



**SERVICE
PROVIDER**

Leveraging the Benefits of Provider Backbone Bridges

The Provider Backbone Bridges standard, IEEE 802.1ah, was developed to address the limitations of the Provider Bridges standard, IEEE 802.1ad, and to add additional capabilities sought by service providers.

When compared to a Provider Bridges (PB) network, a Provider Backbone Bridges (PBB) network deployment offers simplified operations, lower capital expenditures, and overall better scalability for the number of supported customers. PBB also provides advantages when used in conjunction with Virtual Private LAN Service (VPLS), since PBB reduces the overall MAC address learning requirements.

This paper provides an overview of PBB, outlines its advantages over PB, describes common PBB deployment scenarios, and examines its many benefits when deployed in combination with a core MultiProtocol Label Switching (MPLS) network supporting VPLS.

PROVIDER BACKBONE BRIDGES

To address the shortcomings of Provider Bridges, PBB adds a hierarchy view to Ethernet by encapsulating PB frames with a PBB header (which becomes the equivalent of a “Service Provider MAC header”) containing a Backbone Destination MAC Address (B-DA), Backbone Source MAC Address (B-SA), and two new tags (see Figure 1), which are described later in this document. What makes the B-DA and B-SA “backbone” addresses is the fact that these are MAC addresses of service provider’s PBB edge switches. An edge PBB switch encapsulates an ingress PB frame with a PBB header containing the destination MAC address of an appropriate egress edge PBB switch. The egress edge PBB switch removes the PBB header and forwards the frame to an attached PB network. Because PBB adds a PBB header containing new destination and source MAC addresses, it is also known as “MAC-in-MAC.”

By adding the PBB header, PBB isolates the service provider and customer address spaces. This means that Ethernet switches in the core of the service provider network will no longer learn or use customer MAC addresses to forward customer frames to their destinations. This improves the scaling of the service provider network in terms of the number of supported customers, since the number of supported customers is no longer directly tied to the size of the MAC address tables of the core Ethernet switches. In addition, the service provider network is now protected from customer network failures, since frame forwarding is now based on its own PBB header. Moreover, customers benefit from added security, since the customer's MAC addresses are no longer learned or used for frame forwarding decisions in the core of the service provider network.

As additional benefits to the service provider, PBB has the potential to simplify operations, for example, by separating the customer and service provider addressing spaces and to lower capital expenditures by reducing the cost of Ethernet switches used in the core of the network. This can be done since memory and processing power requirements are reduced by limiting MAC address learning to backbone MAC addresses.

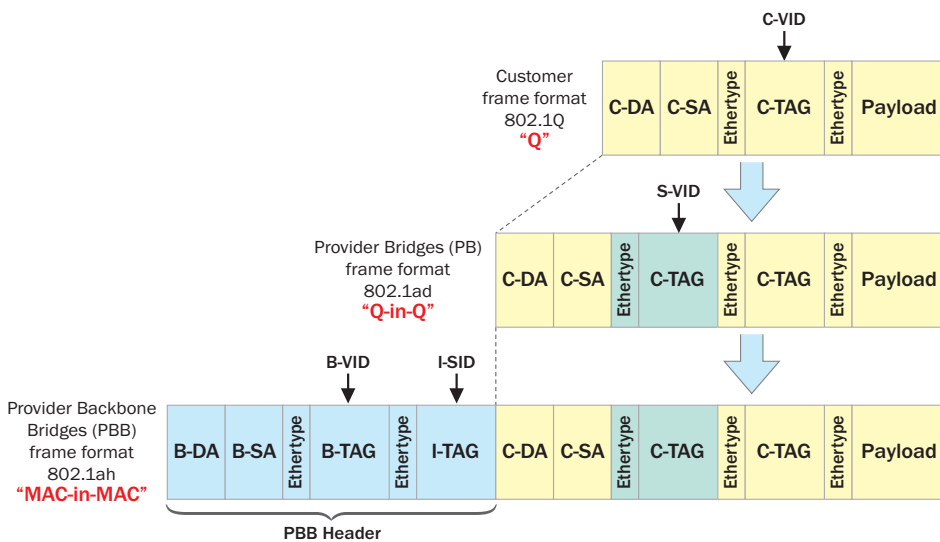


Figure 1. Customer, PB, and PBB frame formats.

