

Internet Engineering Task Force (IETF)  
Request for Comments: 7041  
Category: Informational  
ISSN: 2070-1721

F. Balus, Ed.  
Alcatel-Lucent  
A. Sajassi, Ed.  
Cisco  
N. Bitar, Ed.  
Verizon  
November 2013

Extensions to the Virtual Private LAN Service (VPLS)  
Provider Edge (PE) Model for Provider Backbone Bridging

Abstract

The IEEE 802.1 Provider Backbone Bridges (PBBs) specification defines an architecture and bridge protocols for interconnection of multiple Provider Bridged Networks (PBNs). Provider backbone bridging was defined by IEEE as a connectionless technology based on multipoint VLAN tunnels. PBB can be used to attain better scalability than Provider Bridges (PBs) in terms of the number of customer Media Access Control addresses and the number of service instances that can be supported.

The Virtual Private LAN Service (VPLS) provides a framework for extending Ethernet LAN services, using MPLS tunneling capabilities, through a routed MPLS backbone without running the Rapid Spanning Tree Protocol (RSTP) or the Multiple Spanning Tree Protocol (MSTP) across the backbone. As a result, VPLS has been deployed on a large scale in service provider networks.

This document discusses extensions to the VPLS Provider Edge (PE) model required to incorporate desirable PBB components while maintaining the service provider fit of the initial model.

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/rfc7041>.

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. General Terminology .....4
- 3. PE Reference Model .....6
- 4. Packet Walkthrough .....9
- 5. Control Plane .....11
- 6. Efficient Packet Replication in PBB VPLS .....12
- 7. PBB VPLS OAM .....12
- 8. Security Considerations .....12
- 9. References .....13
  - 9.1. Normative References .....13
  - 9.2. Informative References .....13
- 10. Contributors .....14
- 11. Acknowledgments .....15

## 1. Introduction

The IEEE 802.1 Provider Backbone Bridges specification [PBB] defines an architecture and bridge protocols for interconnection of multiple Provider Bridged Networks (PBNs). PBB can be used to attain better scalability than Provider Bridges [PB] in terms of the number of customer Media Access Control (MAC) addresses and the number of service instances that can be supported. PBB provides a data-plane hierarchy and new addressing designed to achieve such better scalability in Provider Backbone Networks. A number of Ethernet control-plane protocols, such as the Rapid Spanning Tree Protocol (RSTP), the Multiple Spanning Tree Protocol (MSTP), and Shortest Path Bridging (SPB), could be deployed as the core control plane for loop avoidance and load balancing for PBB. The applicability of these control protocols is out of scope for this document.

The Virtual Private LAN Service (VPLS) provides a solution for extending Ethernet LAN services, using MPLS tunneling capabilities, through a routed MPLS backbone without requiring the use of a native Ethernet control-plane protocol across the backbone. VPLS use of the structured FEC 129 [RFC4762] also allows for inter-domain, inter-provider connectivity and enables auto-discovery options across the network, improving the service delivery options.

A hierarchical solution for VPLS was introduced in [RFC4761] and [RFC4762] to provide improved scalability and efficient handling of packet replication. These improvements are achieved by reducing the number of Provider Edge (PE) devices connected in a full-mesh topology through the creation of two-tier PEs. A User-facing PE (U-PE) aggregates all the Customer Edge (CE) devices in a lower-tier access network and then connects to the Network-facing PE (N-PE) device(s) deployed around the core domain. In VPLS, Media Access Control (MAC) address learning and forwarding are done based on Customer MAC addresses (C-MACs); this poses scalability issues on the N-PE devices as the number of VPLS instances (and thus C-MACs) increases. Furthermore, since a set of pseudowires (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.

This document describes how PBB can be integrated with VPLS to allow for useful PBB capabilities while continuing to avoid the use of MSTP in the backbone. The combined solution referred to in this document

as PBB-VPLS results in better scalability in terms of the number of service instances, PWS, and C-MACs that need to be handled in the VPLS PEs.

