

L2VPN Working Group  
Internet Draft  
Intended Status: Informational  
Expires: April 2013

F. Balus (editor)  
Alcatel-Lucent  
  
Ali Sajassi (editor)  
Cisco  
  
Nabil Bitar (editor)  
Verizon

October 22, 2012

**Extensions to VPLS PE model for Provider Backbone Bridging  
draft-ietf-l2vpn-pbb-vpls-pe-model-06.txt**

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

This Internet-Draft will expire on April 22, 2013.

Copyright Notice

Copyright (c) 2012 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.

## Abstract

IEEE 802.1ah standard [[IEEE802.1ah](#)], also known as Provider Backbone Bridges (PBB) defines an architecture and bridge protocols for interconnection of multiple Provider Bridge Networks (PBNs). PBB was defined in IEEE as a connectionless technology based on multipoint VLAN tunnels. MSTP is used as the core control plane for loop avoidance and load balancing. As a result, the coverage of the solution is limited by STP scale in the core of large service provider networks. PBB on the other hand can be used to attain better scalability in terms of number of customer MAC addresses and number of service instances that can be supported.

Virtual Private LAN Service (VPLS) [[RFC4664](#)] provides a framework for extending Ethernet LAN services, using MPLS tunneling capabilities, through a routed MPLS backbone without running (M)STP across the backbone. As a result, VPLS has been deployed on a large scale in service provider networks.

This draft discusses extensions to the VPLS PE model required to incorporate desirable PBB components while maintaining the Service Provider fit of the initial model.

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction.....</a>	<a href="#">3</a>
<a href="#">2.</a>	<a href="#">General terminology.....</a>	<a href="#">4</a>
<a href="#">3.</a>	<a href="#">PE Reference Model.....</a>	<a href="#">5</a>
<a href="#">4.</a>	<a href="#">Packet Walkthrough.....</a>	<a href="#">8</a>
<a href="#">5.</a>	<a href="#">Control Plane.....</a>	<a href="#">10</a>
<a href="#">6.</a>	<a href="#">Efficient Packet replication in PBB VPLS.....</a>	<a href="#">11</a>
<a href="#">7.</a>	<a href="#">PBB VPLS OAM.....</a>	<a href="#">11</a>
<a href="#">8.</a>	<a href="#">Security Considerations.....</a>	<a href="#">11</a>
<a href="#">9.</a>	<a href="#">IANA Considerations.....</a>	<a href="#">11</a>
<a href="#">10.</a>	<a href="#">References.....</a>	<a href="#">11</a>
	<a href="#">10.1. Normative References.....</a>	<a href="#">11</a>
	<a href="#">10.2. Informative References.....</a>	<a href="#">12</a>
<a href="#">11.</a>	<a href="#">Contributors.....</a>	<a href="#">12</a>
<a href="#">12.</a>	<a href="#">Acknowledgments.....</a>	<a href="#">12</a>



## **1. Introduction**

IEEE 802.1ah standard [[IEEE802.1ah](#)], also known as Provider Backbone Bridges (PBB) defines an architecture and bridge protocols for interconnection of multiple Provider Bridge Networks (PBNs). PBB provides data plane hierarchy and new addressing designed to improve the scalability of MAC addresses and service instances in Provider Backbone Networks. MSTP is still used as the core control plane for loop avoidance and load balancing. As a result, the coverage of the solution is limited by STP scale in the core of large service provider networks.

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 (M)STP 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 [[RFC4762](#)] for the purpose of improved scalability and to provide efficient handling of packet replication. These improvements are achieved by reducing the number of PE devices connected in a full-mesh topology through the creation of two-tier PEs. A U-PE aggregates all the CE devices in a lower-tier access network and then connects to the N-PE device(s) deployed around the core domain. In VPLS, MAC address learning and forwarding are done based on customer MAC addresses (C-MACs), which poses scalability issues on the N-PE devices as the number of VPLS instances (and thus customer MAC addresses) 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.

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

[Section 2](#) gives a quick terminology reference. [Section 3](#) covers the reference model for PBB VPLS PE. [Section 4](#) describes the packet walkthrough. [Section 5](#) to 7 discusses the PBB-VPLS usage of existing VPLS mechanisms - control plane, efficient packet replication, OAM.



## **2. General terminology**

Some general terminology is defined here; most of the terminology used is from [[IEEE802.1ah](#)], [[RFC4664](#)] and [[RFC4026](#)]. Terminology specific to this memo is introduced as needed in later sections.

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.

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 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.

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 I and B components.