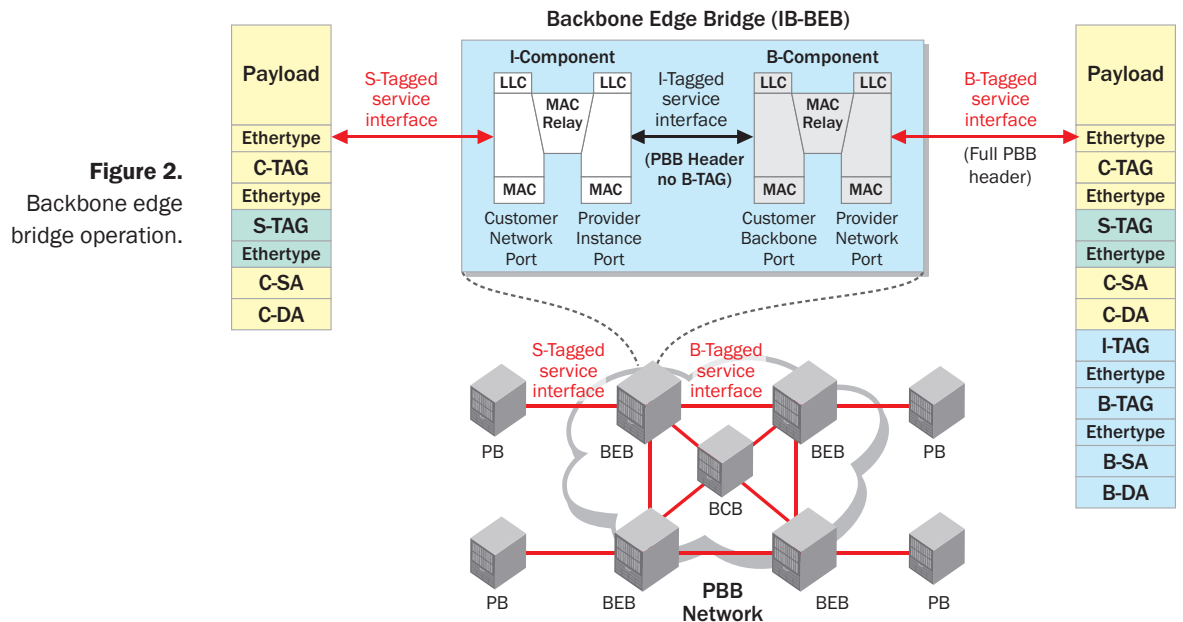
The new Backbone Service Instance Tag (I-TAG) contains a Backbone Service Instance Identifier (I-SID), which is 24 bits long. The I-SID field allows a service provider to identify up to 224, that is, over 16 million, service instances. In other words, over 16 million services or customers can be uniquely identified using the I-SID field. Therefore, the PBB I-TAG allows for highly scalable services by eliminating the 4,094 service instances limitation of PB.

The semantics and the structure of the Backbone VLAN Tag (B-TAG) are identical to that of the PB S-TAG. The B-TAG was designed this way so that core PBB switches do not need to be aware of PBB. In fact, standard PB switches can be used in the core of a PBB network. Only the switches at the edge of the service provider PBB network need to be aware of PBB.

A PBB network uses two types of bridges, as shown in Figure 2:

- Backbone Edge Bridges (BEB)
- Backbone Core Bridges (BCB)

As explained earlier, the functionality required from a BCB is the same as a standard IEEE 802.1ad PB bridge. A BEB is used at the boundary of a PBB network to add and remove the PBB header.



As defined in IEEE 802.1ah, a BEB has two main components: an I-Component and a B-Component. From left to right in Figure 2, the I-Component maps S-VIDs to I-SIDs and adds a PBB header without a B-TAG, while the B-Component maps I-SIDs to B-VIDs and adds a B-TAG. These actions are reverted in the opposite direction.

Also shown in Figure 2, a BEB containing an I-Component and a B-Component is called an IB-BEB. The B-Component of an IB-BEB forwards frames towards the core of a PBB network based on backbone MAC addresses (that is, it learns backbone MAC addresses), while the I-Component forwards frames towards the PB network based on customer MAC addresses (that is, it learns customer MAC addresses). The terms “I-BEB” and “B-BEB” refer to optional BEBs that support a single component type, either an I-Component or a B-Component, respectively. I-BEBs and B-BEBs expose an I-Tagged service interface, which carries frames with a PBB header, but without a B-TAG.