Section 2 provides a quick terminology reference. Section 3 covers the reference model for PBB VPLS PEs. Section 4 describes the packet walkthrough. Sections 5 through 7 discuss the PBB-VPLS usage of existing VPLS mechanisms -- the control plane; efficient packet replication; and Operations, Administration, and Maintenance (OAM).

## 2. General Terminology

Some general terminology is defined here; most of the terminology used is from [PBB], [PB], [RFC4664], and [RFC4026]. Terminology specific to this memo is introduced as needed in later sections.

**B-BEB:** A backbone edge bridge positioned at the edge of a provider backbone bridged network. It contains a B-component that supports bridging in the provider backbone based on Backbone MAC (B-MAC) and B-tag information.

**B-component:** A bridging component contained in backbone edge and core bridges that bridges in the backbone space (B-MAC addresses, B-VLAN).

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

**B-tag:** Field defined in the PBB 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:** The interface between a BEB and a Backbone Core Bridge (BCB) in a provider backbone bridged network. Frames passed through this interface contain a B-tag field.

**B-VID:** The specific VLAN identifier carried inside a B-tag.

**B-VLAN:** The backbone VLAN associated with a B-component.

**B-PW:** The pseudowire used to interconnect B-component instances.

**BEB:** A backbone edge bridge positioned at the edge of a provider backbone bridged network. It can contain an I-component, a B-component, or both I-components and B-components.

**C-VID:** The VLAN identifier in a customer VLAN.

DA: Destination Address.

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, service VLAN IDs).

I-component: A bridging component contained in a backbone edge bridge that bridges in the customer space (customer MAC addresses, service VLAN identifier information (S-VLAN)).

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 PBB provider MAC encapsulation header that conveys the service instance information (I-SID) associated with the frame.

I-Tagged Service Interface: The interface defined between the I-components and B-components inside an IB-BEB or between two B-BEBs. Frames passed through this interface contain an I-tag field.

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, service VLAN IDs) and a B-component for bridging the provider's backbone space (B-MAC, B-tag).

PBs: Provider Bridges (IEEE amendment (802.1ad) to 802.1Q for "QinQ" encapsulation and bridging of Ethernet frames [PB]).

PBBs: Provider Backbone Bridges (IEEE amendment (802.1ah) to 802.1Q for "MAC tunneling" encapsulation and bridging of frames across a provider network [PBB]).

PBBN: Provider Backbone Bridged Network.

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

PSN: Packet-Switched Network.

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

S-Tagged Service Interface: 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.

SA: Source Address.

S-VID: The VLAN identifier in a service VLAN.

Tag: In Ethernet, a header immediately following the Source MAC Address field of the frame.

3. PE Reference Model

The following gives a short primer on the Provider Backbone Bridge (PBB) before describing the PE reference model for PBB-VPLS. The internal components of a PBB bridge module are depicted in Figure 1.

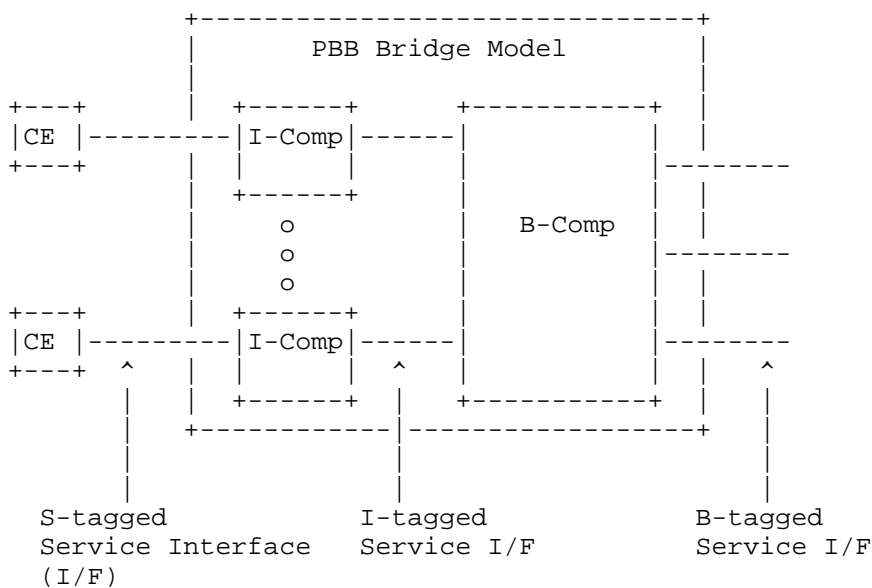


Figure 1: PBB Bridge Model

Provider Backbone Bridges (PBBs) [PBB] offer a scalable solution for service providers to build large bridged networks. The focus of PBB is primarily on improving two main areas with provider Ethernet bridged networks:

