

Internet Engineering Task Force (IETF)
Request for Comments: 7080
Category: Informational
ISSN: 2070-1721

A. Sajassi
S. Salam
Cisco
N. Bitar
Verizon
F. Balus
Nuage Networks
December 2013

Virtual Private LAN Service (VPLS) Interoperability
with Provider Backbone Bridges

Abstract

The scalability of Hierarchical Virtual Private LAN Service (H-VPLS) with Ethernet access networks (RFC 4762) can be improved by incorporating Provider Backbone Bridge functionality in the VPLS access. Provider Backbone Bridging has been standardized as IEEE 802.1ah-2008. It aims to improve the scalability of Media Access Control (MAC) addresses and service instances in Provider Ethernet networks. This document describes different interoperability scenarios where Provider Backbone Bridge functionality is used in H-VPLS with Ethernet or MPLS access network to attain better scalability in terms of number of customer MAC addresses and number of service instances. The document also describes the scenarios and the mechanisms for incorporating Provider Backbone Bridge functionality within H-VPLS with existing Ethernet access and interoperability among them. Furthermore, the document discusses the migration mechanisms and scenarios by which Provider Backbone Bridge functionality can be incorporated into H-VPLS with existing MPLS access.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7080>.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Terminology	3
3. Applicability	5
4. H-VPLS with Homogeneous PBBN Access	6
4.1. Service Interfaces and Interworking Options	8
4.2. H-VPLS with PBBN Access: Type I Service Interface	10
4.3. H-VPLS with PBBN Access: Type II Service Interface	11
5. H-VPLS with Mixed PBBN Access and PBN Access	14
5.1. H-VPLS with Mixed PBBN and PBN Access: Modified PBN PE	15
5.2. H-VPLS with Mixed PBBN and PBN Access: Regular PBN PE	16
6. H-VPLS with MPLS Access	17
6.1. H-VPLS with MPLS Access: PBB U-PE	17
6.2. H-VPLS with MPLS Access: PBB N-PE	19
7. H-VPLS with MPLS Access: PBB Migration Scenarios	21
7.1. 802.1ad Service Frames over VPLS Core	21
7.2. PBB Service Frames over VPLS Core	22
7.3. Mixed 802.1ad and PBB over VPLS Core	23
8. Acknowledgments	24
9. Security Considerations	24
10. References	24
10.1. Normative References	24
10.2. Informative References	25

1. Introduction

The scalability of Hierarchical Virtual Private LAN Service (H-VPLS) with Ethernet access networks [RFC4762] can be improved by incorporating Provider Backbone Bridge (PBB) functionality in the VPLS access. Provider Backbone Bridging has been standardized as IEEE 802.1ah-2008 [802.1ah], which is an amendment to IEEE 802.1Q to improve the scalability of Media Access Control (MAC) addresses and service instances in Provider Ethernet networks. This document describes interoperability scenarios where IEEE 802.1ah functionality is used in H-VPLS with Ethernet or MPLS access network to attain better scalability in terms of the number of customer MAC addresses and the number of services.

This document also covers the interoperability scenarios for deploying H-VPLS with Provider Backbone Bridging Ethernet access when other types of access networks are deployed, including existing 802.1ad Ethernet and MPLS access in either single or multiple service domains. Furthermore, the document explores the scenarios by which an operator can gradually migrate an existing H-VPLS network to Provider Backbone Bridging over VPLS.

Section 2 gives a quick terminology reference and Section 3 highlights the applicability of Provider Backbone Bridging interoperation with VPLS. Section 4 describes H-VPLS with homogeneous Provider Backbone Bridge Access Network. Section 5 discusses H-VPLS with mixed 802.1ah/802.1ad access. Section 6 focuses on Provider Backbone Bridging in H-VPLS with MPLS Access Network including PBB function on U-PE and on N-PE variants. Finally, Section 7 describes gradual migration scenarios from existing H-VPLS to Provider Backbone Bridging over H-VPLS.

2. Terminology

802.1ad: IEEE specification for "QinQ" encapsulation and bridging of Ethernet frames.

802.1ah: IEEE specification for "MAC tunneling" encapsulation and bridging of frames across a Provider Backbone Bridged Network (PBBN).

B-BEB: A backbone edge bridge positioned at the edge of a PBBN. It contains a B-component that supports bridging in the provider backbone based on B-MAC and B-TAG information.

B-MAC: The backbone source or destination MAC address fields defined in the 802.1ah provider MAC encapsulation header.

