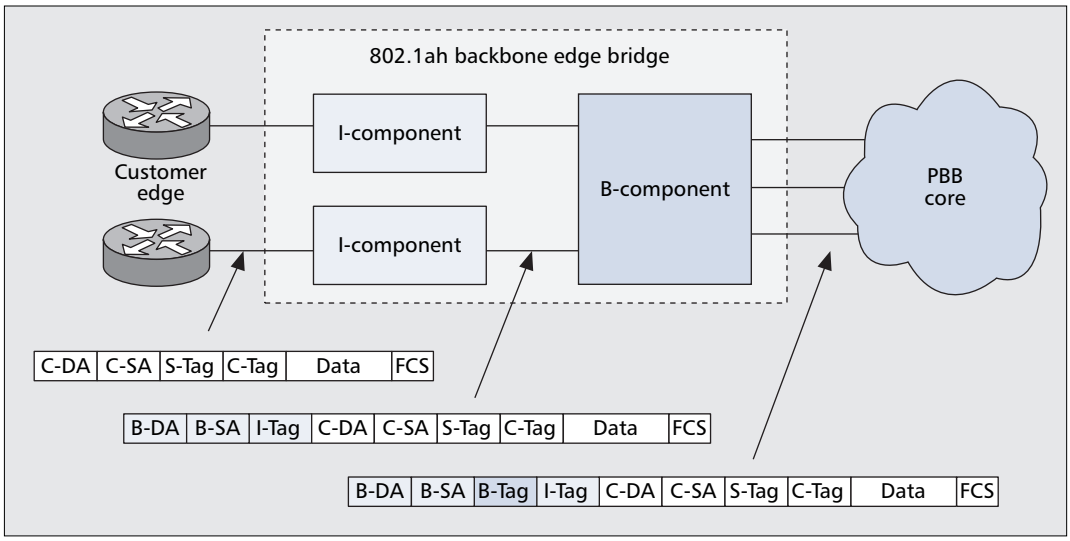


**Figure 3.** *IEEE 802.1ah backbone edge bridge model.*

addresses of devices *BEB1* and *BEB2*. This is two orders of magnitude less than the MAC table size of the 802.1ad bridges. The reason why this is the case is because *BEB1* and *BEB2* encapsulate the customer frames with PBB encapsulation, using their own backbone addresses in the source address field, before forwarding those frames to the PBB core.

To solve the customer service instance scalability problem, the PBB frame format introduces a 24-bit service identifier field known as the I-SID. Each customer service instance is assigned a unique I-SID value that is global within a network administrative domain. Hence, the new PBB frame format effectively expands the number of service instances from 4094 to a theoretical maximum limit of roughly 16 million ($2^{24}$). Note that I-SIDs are visible to backbone edge bridges only and are transparent to the devices in the core of the PBB network.

The PBB frame format also includes a 12-bit backbone VLAN identifier (B-VLAN). This allows the service provider to partition their PBB network into different broadcast domains, for efficient bandwidth utilization, especially in multicast applications [7]. The B-VLAN also enables the service provider to perform load-sharing within the core of the PBB network, by means of bundling different I-SIDs into distinct B-VLANs and mapping the latter into different spanning-tree instances [2].

PBB can be seen as a natural evolution of 802.1ad: for a service provider to upgrade their existing 802.1ad networks to PBB, they need to upgrade only the edge devices to implement backbone edge bridge functionality. The core of the 802.1ad network remains unchanged. Internally, in the broadest sense, a backbone edge bridge comprises of two types of building blocks (e.g., bridge internal components): one or more I-components and a single B-component. The I-component connects to the customer equipment and is responsible for learning customer addresses, for encapsulating customer frames with backbone addresses, and for adding the I-SID. The B-component connects to the PBB core devices. It is responsible for adding the B-VLAN and bridging backbone-encapsulated frames based on backbone addresses. Figure 3 shows the model.

This section covered provider backbone bridging technology and the problems it addresses. In the next section, we shift focus to provide a brief background of hierarchical virtual private LAN service to set the context for discussing the interoperability between the two technologies and the advantages of the combined solution.