As with PB, PBB uses MAC address learning to dynamically build a forwarding database. Since an IB-BEB forwards frames to PB and PBB networks, it has to learn customer and backbone MAC addresses. However, since an IB-BEB is at the edge of the service provider network, it only learns the customer MAC addresses of the local traffic.

Resilient PBB Deployments

Traditionally, PBB relies on the use of a Spanning Tree Protocol (STP), such as the Rapid STP (RSTP) or the Multiple STP (MSTP), to dynamically determine the active topology of the PBB network and to avoid forwarding loops. However, STPs are known to take seconds to reconverge and re-establish network service around failed links or switches in large network deployments. This represents a serious limitation for service providers wanting to deploy resilient PBB networks.

To circumvent Spanning Tree Protocol limitations, an alternative topology discovery technology called Shortest Path Bridging (SPB) is being developed by the IEEE for PB and PBB networks as project 802.1aq. SPB uses a protocol to calculate multiple loop free active topologies. SPB seeks to improve the reconvergence time around failed links or switches and to eliminate the network utilization restrictions imposed by STPs, that is, the confinement of the traffic of each VLAN to a single spanning tree. SPB is being defined for deployment in arbitrary network topologies, for example, mesh and ring topologies, and is still in its early stages of development.

Alternatives to STPs and SPB that are optimized for network deployments in metro ring topologies exist today. Ring topologies are popular in most metro networks as they are ideally suited to cover large metro areas and provide inherent physical redundancy. Optimized ring topology discovery protocols are capable of providing sub-second fault detection, isolation, and traffic restoration around failed links or switches. Such resilient ring topology discovery protocols are available today for immediate deployment in service provider's PB and PBB networks.

The availability of Ethernet Operations Administration and Management (OAM) tools is another required attribute of a resilient PBB network. There are currently two complementing Ethernet OAM standards available for service providers to manage their PBB networks:

- IEEE 802.1ag Connectivity Fault Management
- ITU-T Y.1731 Ethernet OAM

These Ethernet OAM standards specify protocols, procedures, and managed objects to support the detection and isolation of connectivity faults, the propagation of faults, the suppression of alarms, and the management of network performance.

PBB Deployment Scenarios

Figure 3 shows an example deployment of an end-to-end PBB network using metro PBB networks (PBBNs) interconnected by a core PBB network. The switches at the edge of the metro PBBNs can aggregate both IEEE 802.1Q and IEEE 802.1ad (PB) traffic. Local traffic is switched at the metro PBBNs, while end-to-end traffic between metro PBBNs is switched across the core PBB network.