- BCB: A backbone core bridge running in the core of a provider backbone bridged network. It bridges frames based on B-TAG information just as an 802.1ad provider bridge will bridge frames based on a Service VLAN (S-VLAN) identifier.
- B-Component: The backbone component of a Provider Backbone edge bridge as defined in [802.1ah].
- BEB: A backbone edge bridge positioned at the edge of a provider backbone bridged network. It can contain an I-component, B-component, or both.
- B-MACs: Backbone MAC addresses -- outer MAC addresses of a PBB-encapsulated frame.
- B-TAG: A field defined in the 802.1ah provider MAC encapsulation header that conveys the backbone VLAN identifier information. The format of the B-TAG field is the same as that of an 802.1ad S-TAG field.
- B-Tagged Service Interface: This is the interface between a BEB and BCB in a PBB network. Frames passed through this interface contain a B-TAG field.
- B-VID: This is the specific VLAN identifier carried inside a B-TAG
- C-MACs: Customer MAC addresses are the inner MAC addresses of a PBB-encapsulated frame.
- H-VPLS: Hierarchical Virtual Private LAN Service.
- I-component: A bridging component contained in a backbone edge bridge that bridges in the customer space (customer MAC addresses, S-VLAN).
- IB-BEB: A backbone edge bridge positioned at the edge of a provider backbone bridged network. It contains an I-component for bridging in the customer space (customer MAC addresses, S-VLAN IDs) and a B-component for bridging the provider's backbone space (B-MAC, B-TAG).
- I-BEB: A backbone edge bridge positioned at the edge of a provider backbone bridged network. It contains an I-component for bridging in the customer space (customer MAC addresses, S-VLAN IDs).
- I-SID: The 24-bit service instance field carried inside the I-TAG. I-SID defines the service instance that the frame should be "mapped to".

I-TAG: A field defined in the 802.1ah provider MAC encapsulation header that conveys the service instance information (I-SID) associated with the frame.

I-Tagged Service Interface: This is the interface defined between the I-component and B-component inside an IB-BEB or between two B-BEBs. Frames passed through this interface contain an I-TAG field.

N-PE: Network-facing Provider Edge

PBB: Provider Backbone Bridge

PBBN: Provider Backbone Bridged Network

PBN: Provider Bridged Network. A network that employs 802.1ad (QinQ) technology.

S-TAG: A field defined in the 802.1ad QinQ encapsulation header that conveys the Service VLAN (S-VLAN) identifier information.

S-Tagged Service Interface: This the interface defined between the customer (CE) and the I-BEB or IB-BEB components. Frames passed through this interface contain an S-TAG field.

S-VLAN: The specific Service VLAN identifier carried inside an S-TAG

U-PE: User-facing Provider Edge

VPLS: Virtual Private LAN Service

3. Applicability

[RFC4762] describes a two-tier hierarchical solution for VPLS for the purpose of improved pseudowire (PW) scalability. This improvement is achieved by reducing the number of PE devices connected in a full-mesh topology through connecting CE devices via the lower-tier access network, which in turn is connected to the top-tier core network.

[RFC4762] describes two types of H-VPLS network topologies -- one with an MPLS access network and another with an IEEE 802.1ad (QinQ) Ethernet access network. In both types of H-VPLS, the learning and forwarding of MAC addresses is based on customer MAC addresses (C-MACs), which poses scalability issues as the number of VPLS instances (and thus C-MACs) increases. Furthermore, since a set of PWs is maintained on a per customer service instance basis, the number of PWs required at N-PE devices is proportional to the number of customer service instances multiplied by the number of N-PE devices in the full-mesh set. This can result in scalability issues

(in terms of PW manageability and troubleshooting) as the number of customer service instances grows.

In addition to the above, H-VPLS with an 802.1ad Ethernet access network has another scalability issue in terms of the maximum number of service instances that can be supported in the access network as described in [RFC4762]. Since the number of provider VLANs (S-VLANs) is limited to 4094 and each S-VLAN represents a service instance in an 802.1ad network, then the maximum number of service instances that can be supported is 4094. These issues are highlighted in [RFC6246].

This document describes how IEEE 802.1ah (aka Provider Backbone Bridges) can be integrated with H-VPLS to address these scalability issues. In the case of H-VPLS with 802.1ah Ethernet access, the solution results in better scalability in terms of both number of service instances and number of C-MACs in the Ethernet access network and the VPLS core network, as well as number of PWs in VPLS core network. And in the case of H-VPLS with MPLS access, Provider Backbone Bridging functionality can be used at the U-PE or N-PE, which results in reduction of customer MAC addresses and the number of PWs in the VPLS core network.

The interoperability scenarios depicted in this document fall into the following two categories:

- Scenarios in which Provider Backbone Bridging seamlessly works with current VPLS implementations (e.g., Section 4.2).
- Scenarios in which VPLS-PE implementations need to be upgraded in order to work with Provider Backbone Bridging (e.g., Sections 4.3 and 5.1).