## Hierarchical Virtual Private LAN Service

H-VPLS provides a multi-tier hierarchical architecture to implement multipoint Ethernet-based Layer 2 VPN services over IP/MPLS through the use of pseudowires [8]. A pseudowire is an emulated point-to-point link that consists of two unidirectional label switched paths (LSPs) that enable the extension of an Ethernet (and other technologies) physical wire over a packet-based network [9].

In its simplest form, the H-VPLS architecture (Fig. 4) consists of a top-tier IP/MPLS core comprised of network provider edge (N-PE) nodes connected via a full mesh of pseudowires. The second tier is a set of access networks that connect user-facing provider edge (U-PE) nodes to the N-PEs. These U-PE nodes connect directly to the customer equipment. The access network may employ either MPLS pseudowires (the right-hand side access network in Fig. 4) or IEEE 802.1ad Ethernet transport (the left-hand side access network in Fig. 4). The U-PEs, if bridging-capable, may attempt to locally switch the customer traffic; otherwise, they would forward the traffic to an N-PE node. The latter would then switch the customer frames by means of bridging them at the MAC layer. In this scheme, the N-PE is forced to learn and cache all customer addresses from frames received through all local U-PE nodes, as well as all remote N-PE and U-PE nodes (i.e., those on the other side of the core). The learning is performed on a per customer basis. This mandates having a dedicated forwarding table at each of the N-PE and U-PE devices for each customer

instance, as well as having a dedicated set of pseudowires for each customer instance. The net result is a substantial number of pseudowires required in the network and a substantial number of addresses to be stored in N-PE devices.

The U-PE nodes offer a standard IEEE 802.1Q or IEEE 802.1ad service interface to the customer equipment. To guarantee customer VLAN transparency, the U-PEs can implement IEEE 802.1ad provider edge functionality so that a service VLAN is imposed on ingress customer traffic. This service VLAN represents a customer service instance throughout the provider's network and can be associated with a VPLS instance on the N-PE nodes. U-PE nodes are responsible for removing the service VLAN from the frames before delivering them to customer equipment.

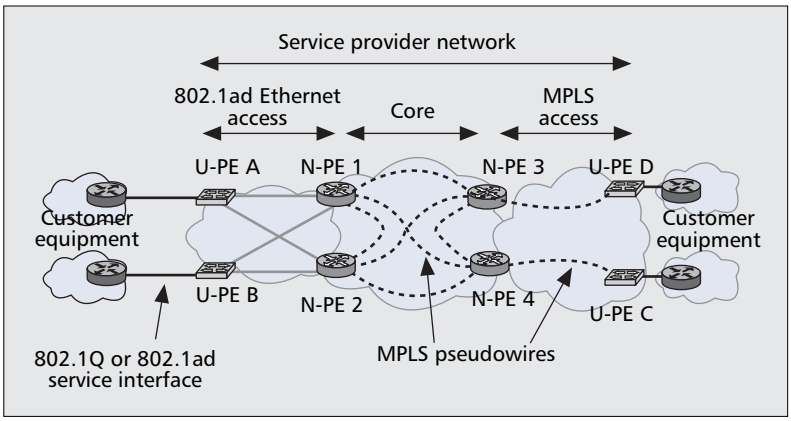


■ **Figure 4.** *Hierarchical virtual private LAN service (H-VPLS) architecture*

## PBB AND IP/MPLS INTEROPERABILITY

By combining PBB technology with IP/MPLS, service providers can leverage the strengths of each technology at the proper locality within the network. The advantages over existing solutions will become evident in the course of the discussion on two illustrative deployment topologies: H-VPLS with Ethernet access and H-VPLS with MPLS access networks.

### PBB IN H-VPLS WITH AN ETHERNET ACCESS NETWORK

Service providers can overcome the scalability limitations of traditional H-VPLS by employing PBB technology instead of IEEE 802.1ad in the access networks. Customer edge devices can connect to these PBB access networks via one of the existing Ethernet interfaces: IEEE 802.1Q or IEEE 802.1ad.