- MAC-address table scalability
- Service instance scalability

To obviate the above two limitations, PBB introduces a hierarchical network architecture with associated new frame formats that extend the work completed by Provider Bridges (PBs). In the PBBN architecture, customer networks (using PBs) are aggregated into PBBNs, which utilize the IEEE PBB frame format. The frame format employs a MAC tunneling encapsulation scheme for tunneling customer Ethernet frames within provider Ethernet frames across the PBBN. A VLAN identifier (B-VID) is used to segregate the backbone into broadcast domains, and a new 24-bit service identifier (I-SID) is defined and used to associate a given customer MAC frame with a provider service instance (also called the service delimiter). It should be noted that in [PBB] there is a clear segregation between provider service instances (represented by I-SIDs) and provider VLANs (represented by B-VIDs), which was not the case for PBs.

As shown in Figure 1, a PBB bridge may consist of a single B-component and one or more I-components. In simple terms, the B-component provides bridging in the provider space (B-MAC, B-VLAN), and the I-component provides bridging in the customer space (C-MAC, S-VLAN). The customer frame is first encapsulated with the provider backbone header (B-MAC, B-tag, I-tag); then, the bridging is performed in the provider backbone space (B-MAC, B-VLAN) through the network till the frame arrives at the destination BEB, where it gets decapsulated and passed to the CE. If a PBB bridge consists of both I-components and B-components, then it is called an IB-BEB, and if it only consists of either B-components or I-components, then it is called a B-BEB or an I-BEB, respectively. The interface between an I-BEB or IB-BEB and a CE is called an S-tagged service interface, and the interface between an I-BEB and a B-BEB (or between two B-BEBs) is called an I-tagged service interface. The interface between a B-BEB or IB-BEB and a Backbone Core Bridge (BCB) is called a B-tagged service interface.

To accommodate the PBB components, the VPLS model defined in [RFC4664] is extended as depicted in Figure 2.