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

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

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, service 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, service 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 the interface defined between the I and B components inside an IB-BEB or between two B-BEB. Frames passed through this interface contain an I-TAG field

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 identifier information (S-VLAN).

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

### **3. PE Reference Model**

The following gives a short primer on 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) [IEEE 802.1ah] offers 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 which extend the work completed by Provider Bridges (IEEE 802.1ad). In the PBBN architecture, customer networks (using IEEE 802.1Q or 802.1ad bridging) are aggregated into Provider Backbone Bridge Networks (PBBNs) which utilize the IEEE 802.1ah 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 [802.1ah] 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 802.1ad.



As shown in the 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 provider space (B-MAC, B-VLAN) and the I-component provides bridging in 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 de-encapsulated and passed to the CE. If a PBB bridge consists of both I & B components, then it is called IB-BEB and if it only consists of either B-component or I-component, then it is called B-BEB or I-BEB respectively. The interface between an I-BEB or IB-BEB and a CE is called S-tagged service interface and the interface between an I-BEB and a B-BEB (or between two B-BEBs) is called I-tagged service interface. The interface between a B-BEB or IB-BEB and a Backbone Core Bridge (BCB) is called B-Tagged service interface.

To accommodate the PBB components the VPLS model defined in [\[RFC4664\]](#) is extended as depicted in figure 1.

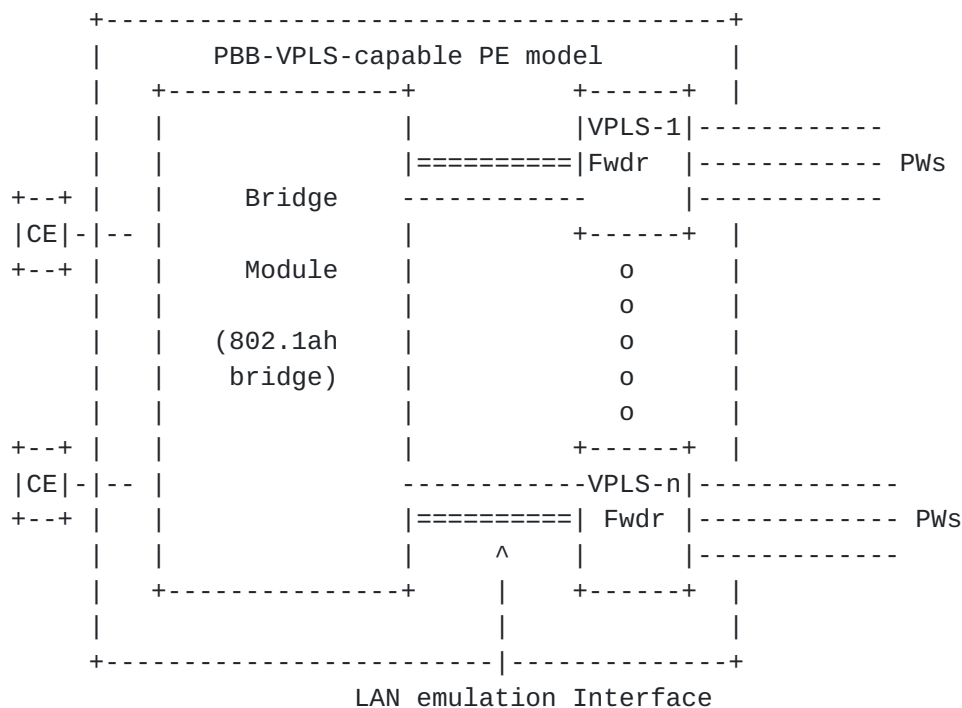


Figure 2: PBB-VPLS capable PE Model



The PBB Module as defined in [[IEEE802.1ah](#)] 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 (HVPLS topologies). The VPLS instances are represented externally in the MPLS context by a L2FEC which 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. Same 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 HVPLS, 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 HVPLS 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 reducing substantially resource consumption in the VPLS core.

#### **4. Packet Walkthrough**

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

Furthermore, the B-component is connected via several virtual interfaces to the PW Forwarder module. The function of 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 PBB bridge module where all the [[IEEE 802.1ah](#)] procedures including the C-MAC/B-MAC forwarding and PBB encapsulation/decapsulation takes place and thus avoids specifying any of these functions in here.



Because of text-based graphics, the Figure 3 only shows PWs on the core-facing side; however, in 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, the access-side PW Forwarder is not depicted without loss of any generality.