The interoperability technique is best explained with an example. Consider the network in Fig. 5. The device *U-PE A*, which interfaces with the customer equipment, is equipped with PBB backbone edge bridge functionality. It is responsible for encapsulating customer frames (e.g., frames from *H1* to *H4*) with the PBB header and bridging those encapsulated frames to *N-PE A*. Note that *U-PE A* uses its own backbone address (*A1*) as the source address and *U-PE B*'s backbone address (B1) as the destination address in the PBB encapsulation. *N-PE A* then forwards the PBB-encapsulated customer frames onto the VPLS core network. Upon traversing the core and reaching the remote *N-PE B*, the frames are bridged within the remote PBB network. The frames eventually reach *U-PE B*, which removes the PBB encapsulation and delivers the customer frames to the remote customer edge device (*CE 4*). Those frames eventually reach the target end-station *H4*. Given that the PBB header adds provider source and destination backbone addresses to the customer frame, all forwarding within the service provider network will be performed based on the backbone addresses instead of the customer addresses.

The key to PBB and H-VPLS interoperability is in the implementation of the N-PE. Existing N-PE implementations do not understand the PBB frame format. However, that does not mean that those implementations cannot be used for interoperability with PBB. As long as it is not required to look at the 24-bit I-SID on the N-PE, the existing implementation can be leveraged for interoperability. The N-PE, in this case,

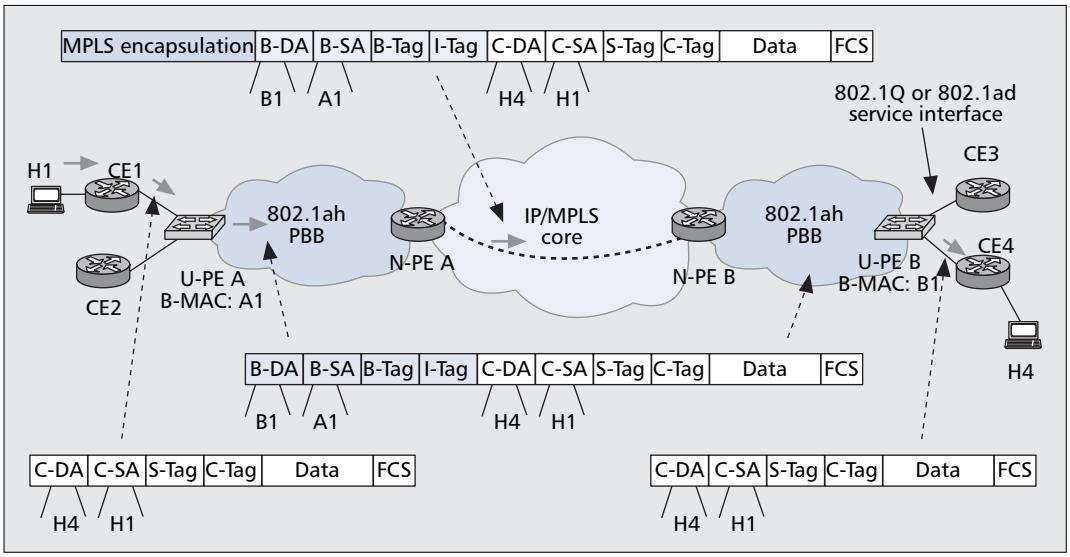


■ **Figure 5.** *H-VPLS with PBB access network.*

*For H-VPLS with MPLS access networks, incorporating PBB functions at the U-PE serves to improve the scalability of the network in terms of both the numbers of MAC addresses and the number of service instances that can be supported.*