4. H-VPLS with Homogeneous PBBN Access

PBBN access offers MAC-address-table scalability for H-VPLS PE nodes. This is due to the MAC tunneling encapsulation scheme of PBB, which only exposes the provider's own MAC addresses to PE nodes (B-MACs of Provider's PBB-capable devices in the access network), as opposed to customers' MAC addresses in conventional H-VPLS with MPLS or 802.1ad access.

PBBN access also offers service-instance scalability when compared to H-VPLS with 802.1Q/802.1ad access networks. This is due to the new 24-bit service identifier (I-SID) used in PBB encapsulation, which allows up to 16M services per PBB access network, compared to 4094 services per 802.1Q/802.1ad access network.

use the same B-VID value, in practice, they often use the same value because once the I-SID grouping is made identical in two PBBNs. It is very easy to make the values of the corresponding B-VIDs identical also.

Consequently, three different kinds of "Service Domains" are defined in the following manner:

- Tightly Coupled Service Domain - Different PBBNs' access belonging to the same I-SID Domain and B-VLAN Domain. However, the network control protocols (e.g., xSTP) run independently in each PBBN access.
- Loosely Coupled Service Domain - Different PBBNs' access belonging to the same I-SID Domain. However, each PBBN access maintains its own independent B-VLAN Domain. Again, the network control protocols (e.g., xSTP) run independently in each PBBN access.
- Different Service Domain - In this case, each PBBN access maintains its own independent I-SID Domain and B-VLAN Domain, with independent network control protocols (e.g., xSTP) in each PBBN access.

In general, correct service connectivity spanning networks in a Tightly Coupled Service Domain can be achieved via B-VID mapping between the networks (often even without B-VID translation). However, correct service connectivity spanning networks in a Loosely Coupled Service Domain requires I-SID to B-VID remapping (i.e., unbundling and rebundling of I-SIDs into B-VIDs). Furthermore, service connectivity spanning networks in Different Service Domains requires both I-SID translation and I-SID to B-VID remapping.

4.1. Service Interfaces and Interworking Options

Customer devices will interface with PBBN edge bridges using existing Ethernet interfaces including IEEE 802.1Q and IEEE 802.1ad. At the PBBN edge, C-MAC frames are encapsulated in a PBB header that includes service provider source and destination MAC addresses (B-MACs) and are bridged up to the VPLS-PE. The PBB-encapsulated C-MAC frame is then injected into the VPLS backbone network, delivered to the remote VPLS-PE node(s), and switched onto the remote PBBN access. From there, the PBBN bridges the encapsulated frame to a PBBN edge bridge where the PBB header is removed and the customer frame is sent to the customer domain.

Interoperating between PBBN devices and VPLS-PE nodes can leverage the BEB functions already defined in [802.1ah]. When I-SID visibility is required at the VPLS-PE nodes, a new service interface based on I-SID tag will need to be defined.

Moreover, by mapping a bridge domain (e.g., B-VLAN) to a VPLS instance, and bundling multiple end-customer service instances (represented by I-SID) over the same bridge domain, service providers will be able to significantly reduce the number of full-mesh PWs required in the core. In this case, I-SID visibility is not required on the VPLS-PE and the I-SID will serve as the means of multiplexing/de-multiplexing individual service instances in the PBBN over a bundle (e.g., B-VLAN).

When I-SID visibility is expected across the service interface at the VPLS-PE, the VPLS-PE can be considered to offer service-level interworking between PBBN access and the IP/MPLS core. Similarly, when the PE is not expected to have visibility of the I-SID at the service interface, the VPLS-PE can be considered to offer network-level interworking between PBBN access and the MPLS core.

A VPLS-PE is always part of the IP/MPLS core, and it may optionally participate in the control protocols (e.g., xSTP) of the access network. When connecting to a PBBN access, the VPLS-PE needs to support one of the following two types of service interfaces:

- Type I: B-Tagged Service Interface with B-VID as Service Delimiter

The PE connects to a Backbone Core Bridge (BCB) in the PBBN access. The handoff between the BCB and the PE is B-Tagged PBB-encapsulated frames. The PE is transparent to PBB encapsulations and treats these frames as 802.1ad frames since the B-VID EtherType is the same as the S-VID EtherType. The PE does not need to support PBB functionality. This corresponds to conventional VPLS-PEs' tagged service interface. When using Type I service interface, the PE needs to support either raw mode or tagged mode Ethernet PW. Type I service interface is described in detail in Section 4.2.

- Type II: I-Tagged Service Interface with I-SID as Service Delimiter

The PE connects to a B-BEB (backbone edge bridge with B-component) in the PBBN access. The PE itself also supports the B-BEB functionality of [802.1ah]. The handoff between the B-BEB in the PBBN access and the PE is an I-Tagged PBB-encapsulated frame. With Type II service

interface, the PE supports the existing raw mode and tagged mode PW types. Type II service interface is described in detail in Section 4.3.