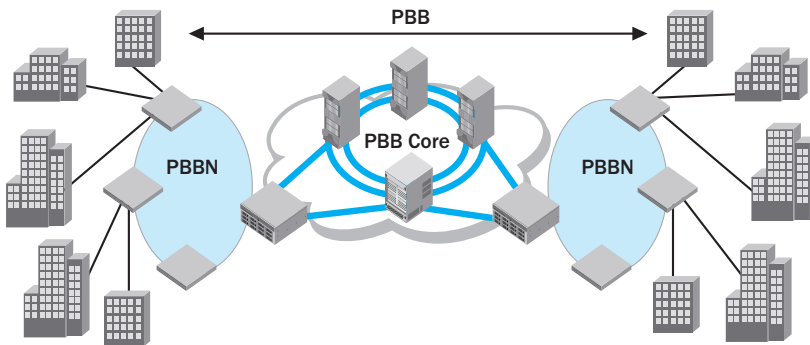


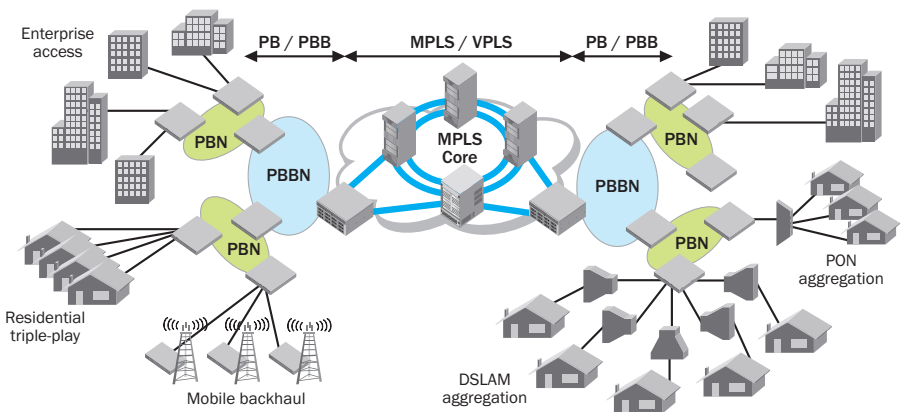
Figure 3.
Building an end-to-end PBB network.

Figure 4 shows an example of how a service provider can support end-to-end Carrier Ethernet services using edge PB Networks (PBNs) and metro PBB networks, while preserving investment in an existing core MPLS network. Services supported by this Carrier Ethernet network include:

- End-to-end Enterprise services that scale beyond the metros
- High-density aggregation of access devices such as DSLAMs, PON OLTs, or CMTS systems
- FTTx with stringent Service Level Agreements (SLA) for residential triple play services
- Mobile Backhaul

In Figure 4, edge PB networks provide access aggregation. Metro PBB networks aggregate traffic from edge PBNs, while a core MPLS network provides for wide area transport and connectivity of metro PBBNs. In addition to supporting the delivery of MPLS L3 VPN services, the core MPLS network enables edge-to-edge Enterprise services such as E-Line and E-LAN services using pseudowires and VPLS.

Figure 4.
True convergence
in a service provider
network.



Interworking of PBB with VPLS provides several advantages for scaling VPLS deployments. PBB effectively hides customer MAC address from the VPLS service by adding a PBB header to customer frames. This improves the scalability of VPLS, since the Provider Edge (PE) router responsible for initiating the VPLS instance no longer has to learn customer MAC addresses. This is similar to the scaling benefits that PBB core switches enjoy with PBB encapsulation. In summary, PBB interworking with VPLS allows for:

- **Simplification of operations.** This is done by significantly simplifying network management with the separation of customer and service provider addressing spaces.
- **Lower capital expenditures.** This approach reduces the cost of MPLS routers used in the core of the network, since memory and processing power requirements are reduced by limiting VPLS frame addressing to the service provider space only.
- **VPLS scalability.** By limiting frame addressing to the service provider space, PBB encapsulation allows VPLS to support more customers. In addition, as PBB provides a first level of customer aggregation, the number of PEs may be reduced, as compared to a deployment where PEs also provide customer aggregation. With fewer PEs, the total number of full mesh pseudowires required by VPLS is reduced, which translates into a lower total signaling load on each PE.

PBB Support with Brocade Products

Brocade's service provider vision is to deliver a high-performance, highly flexible, and efficient network architecture for service providers to deliver core, aggregation, and edge services with compelling scale and economics. This vision is enabled through a consistent and high-quality 10 Gigabit Ethernet (GbE) product portfolio supporting IPv4/v6 routing, MPLS, and PB/PBB switching, and comprised of both high-density modular systems as well as compact 1U (rack unit) systems for deployment flexibility.

The Brocade MEF-certified Carrier Ethernet product portfolio includes the modular NetIron XMR Series of core routers, the NetIron MLX Series of aggregation/metro routers, and the NetIron CES 2000 Series of fixed form factor edge switches. The NetIron XMR and MLX families feature advanced IPv4/v6/MPLS/L2 capabilities with support for 10/100/1000 Megabits per second (Mbps) Ethernet, 10 GbE, and OC-12/-48/-192 interfaces. The NetIron MLX Series supports up to 256 x 10 GbE wire-speed ports in a single chassis while the NetIron XMR Series supports up to 128 x 10 GbE wire-speed ports. The NetIron CES 2000 Series is the industry's first 1U PBB-capable

10 GbE switch.

The Brocade Carrier Ethernet product portfolio, which has achieved MEF 9 and MEF 14 certification, enables service providers to meet or exceed the stringent requirements of Carrier Ethernet attributes defined by the Metro Ethernet Forum (MEF):