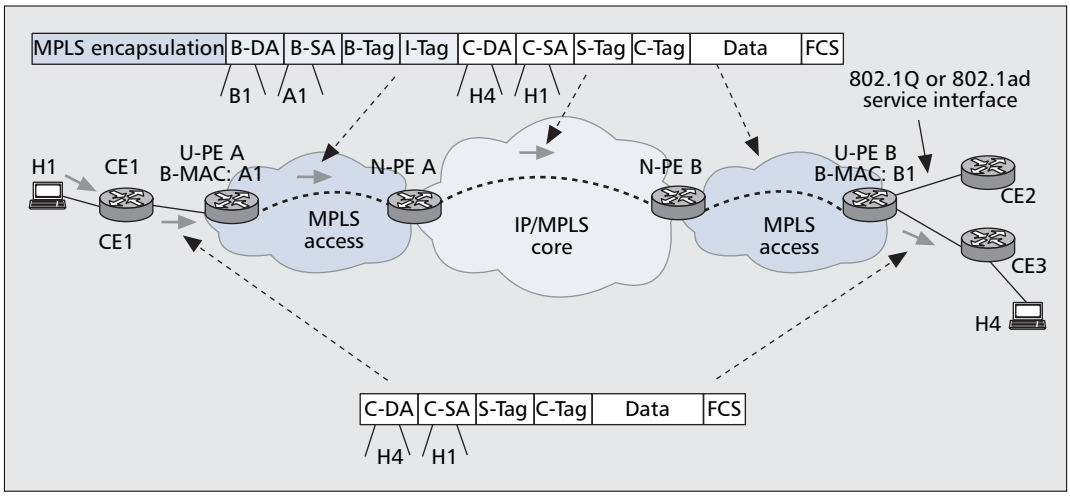


**Figure 6.** *PBB in H-VPLS with MPLS access network.*

treats the B-VLAN as if it were an 802.1ad service VLAN tag and forwards these frames to the VPLS core using existing H-VPLS mechanisms as defined in [8]. The N-PE is completely oblivious to the fact that it is forwarding PBB frames; it considers them regular 802.1ad frames. When I-SID visibility is required at the N-PE, then additional capability to parse the PBB encapsulation and identify the I-SID is required, along with a new pseudowire type to transport PBB-encapsulated frames [10]. Furthermore, new service interfaces based on an I-SID tag also may be required depending on the position of the N-PE within the provider's administrative service boundary. In broad terms, I-SID visibility is a must on the N-PE when either a VPLS instance is defined on a per I-SID basis in the core, or when the N-PE is connected to the PBB network via an I-SID tagged service interface marking the boundary of an administrative service domain. When the N-PE is in a different administrative service domain than the PBB access network, then I-SID translation is required between the two service domains [11]. When I-SID visibility is required on the service interface of the N-PE, the latter is viewed to offer service-level interworking between the PBB network and the MPLS domain. Conversely, when the N-PE is not required to have I-SID awareness on its service interface, the N-PE is considered to offer network-level interworking between PBB and the MPLS domain.

## PBB in H-VPLS with MPLS Access Network

For H-VPLS with MPLS access networks, incorporating PBB functions at the U-PE serves to improve the scalability of the network in terms of both the numbers of MAC addresses and the number of service instances that can be supported.

Customer edge devices connect to a U-PE node using standard Ethernet interfaces: IEEE 802.1Q or IEEE 802.1ad. The U-PE is connected upstream by MPLS pseudowires to one or more N-PE nodes. The N-PE nodes are connected by a full mesh of pseudowires (per VPLS instance) spanning the IP/MPLS backbone. A U-PE is outfitted with PBB backbone edge bridge functions so it can encapsulate and de-capsulate customer frames using the PBB encapsulation and perform I-SID translation if required. For example, in the network of Fig. 6, customer frames from *H1* destined to *H4* are first encapsulated by U-PE A with PBB encapsulation. The header of the PBB encapsulation includes the backbone destination address of *U-PE B* (*B1*) and the backbone source address of *U-PE A* (*A1*). *U-PE A* also imposes the MPLS pseudowire encapsulation to the frame before forwarding it to *N-PE A*. *N-PE A* will perform a lookup on the backbone destination address and forward the frame on the right pseudowire over the VPLS core to *N-PE B*. The latter then performs a MAC address lookup (again, on the backbone address) and forwards the frame to *U-PE A*. Both *N-PE A* and N-PE B retain the header of the PBB encapsulation intact. *U-PE B* then disposes of the MPLS encapsulation, as well as the PBB encapsulation, and forwards the original customer frame toward the customer device.