4.2. H-VPLS with PBBN Access: Type I Service Interface

This is a B-Tagged service interface with B-VID as service delimiter on the VPLS-PE. It does not require any new functionality on the VPLS-PE. As shown in Figure 2, the PE is always part of the IP/MPLS core. The PE may also be part of the PBBN access (e.g., VPLS-PE on right side of Figure 2) by participating in network control protocols (e.g., xSTP) of the PBBN access.




Figure 2: H-VPLS with PBBN Access and Type I Service Interface

Type I service interface is applicable to networks with Tightly Coupled Service Domains, where both I-SID Domains and B-VLAN Domains are the same across all PBBN access networks.

The BCB and the VPLS-PE will exchange PBB-encapsulated frames that include source and destination B-MACs, a B-VID, and an I-SID. The service delimiter, from the perspective of the VPLS-PE, is the B-VID; in fact, this interface operates exactly as a current 802.1Q/ad interface into a VPLS-PE does today. With Type I service interface, the VPLS-PE can be considered to provide network-level interworking between PBBN and MPLS domains, since VPLS-PE does not have visibility of I-SIDs.

The main advantage of this service interface, when compared to other types, is that it allows the service provider to save on the number of full-mesh PWs required in the core. This is primarily because

multiple service instances (I-SIDs) are bundled over a single full-mesh PW corresponding to a bridge domain (e.g., B-VID), instead of requiring a dedicated full-mesh PW per service instance. Another advantage is the MAC address scalability in the core since the core is not exposed to C-MACs.

The disadvantage of this interface is the comparably excessive replication required in the core: since a group of service instances share the same full-mesh of PWs, an unknown unicast, multicast, or broadcast on a single service instance will result in a flood over the core. This, however, can be mitigated via the use of flood containment per I-SID (B-MAC multicast pruning).

Three different modes of operation are supported by Type I service interface:

- Port Mode: All traffic over an interface in this mode is mapped to a single VPLS instance. Existing PW signaling and Ethernet raw mode (0x0005) PW type, defined in [RFC4447] and [RFC4448], are supported.
- VLAN Mode: all traffic associated with a particular VLAN identified by the B-VID is mapped to a single VPLS instance. Existing PW signaling and Ethernet raw mode (0x0005) PW type, defined in [RFC4447] and [RFC4448], are supported.
- VLAN Bundling Mode: all traffic associated with a group or range of VLANs or B-VIDs is mapped to a single VPLS instance. Existing PW signaling and Ethernet raw mode (0x0005) PW type, defined in [RFC4447] and [RFC4448], are supported.

For the VLAN mode, it is also possible to use Ethernet tagged mode (0x0004) PW, as defined in [RFC4447] and [RFC4448], for interoperability with equipment that does not support raw mode. The use of raw mode is recommended to be the default though.

4.3. H-VPLS with PBBN Access: Type II Service Interface

This is an I-Tagged service interface with I-SID as service delimiter on the VPLS-PE. It requires the VPLS-PE to include the B-component of PBB BEB for I-SID processing in addition to the capability to map an I-SID Bundle to a VPLS instance. As shown in Figure 3, the PE is always part of the IP/MPLS core and connects to one or more B-BEBs in the PBBN access.




Figure 3: H-VPLS with PBBN Access and Type II Service Interface

Type II service interface is applicable to Loosely Coupled Service Domains and Different Service Domains. B-VLAN Domains can be independent and the B-VID is always locally significant in each PBBN access: it does not need to be transported over the IP/MPLS core. Given the above, it should be apparent that Type II service interface is applicable to Tightly Coupled Service Domains as well.

By definition, the B-BEB connecting to the VPLS-PE will remove any B-VLAN tags for frames exiting the PBBN access. The B-BEB and VPLS-PE will exchange PBB-encapsulated frames that include source and destination B-MACs and an I-SID. The service delimiter, from the perspective of the VPLS-PE, is the I-SID. Since the PE has visibility of I-SIDs, the PE provides service-level interworking between PBBN access and IP/MPLS core.

Type II service interface may operate in I-SID Bundling Mode: all traffic associated with a group or range of I-SIDs is mapped to a single VPLS instance. The PE maintains a mapping of I-SIDs to a PE local bridge domain (e.g., B-VID). The VPLS instance is then associated with this bridge domain. With Loosely Coupled service Domains, no I-SID translation needs to be performed. Type II Service interface also supports Different Service Domains in this mode, since the B-BEB link in the PE connecting to the local PBBN can perform the translation of PBBN-specific I-SID to a local I-SID within the IP/MPLS core, which may then be translated to the other PBBN-specific I-SID on the egress PE. Such translation can also occur in the B-BEB of PBBN access. Existing PW signaling and Ethernet raw mode (0x0005), defined in [RFC4447] and [RFC4448], is supported. It is

also possible to use a tagged mode (0x0004) PW for purpose of interoperability with equipment that does not support raw mode.