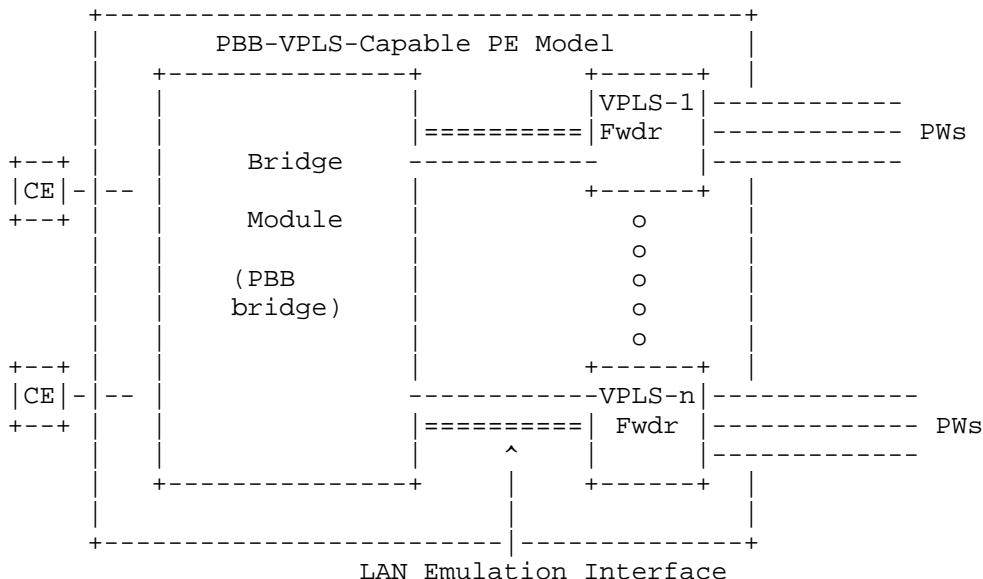


Figure 2: PBB-VPLS-Capable PE Model

The PBB module as defined in the [PBB] specification is expanded to interact with VPLS Forwarders. The VPLS Forwarders are used in [RFC4762] to build a PW mesh or a set of spoke PWs (Hierarchical VPLS (H-VPLS) topologies). The VPLS instances are represented externally in the MPLS context by a Layer 2 Forwarding Equivalence Class (L2FEC) that binds related VPLS instances together. VPLS Signaling advertises the mapping between the L2FEC and the PW labels and implicitly associates the VPLS bridging instance to the VPLS Forwarders [RFC4762].

In the PBB-VPLS case, the backbone service instance in the B-component space (B-VID) is represented in the backbone MPLS network using a VPLS instance. In the same way as for the regular VPLS case, existing signaling procedures are used to generate through PW labels the linkage between VPLS Forwarders and the backbone service instance.

Similarly, with the regular H-VPLS, another L2FEC may be used to identify the customer service instance in the I-component space. This will be useful, for example, to address the PBB-VPLS N-PE case where H-VPLS spokes are connecting the PBB-VPLS N-PE to a VPLS U-PE.



It is important to note that the PBB-VPLS solution inherits the PBB service aggregation capability where multiple customer service instances may be mapped to a backbone service instance. In the PBB-VPLS case, this means multiple customer VPNs can be transported using a single VPLS instance corresponding to the backbone service instance, thus substantially reducing resource consumption in the VPLS core.

#### 4. Packet Walkthrough

Since the PBB bridge module inherently performs forwarding, the PE reference model of Figure 2 can be expanded as shown in Figure 3.

Furthermore, the B-component is connected via several virtual interfaces to the PW Forwarder module. The function of the PW Forwarder is defined in [RFC3985]. In this context, the PW Forwarder simply performs the mapping of the PWs to the virtual interface on the B-component, without the need for any MAC lookup.

This simplified model takes full advantage of the PBB module -- where all the [PBB] procedures, including C-MAC/B-MAC forwarding and PBB encapsulation/decapsulation, take place -- and thus avoids the need to specify any of these functions in this document.

Because of text-based graphics, Figure 3 only shows PWs on the core-facing side; however, in the case of MPLS access with spoke PWs, the PE reference model is simply extended to include the same PW Forwarder function on the access-facing side. To avoid cluttering the figure, but without losing any generality, the access-side PW Forwarder (Fwdr) is not depicted.