- **Standardized Services.** The Brocade Carrier Ethernet product portfolio provides rich Carrier Ethernet services using Layer 2 VLANs (802.1Q), PB (802.1ad), and PBB (802.1ah) and supports E-LINE (Ethernet Private Line and Ethernet Virtual Private Line), E-LAN (Ethernet LAN), and E-TREE (Ethernet Tree) services.
- **Scalability.** The Brocade product portfolio includes PBB functionality with 10 GbE capability, best in class MAC scalability, and unparalleled scalability in Carrier Ethernet services when combined with VPLS on the NetIron XMR/MLX.
- **Reliability.** The Brocade product portfolio supports standard Layer 2 protocols, Layer 2 tunneling of customer protocols, and the integration of Layer 2 resiliency protocols with VPLS (when using NetIron XMR/MLX). The Brocade Metro Ring Protocol (MRP and MRP-II) and Virtual Switch Redundancy Protocol (VSRP) are available for PBB and PB networks. Brocade MRP is specifically designed for metro rings as an alternative to Spanning Tree Protocols. Ring topologies are popular in most metro networks as they are ideally suited to covering large metro areas and providing inherent physical redundancy. MRP offers sub-second fault detection, isolation, and failover in overlapping rings. VSRP is a sister technology to MRP, but is specifically designed for mesh topologies. Combining MRP and VSRP provides an extremely robust network design. MRP and VSRP complement the Brocade Rapid STP and SuperSpan feature set.
- **Quality of Service (QoS).** The Brocade product portfolio supports high-touch QoS management with advanced QoS controls, ingress and egress bandwidth profiles with 2-rate 3-color policers, egress traffic shaping, and advanced priority remarking capabilities.
- **Service Management.** The Brocade product portfolio supports IEEE 802.1ag Connectivity Fault Management, including Continuity Check Messages (CCM), linktrace, and loopback messages, MAC ping and MAC traceroute, and LACP for ensuring bi-directionality. In addition, the Brocade product portfolio supports the management of end-to-end services using MEF 17 Service OAM Framework.

SUMMARY

This paper provides an overview of PBB, described its advantages over PB, examines common PBB deployment scenarios, and explains how the interworking of PBB with MPLS/VPLS can provide scaling benefits. Brocade's rich Carrier Ethernet product portfolio supports a high-performance, highly flexible, and efficient network architecture for service providers to deliver core, aggregation, and edge services with compelling scale and economics. Brocade offers service providers a wide range of technologies spanning the breadth of IPv4/v6 routing, MPLS, and PB/PBB switching, and is comprised of both high-density modular systems as well as compact 1U systems for deployment flexibility.

To learn more, contact your Brocade sales representative or visit www.brocade.com.

ACRONYMS

C-DA—Customer Destination MAC address

C-SA—Customer Source MAC address

C-Tag—Customer VLAN Tag

C-VID—Customer VID (12 bits)

B-DA—Backbone Destination MAC address

B-SA—Backbone Source MAC address

B-TAG—Backbone VLAN Tag

B-VID—Backbone VID (12 bits)

I-SID—Backbone Service Instance Identifier (24 bits)

I-TAG—Backbone Service Instance Tag

OAM—Operations Administration and Management

S-TAG—Service Provider VLAN Tag

S-VID—Service Provider VID (12 bits)

VID—VLAN ID

VLAN—Virtual Local Area Network

Corporate Headquarters

San Jose, CA USA
T: +1-408-333-8000
info@brocade.com

European Headquarters

Geneva, Switzerland
T: +41-22-799-56-40
emea-info@brocade.com

Asia Pacific Headquarters

Singapore
T: +65-6538-4700
apac-info@brocade.com

© 2010 Brocade Communications Systems, Inc. All Rights Reserved. 05/10 GA-WP-1380-01

Brocade, the B-wing symbol, BigIron, DCX, Fabric OS, FastIron, IronView, NetIron, SAN Health, ServerIron, and Turbolron are registered trademarks, and Brocade Assurance, DCFM, Extraordinary Networks, and Brocade NET Health are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned are or may be trademarks or service marks of their respective owners.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

**BROCADE**