Type II service interface provides operators with the flexibility to trade off PW state for multicast flooding containment, since a full-mesh of PWs can be set up:

- a. per I-SID,
- b. per group of I-SIDs, or
- c. for all I-SIDs.

For (a) and (b), the advantage that Type II service interface has compared to Type I is that it can reduce replication in the core without the need for a mechanism that provides flood containment per-I-SID (B-MAC multicast pruning). This is mainly due to the increased segregation of service instances over disjoint full meshes of PWs. For (c), both Type II and Type I service interfaces are at par with regard to flood containment.

For (a) and (b), the disadvantage of this service interface, compared to Type I, is that it may require a larger number of full-mesh PWs in the core. For (c), both Type II and Type I service interfaces are at par with regard to PW state. However, for all three scenarios, the number of full-mesh PWs can still be fewer than the number required by H-VPLS without PBBN access, since an I-SID can multiplex many S-VLANs.

It is expected that this interface type will be used for customers with significant multicast traffic (but without multicast pruning capability in the VPLS-PE) so that a separate VPLS instance is set up per group of customers with similar geographic locality (per I-SID group).

Note: Port mode is not called out in Type II service interface since it requires the mapping of I-SIDs to be identical on different I-Tagged interfaces across VPLS network. If this is indeed the case, Port mode defined in Type I service interface (Section 4.2) can be used.

5.1. H-VPLS with Mixed PBBN and PBN Access: Modified PBN PE

As shown in Figure 5, the operation of VPLS-PE2 (connecting to the PBBN access on the right) is no different from what was discussed in Section 4. Type II service interface, as discussed in the above section, is applicable. It is the behavior of VPLS-PE1 (connecting to the PBN access on the left) that is the focus of this section.




Figure 5: H-VPLS with Mixed PBN and PBBN Access: Modified PBN PE

Some assumptions made for this topology include:

- CE is directly connected to PBBN via S-Tagged or port-based interface.
- I-SID in PBBN access represents the same customer as S-VID in PBN access.
- At S-Tagged service interface of PE with IB-BEB functionality (e.g., PE1 in Figure 5), the only viable service is 1:1 mapping of S-VID to I-SID. However, towards the core network side, the same PE can support I-SID bundling into a VPLS instance.
- PE1 participates in the local I-SID Domain of the IP/MPLS core so the model accommodates for the rest of the PBB network any of the three domain types described in Section 4 -- Tightly Coupled, Loosely Coupled, and Different Service Domains.

- For ease of provisioning in these disparate access networks, it is recommended to use the same I-SID Domain among the PBBN access networks and the PEs with IB-BEB functionality (those connecting to PBN).

This topology operates in I-SID Bundling Mode: at a PE connecting to PBN access, each S-VID is mapped to an I-SID and subsequently a group of I-SIDs is mapped to a VPLS instance. Similarly, at a PE connecting to PBBN access, each group of I-SIDs is mapped to a VPLS instance. Similar to Type II service interface, no I-SID translation is performed for the I-SID bundling case. Existing PW signaling and Ethernet raw mode (0x0005) PW type, defined in [RFC4447] and [RFC4448], are supported. It is possible to use tagged mode (0x0004) PW for backward compatibility as well.

5.2. H-VPLS with Mixed PBBN and PBN Access: Regular PBN PE

As shown in Figure 6, the operation of VPLS-PE1 (connecting to the PBN access on the left) is no different from existing VPLS-PEs. It is the behavior of VPLS-PE2 (connecting to the PBBN access on the right) that is the focus of this section.




Figure 6: H-VPLS with Mixed PBN and PBBN Access: Regular PBN PE

Some assumptions made for this topology include:

- The CE is directly connected to the PBBN access via an S-Tagged or port-based Interface.

- The I-SID in the PBBN access represents the same customer as the S-VID in the PBN access.
- There is 1:1 mapping between the I-SID and the VPLS instance.
- At the S-Tagged service interface of the PE connecting to PBN (e.g., PE1 in Figure 6), the PE only provides 1:1 mapping of S-VID to the VPLS instance. S-VID bundling is not a viable option since it does not correspond to anything in the PBBN access.
- The PE connecting to the PBBN access (e.g., PE2 in Figure 6), supports IB-BEB functionality and the I-component is connected to the VPLS Forwarder (i.e., the I-component faces the IP/MPLS core whereas the B-component faces the PBBN access network). One or more I-SIDs can be grouped into a B-VID in the PBBN access.
- Since C-VID grouping in different PBBN access networks must be consistent, it is assumed that same I-SID Domain is used across these PBBN access networks.

Unlike the previous case, I-SID bundling mode is not supported in this case. This is primarily because the VPLS core operates in the same manner as today. The PE with IB-BEB functionality connecting to PBBN access performs the mapping of each VPLS instance to an I-SID and one or more of these I-SIDs may be mapped onto a B-VID within the PBBN access network.