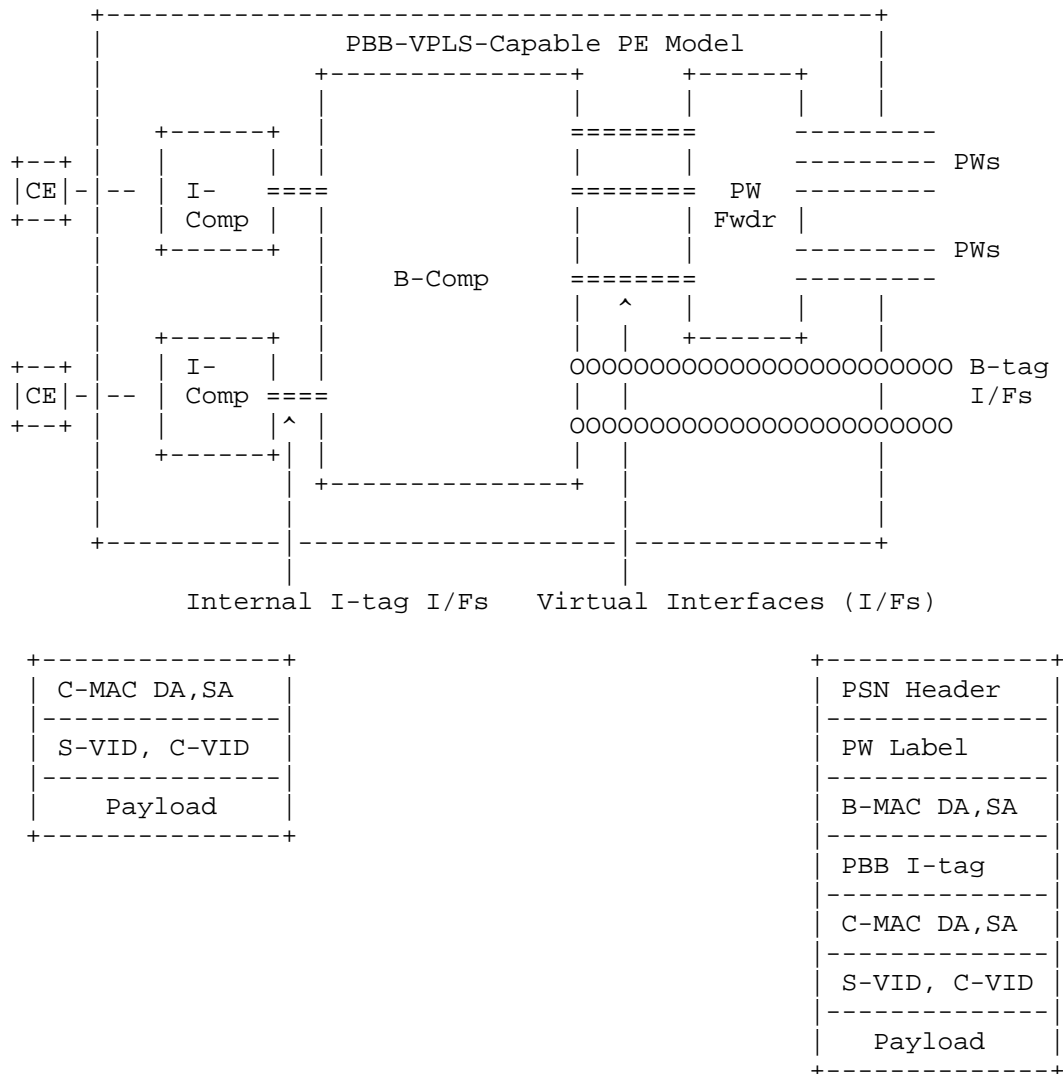


Figure 3: Packet Walkthrough for PBB VPLS PE

In order to better understand the data-plane walkthrough, let us consider the example of a PBB packet arriving over a Backbone pseudowire (B-PW). The PSN header is used to carry the PBB encapsulated frame over the backbone while the PW label will point to the related Backbone Service Instance (B-SI), in the same way as for regular VPLS. The PW label has in this case an equivalent role with the backbone VLAN identifier on the PBB B-tagged interface.

An example of the PBB packet for the regular Ethernet PW is depicted on the right-hand side of Figure 3. The MPLS packet from the MPLS core network is received by the PBB-VPLS PE. The PW Forwarder function of the PE uses the PW label to derive the virtual interface-id on the B-component, and then, after removing the PSN and PW encapsulation, it passes the packet to the B-component. From there on, the processing and forwarding are performed according to [PBB], where bridging based on the Backbone MAC (B-MAC) Destination Address (DA) is performed. This scenario results in one of the following outcomes:

1. The packet is forwarded to a physical interface on the B-component. In this case, the PBB Ethernet frame is forwarded as is.
2. The packet is forwarded to a virtual interface on the B-component. This is not typically the case, because of a single split-horizon group within a VPLS instance; however, if there is more than one split-horizon group, then such forwarding takes place. In this case, the PW Forwarder module adds the PSN and PW labels before sending the packet out.
3. The packet is forwarded toward the access side via one of the I-tagged service interfaces connected to the corresponding I-components. In this case, the I-component removes the B-MAC header according to [PBB] and bridges the packet using the C-MAC DA.

If the destination B-MAC is an unknown MAC address or a Group MAC address (multicast or broadcast), then the B-component floods the packet to one or more of the three destinations described above.