In current H-VPLS design, the N-PE is forced to learn the customer addresses of all VPLS instances in which it participates. As the number of customers increases, this can easily add up to millions of addresses at the N-PE. However, if the U-PE performs PBB encapsulation, then the N-PE would be required to learn only the addresses of the U-PEs, which leads to a significant reduction in the MAC addresses. In addition, when PBB encapsulation is used, the U-PE may multiplex many I-SIDs into a single B-VLAN. If the VPLS instance is set up per B-VLAN, then it is possible to achieve a significant reduction in the number of pseudowires. It should be noted that this reduction in pseudowires comes at the cost of potentially increased replication over the pseudowire full mesh. A given customer's multicast and/or broadcast frames are effectively broadcasted within the B-VLAN. This may result in additional frame replication because the full mesh of pseudowires corresponding to a B-VLAN most likely spans more N-PEs than a full mesh of pseudowires corresponding to a single I-SID.

However, if one supports VPLS multicast data via MPLS point-to-multipoint tunnels [12], this drawback is rendered inconsequential.

### Advantages of the Solution

There are a number of advantages associated with interoperating PBB technology with H-VPLS. First, the introduction of the backbone address space enables the systems in the network to scale better because bridging is based on backbone addresses instead of the numerous end customer addresses. Second, the larger 24-bits I-SID space, compared to the 12-bit VLAN identifier space, increases the number of service identifiers by several orders of magnitude. Instead of being confined to 4094 service instances per access network, the provider can, at least in theory, accommodate $2^{24}$ service instances. Third, by having an N-PE map a B-VLAN to a VPLS instance and bundle multiple end-customer service instances over the same B-VLAN, it is possible to significantly reduce the number of full-mesh pseudowires required within the core. In this case, I-SID visibility is not required on the N-PE, and the I-SID serves the purpose of multiplexing/de-multiplexing customer service instances within a bundle (B-VLAN). Hence, instead of maintaining a full mesh of pseudowires per service instance as with the current H-VPLS, it is possible to maintain a full mesh per group of service instances. Finally, the scaling advantages associated with H-VPLS, including reduced pseudowire signaling overhead, remain in effect.

## Related Work

An alternative solution to the discussion in this article involves deploying Ethernet bridges end-to-end in the provider's network. In such an architecture, the access tier of the network employs IEEE 802.1ad technology, and the core uses PBB instead of MPLS [2, 7]. This is not a very attractive proposition for service providers that already have invested in MPLS for their core.

Another related technology, under study in the Internet Engineering Task Force (IETF), is transparent interconnection of lots of links (TRILL). It introduces the notion of routing bridges that encapsulate end-station Ethernet frames in a second Ethernet header, and it uses the intermediate system-to-intermediate system (IS-IS) routing protocol in lieu of the Spanning Tree Protocol (STP) to circumvent the issues with STP. This bears some conceptual similarities to PBB; however, it is intended for LAN deployments rather than service provider networks [13, 14].

## Conclusion

By combining PBB technology in the access with IP/MPLS in the core, service providers can leverage the strengths of both technologies and overcome scalability limitations in their networks. The combined solution addresses the two problems of service instance, as well as MAC address scalability in carrier networks. At the time of this writing, provider backbone bridging is still under standardization as IEEE 802.1ah. In addition, interoperability of provider backbone bridging with VPLS is being discussed in the IETF Layer 2 VPN workgroup [10, 11].