6. H-VPLS with MPLS Access

In this section, the case of H-VPLS with MPLS access network is discussed. The integration of PBB functionality into VPLS-PE for such access networks is described to improve the scalability of the network in terms of the number of MAC addresses and service instances that are supported.

For this topology, it is possible to embed PBB functionality in either the U-PE or the N-PE. Both of these cases are described in the following subsections.

6.1. H-VPLS with MPLS Access: PBB U-PE

As stated earlier, the objective for incorporating PBB function at the U-PE is to improve the scalability of H-VPLS networks in terms of the number of MAC addresses and service instances that are supported.

In current H-VPLS, the N-PE must learn customer MAC addresses (C-MACs) of all VPLS instances in which it participates. This can easily add up to hundreds of thousands or even millions of C-MACs at

the N-PE. When the U-PE performs PBB encapsulation, the N-PE only needs to learn the MAC addresses of the U-PEs, which is a significant reduction. Furthermore, when PBB encapsulation is used, many I-SIDs are multiplexed within a single bridge domain (e.g., B-VLAN). If the VPLS instance is set up per B-VLAN, then one can also achieve a significant reduction in the number of full-mesh PWs. It should be noted that this reduction in full-mesh PWs comes at the cost of potentially increased replication over the full-mesh of PWs: customer multicast and/or broadcast frames are effectively broadcasted within the B-VLAN. This may result in additional frame replication because the full-mesh PWs corresponding to a B-VLAN are most likely bigger than the full-mesh PWs corresponding to a single I-SID. However, flood containment per I-SID (B-MAC multicast pruning) can be used to remedy this drawback and have multicast traffic replicated efficiently for each customer (i.e., for each I-SID).

Figure 7 below illustrates the scenario for H-VPLS with MPLS access. As illustrated, customer networks or hosts (CE) connect into the U-PE nodes using standard Ethernet interfaces [802.1D-REV], [802.1Q], or [802.1ad]. The U-PE is connected upstream to one or more VPLS N-PE nodes by MPLS PWs (per VPLS instance). These, in turn, are connected via a full mesh of PWs (per VPLS instance) traversing the IP/MPLS core. The U-PE is outfitted with PBB Backbone Edge Bridge (BEB) functions where it can encapsulate/decapsulate customer MAC frames in provider B-MACs and perform I-SID translation if needed.




Figure 7: H-VPLS with MPLS Access Network and PBB U-PE

The U-PE still provides the same type of services toward its customers as before and they are:

- Port mode (either 802.1D, 802.1Q, or 802.1ad)
- VLAN mode (either 802.1Q or 802.1ad)
- VLAN-bundling mode (either 802.1Q or 802.1ad)

By incorporating a PBB function, the U-PE maps each of these services (for a given customer) onto a single I-SID based on the configuration at the U-PE. Many I-SIDs are multiplexed within a single bridge domain (e.g., B-VLAN). The U-PE can then map a bridge domain onto a VPLS instance and the encapsulated frames are sent over the PW associated with that VPLS instance. Furthermore, the entire Ethernet bridging operation over the VPLS network is performed as defined in [RFC4762]. In other words, MAC forwarding is based on the B-MAC address space and service delimiter is based on VLAN ID, which is B-VID in this case. There is no need to inspect or deal with I-SID values in the VPLS N-PEs.

In the case of PBB U-PEs in a single I-SID Domain, I-SID assignment is performed globally across all MPLS access networks and therefore there is no need for I-SID translation. This scenario supports I-SID bundling mode, and it is assumed that the mapping of the I-SIDs to the bridge domain (e.g., B-VLAN) is consistent across all the participating PE devices. In the case of the I-SID bundling mode, a bridge domain (e.g., B-VLAN) is mapped to a VPLS instance and an existing Ethernet raw mode (0x0005) or tagged mode (0x0004) PW type is used as defined in [RFC4447] and [RFC4448].

I-SID mode can be considered to be a degenerate case of I-SID bundling where a single bridge domain is used per I-SID. However, that results in an increased number of bridge domains and PWs in the PEs. PBB flood containment (B-MAC multicast pruning) per I-SID can be used in conjunction with I-SID bundling mode to limit the scope of flooding per I-SID thus removing the need for I-SID mode.

6.2. H-VPLS with MPLS Access: PBB N-PE

In this case, the PBB function is incorporated at the N-PE to improve the scalability of H-VPLS networks in terms of the numbers of MAC addresses and service instances that are supported.

Customer networks or hosts (CE) connect into the U-PE nodes using standard Ethernet interfaces [802.1D-REV], [802.1Q], or [802.1ad]. The U-PE is connected upstream to one or more VPLS N-PE nodes by MPLS PWs (per customer). These, in turn, are connected via a full mesh of PWs (per customer or group of customers) traversing the IP/MPLS core.