## 5. Control Plane

The control-plane procedures described in [RFC6074], [RFC4761], and [RFC4762] can be reused in a PBB-VPLS to set up the PW infrastructure in the service provider and/or customer bridging space. This allows porting the existing control-plane procedures (e.g., BGP Auto-Discovery (BGP-AD), PW setup, VPLS MAC flushing, PW OAM) for each domain.

## 6. Efficient Packet Replication in PBB VPLS

The PBB VPLS architecture takes advantage of the existing VPLS features addressing packet replication efficiency. The H-VPLS hierarchy may be used in both customer and backbone service instances to reduce the redundant distribution of packets over the core. IGMP and PIM snooping may be applied on a "per customer service instance" basis to control the distribution of the multicast traffic to non-member sites.

[IEEE-802.1Q] specifies the use of the Multiple MAC Registration Protocol (MMRP) for flood containment in the backbone instances. The same solution can be ported in the PBB-VPLS solution.

Further optimizations of the packet replication in PBB-VPLS are out of the scope of this document.

## 7. PBB VPLS OAM

The existing VPLS, PW, and MPLS OAM procedures may be used in each customer service instance or backbone service instance to verify the status of the related connectivity components.

PBB OAM procedures make use of the IEEE Ethernet Connectivity Fault Management [CFM] and ITU-T Y.1731 [Y.1731] tools in both I-components and B-components.

Both sets of tools (PBB and VPLS) may be used for the combined PBB-VPLS solution.

## 8. Security Considerations

No new security issues are introduced beyond those described in [RFC4761] and [RFC4762].

## 9. References

### 9.1. Normative References

- [RFC4761] Kompella, K., Ed., and Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", RFC 4761, January 2007.
- [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.

### 9.2. Informative References

- [RFC3985] Bryant, S., Ed., and P. Pate, Ed., "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, March 2005.
- [RFC4664] Andersson, L., Ed., and E. Rosen, Ed., "Framework for Layer 2 Virtual Private Networks (L2VPNs)", RFC 4664, September 2006.
- [PBB] Clauses 25 and 26 of "IEEE Standard for Local and metropolitan area networks - Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", IEEE Std 802.1Q-REV, 2013.
- [PB] Clauses 15 and 16 of "IEEE Standard for Local and metropolitan area networks - Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", IEEE Std 802.1Q-REV, 2013.
- [CFM] CFM clauses of "IEEE Standard for Local and metropolitan area networks - Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", IEEE Std 802.1Q-REV, 2013.
- [IEEE-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-REV, 2013.

[Y.1731] ITU-T Recommendation Y.1731, "OAM functions and mechanisms for Ethernet based networks", July 2011.

[RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, March 2005.

## 10. Contributors

The following people made significant contributions to this document:

Matthew Bocci  
Alcatel-Lucent  
Voyager Place  
Shoppenhangers Road  
Maidenhead  
Berks, UK

EMail: matthew.bocci@alcatel-lucent.com

Raymond Zhang  
Alcatel-Lucent

EMail: raymond.zhang@alcatel.com

Geraldine Calvignac  
Orange  
2, avenue Pierre-Marzin  
22307 Lannion Cedex  
France

EMail: geraldine.calvignac@orange.com

John Hoffmans  
KPN  
Regulusweg 1  
2516 AC Den Haag  
The Netherlands

EMail: john.hoffmans@kpn.com

Olen Stokes  
Extreme Networks  
PO Box 14129  
RTP, NC 27709  
USA

EEmail: ostokes@extremenetworks.com

## 11. Acknowledgments

The authors would like to thank Wim Henderickx, Mustapha Aissaoui, Dimitri Papadimitriou, Pranjal Dutta, Jorge Rabadan, Maarten Vissers, and Don Fedyk for their insightful comments and probing questions.

## Authors' Addresses

Florin Balus (editor)  
Alcatel-Lucent  
701 E. Middlefield Road  
Mountain View, CA 94043  
USA

EEmail: florin.balus@alcatel-lucent.com

Ali Sajassi (editor)  
Cisco  
170 West Tasman Drive  
San Jose, CA 95134  
USA

EEmail: sajassi@cisco.com

Nabil Bitar (editor)  
Verizon  
60 Sylvan Road  
Waltham, MA 02145  
USA

EEmail: nabil.n.bitar@verizon.com