*By combining PBB technology in the access with IP/MPLS in the core, service providers can leverage the strengths of both technologies and overcome scalability limitations in their networks.*

## References


[1] IEEE 802.1Q-2005, "Local and Metropolitan Area Networks Virtual Bridged Local Area Networks," May 2006.
[2] IEEE P802.1ah/D3.6, "Virtual Bridged Local Area Networks — Amendment 6: Provider Backbone Bridges," June 2007.
[3] G. Swallow, "MPLS Advantages for Traffic Engineering," *IEEE Commun. Mag.*, vol. 37, no. 12, Dec. 1999, pp. 54–57.
[4] V. Rajaravivarma, "Virtual Local Area Network Technology and Applications," *Proc. 29th Southeastern Symp. Sys. Theory*, Mar. 1997, pp. 49–52.
[5] S. Halabi, *Metro Ethernet*, Cisco Press, Oct. 2003.
[6] P. Wang, C. Chan, and P. Lin, "MAC Address Translation for Enabling Scalable Virtual Private LAN Services," *Proc. 21st Int'l. Conf. Adv. Info. Net. and Apps. Wksp.*, vol. 1, May 2007, pp. 870–75.
[7] A. Elangovan, "Efficient Multicasting and Broadcasting in Layer 2 Provider Backbone Networks," *IEEE Commun. Mag.*, vol. 43, no. 11, Nov. 2005, pp. 166–70.
[8] M. Lasserre and V. Kompella, "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling," IETF RFC 4762, Jan. 2007; http://tools.ietf.org/html/rfc4762
[9] P. Knight and C. Lewis, "Layer 2 and 3 Virtual Private Networks: Taxonomy, Technology, and Standardization Efforts," *IEEE Commun. Mag.*, vol. 42, no. 6, June 2004, pp. 124–31.
[10] L. Martini and A. Sajassi, "802.1ah Ethernet Pseudowire," IETF Internet draft, Nov. 2007, work in progress.
[11] A. Sajassi *et al.*, "VPLS Interoperability with Provider Backbone Bridges," IETF Internet draft, Nov. 2007, work in progress.
[12] S. Yasukawa, Signaling Requirements for Point-to-Multipoint Traffic-Engineered MPLS Label Switched Paths (LSPs), IETF RFC 4461, Apr. 2006; http://tools.ietf.org/html/rfc4461
[13] R. Perlman *et al.*, "Rbridges: Base Protocol Specification," IETF Internet draft, Nov. 2007, work in progress.
[14] E. Gray, "The Architecture of an RBridge Solution to TRILL," IETF Internet draft, Nov. 2007, work in progress.


## Biographies


SAMER SALAM (ssalam@cisco.com) holds an M.S. in computer engineering from the University of Southern California and a B.E. in computer and communications engineering from the American University of Beirut. He is a senior technical leader at Cisco Systems. His responsibilities include system/software architecture and product development for Metro Ethernet. He is one of the leads in the definition of Cisco's Carrier Ethernet services and OAM software architectures. Previously, he worked on the system architecture of the multiservice edge platforms and product development for broadband and dial solutions.

ALI SAJASSI (sajassi@cisco.com) received a B.S. in electrical engineering and a B.S. in computer science from the University of Wisconsin-Madison. He holds an M.S. in electrical engineering from Purdue University, and he completed his Ph.D. studies at George Washington University in the area of networking and communications. He is a senior system/solution architect at Cisco Systems, where he has been working on end-to-end solutions in the areas of VoIP, layer 2 virtual private networks (L2VPN), and metro Ethernet networks. He is an active participant in IETF, IEEE, and ITU standards groups, and he has made numerous contributions to these groups in the area of L2VPN. Prior to joining Cisco Systems, he worked at several companies in both the data communications and telecommunications areas. He was the director of system engineering at AvalCom where he was responsible for design and development of W-CDMA and IS136 base stations systems. He was the principal architect at Sentient Networks where he designed next-generation multiservice access switches capable of supporting any protocol (ATM, FR, PPP, IP, etc.) at any of its high-speed interfaces at any time (ASAP). Prior to Sentient, he worked at Hughes Network Systems where as principal engineer he was responsible for design and development of base station systems for mobile satellite systems, and IS-136 and E-TDMA cellular systems.