The U-PE still provides the same type of services toward its customers as before and they are:

- Port mode (either 802.1D, 802.1Q, or 802.1ad)
- VLAN mode (either 802.1Q or 802.1ad)
- VLAN-bundling mode (either 802.1Q or 802.1ad)

The spoke PW from the U-PE to the N-PE is not service multiplexed because there is no PBB functionality on the U-PE, i.e., one service per PW.




Figure 8: H-VPLS with MPLS Access Network and PBB N-PE

By incorporating a PBB function, the N-PE maps each of these services (for a given customer) onto a single I-SID based on the configuration at the N-PE. Many I-SIDs can be multiplexed within a single bridge domain (e.g., B-VLAN). The N-PE can, then, either map a single I-SID into a VPLS instance or map a bridge domain (e.g., B-VLAN) onto a VPLS instance, according to its configuration. Next, the encapsulated frames are sent over the set of PWs associated with that VPLS instance.

In the case of PBB N-PEs in a single I-SID Domain, I-SID assignment is performed globally across all MPLS access networks and therefore there is no need for I-SID translation. This scenario supports I-SID bundling mode, and it is assumed that the mapping of the I-SIDs to the bridge domain (e.g., B-VLAN) is consistent across all the participating PE devices. In the case of the I-SID bundling mode, a bridge domain (e.g., B-VLAN) is mapped to a VPLS instance and an existing Ethernet raw mode (0x0005) or tagged mode (0x0004) PW type as defined in [RFC4447] and [RFC4448], can be used.

I-SID mode can be considered to be a degenerate case of I-SID bundling where a single bridge domain is used per I-SID. However, that results in an increased number of bridge domains and PWs in the PE. PBB flood containment (B-MAC multicast pruning) per I-SID can be used in conjunction with I-SID bundling mode to limit the scope of flooding per I-SID thus removing the need for I-SID mode.

7. H-VPLS with MPLS Access: PBB Migration Scenarios

Operators and service providers that have deployed H-VPLS with either MPLS or Ethernet are unlikely to migrate to PBB technology over night because of obvious cost implications. Thus, it is imperative to outline migration strategies that will allow operators to protect investments in their installed base while still taking advantage of the scalability benefits of PBB technology.

In the following subsections, we explore three different migration scenarios that allow a mix of existing H-VPLS access networks to coexist with newer PBB-based access networks. The scenarios differ in whether or not the Ethernet service frames passing over the VPLS core are PBB-encapsulated. The first scenario, in Section 7.1, involves passing only frames that are not PBB-encapsulated over the core. The second scenario, in Section 7.2, stipulates passing only PBB-encapsulated frames over the core. Whereas, the final scenario, in Section 7.3, depicts a core that supports a mix of PBB-encapsulated and non-PBB-encapsulated frames. The advantages and disadvantages of each scenario will be discussed in the respective sections.

7.1. 802.1ad Service Frames over VPLS Core

In this scenario, existing access networks are left unchanged. All N-PEs would forward frames based on C-MACs. In other words, Ethernet frames that are traversing the VPLS core (within PWs) would use the 802.1ad frame format, as in current VPLS. Hence, the N-PEs in existing access networks do not require any modification. For new MPLS access networks that have PBB functions on the U-PE, the corresponding N-PE must incorporate built-in IB-BEB functions in order to terminate the PBB encapsulation before the frames enter the core. A key point here is that while both the U-PE and N-PE nodes implement PBB IB-BEB functionality, the former has the I-component facing the customer (CE) and the B-component facing the core; whereas the latter has the I-component facing the core and the B-component facing the customer (i.e., access network).




Figure 9: Migration with 802.1ad Service Frames over VPLS Core

The main advantage of this approach is that it requires no change to existing access networks or existing VPLS N-PEs. The main disadvantage is that these N-PEs will not leverage the advantages of PBB in terms of MAC address and PW scalability. It is worth noting that this migration scenario is an optimal option for an H-VPLS deployment with a single PBB-capable access network. When multiple PBB-capable access networks are required, then the scenario in Section 7.3 is preferred, as it provides a more scalable and optimal interconnect amongst the PBB-capable networks.

7.2. PBB Service Frames over VPLS Core

This scenario requires that the VPLS N-PE connecting to existing MPLS access networks be upgraded to incorporate IB-BEB functions. All Ethernet service frames passing over the VPLS core would be PBB-encapsulated. The PBB over MPLS access networks would require no special requirements beyond what is captured in Section 6 of this document. In this case, both the U-PE and N-PE, which implement IB-BEB functionality, have the I-component facing the customer and the B-component facing the core.




Figure 10: Migration with PBB Service Frames over VPLS Core

The main advantage of this approach is that it allows better scalability of the VPLS N-PEs in terms of MAC address and pseudowire counts. The disadvantage is that it requires upgrading the VPLS N-PEs of all existing MPLS access networks.

7.3. Mixed 802.1ad and PBB over VPLS Core

In this scenario, existing access networks are left unchanged, and they exchange Ethernet frames with 802.1ad format over the PWs in the core. The newly added access networks, which incorporate PBB functionality exchange Ethernet frames that are PBB-encapsulated amongst each other over core PWs. For service connectivity between existing access network (non-PBB-capable) and new access network (PBB based), the VPLS N-PE of the latter network employs IB-BEB functionality to decapsulate the PBB header from frames outbound to the core and encapsulate the PBB header for frames inbound from the core. As a result, a mix of PBB-encapsulated and 802.1ad Ethernet service frames are exchanged over the VPLS core.

This mode of operation requires new functionality on the VPLS N-PE of the PBB-capable access network, so that the PE can send frames in 802.1ad format or PBB format, on a per PW basis, depending on the capability of the destination access network. Effectively, the PE would have to incorporate B-BEB as well as IB-BEB functions.

A given PE needs to be aware of the capability of its remote peer in order to determine whether it connects to the right PW Forwarder. This can be achieved either via static configuration or by extending the VPLS control plane (BGP-based auto-discovery and LDP Signaling) discussed in [RFC6074]. The latter approach and the details of the extensions required are out of scope for this document.




Figure 11: Migration with Mixed 802.1ad and PBB Service Frames over VPLS Core

The U-PE and N-PE of the PBB-capable access network both employ BEB functionality. The U-PE implements IB-BEB functionality where the I-component faces the customer (CE) and the B-component faces the core. The N-PE, on the other hand, implements IB-BEB functionality with the I-component facing the core and the B-component facing the customer (access network). In addition, the N-PE implements standalone B-BEB functionality.

This scenario combines the advantages of both previous scenarios without any of their shortcomings, namely: it does not require any changes to existing access networks and it allows the N-PE to leverage the scalability benefits of 802.1ah for PBBN to PBBN connectivity. The disadvantage of this option is that it requires new functionality on the N-PE of the PBBN access. A second disadvantage is that this option requires two P2MP LSPs to be set up at the ingress N-PE: one for the N-PEs that support PBB encapsulation and another one for the N-PEs that don't support PBB encapsulation.

8. Acknowledgments

The authors would like to thank Chris Metz and Dinesh Mohan for their valuable feedback and contributions.

9. Security Considerations

This document does not introduce any additional security aspects beyond those applicable to VPLS/H-VPLS. VPLS/H-VPLS security considerations are already covered in [RFC4111] and [RFC4762].

10. References

10.1. Normative References

- [802.1ad] "IEEE Standard for and metropolitan area networks -- Virtual Bridged Local Area Networks -- Provider Bridges", 802.1ad-2005, August 2005.
- [802.1ah] "IEEE Standard for Local and metropolitan area networks -- Virtual Bridged Local Area Networks Amendment 7: Provider Backbone Bridges", IEEE Std. 802.1ah-2008, August 2009.
- [RFC4447] Martini, L., Ed., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", RFC 4447, April 2006.

- [RFC4448] Martini, L., Ed., Rosen, E., El-Aawar, N., and G. Heron, "Encapsulation Methods for Transport of Ethernet over MPLS Networks", RFC 4448, April 2006.
- [RFC4762] Lasserre, M., Ed., and V. Kompella, Ed., "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", RFC 4762, January 2007.
- [RFC6074] Rosen, E., Davie, B., Radoaca, V., and W. Luo, "Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)", RFC 6074, January 2011.

10.2. Informative References

- [802.1Q] "IEEE Standard for Local and metropolitan area networks - Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", IEEE Std 802.1Q(tm), 2012 Edition, October 2012.
- [802.1D-REV] "IEEE Standard for Local and metropolitan area networks Media Access Control (MAC) Bridges", IEEE Std. 802.1D, June 2004.
- [RFC6246] Sajassi, A., Ed., Brockners, F., Mohan, D., Ed., and Y. Serbest, "Virtual Private LAN Service (VPLS) Interoperability with Customer Edge (CE) Bridges", RFC 6246, June 2011.
- [RFC4111] Fang, L., Ed., "Security Framework for Provider-Provisioned Virtual Private Networks (PPVPNs)", RFC 4111, July 2005.

Authors' Addresses

Ali Sajassi
Cisco
EMail: sajassi@cisco.com

Samer Salam
Cisco
EMail: ssalam@cisco.com

Nabil Bitar
Verizon Communications
EMail : nabil.n.bitar@verizon.com

Florin Balus
Nuage Networks
EMail: florin.balus@nuagenetworks.net