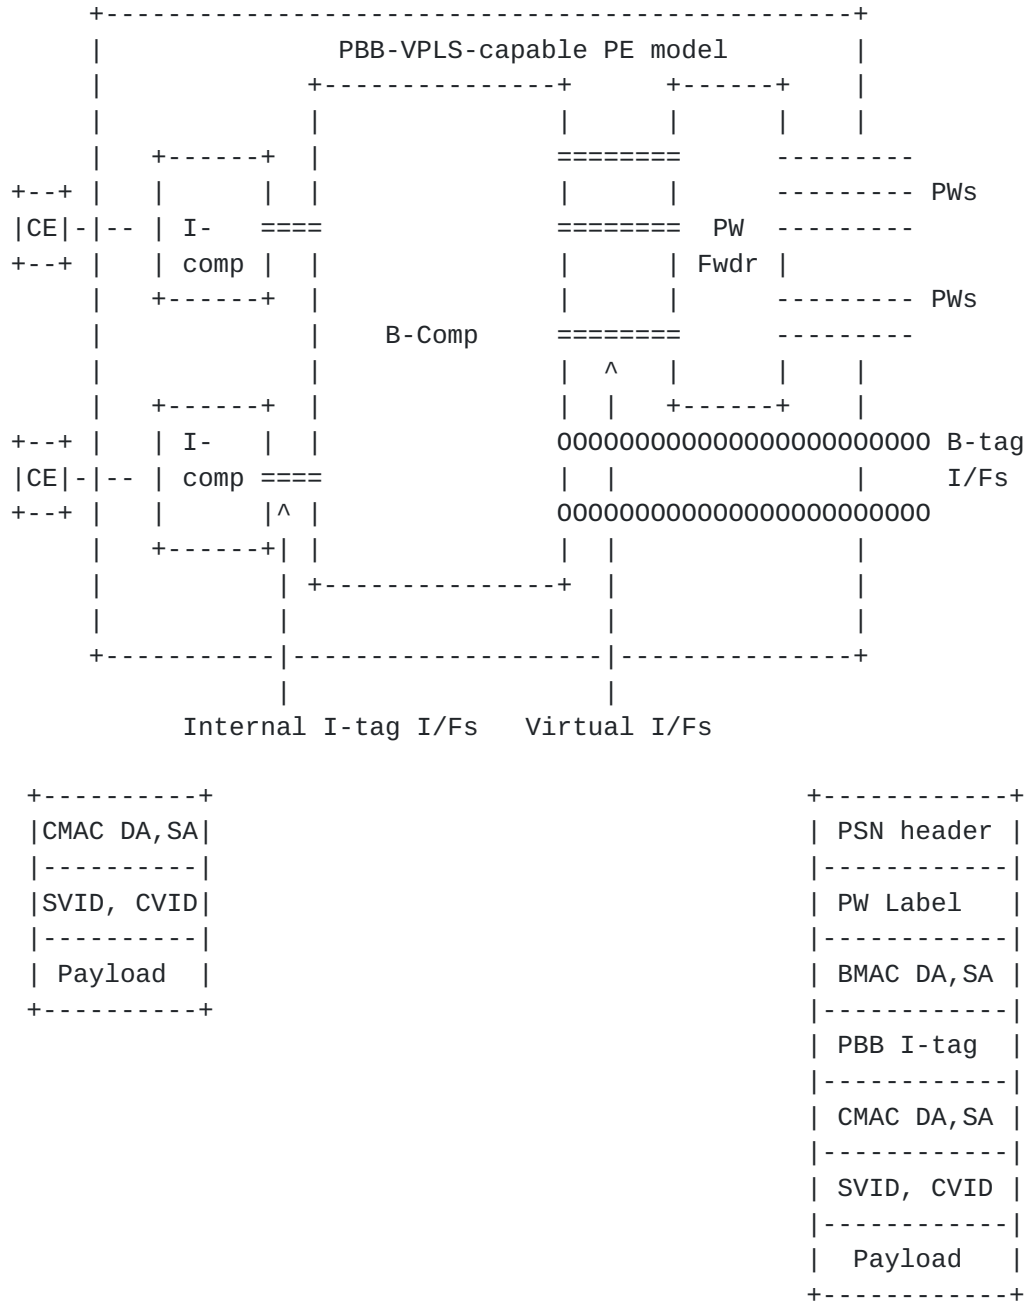


Figure 3: Packet Walkthrough for PBB VPLS PE





In order to better understand the data plane walkthrough let us consider the example of a PBB packet arriving over a 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), same as for regular VPLS. The PW Label has in this case an equivalent role with the Backbone VLAN id on the PBB B-tagged interface.

An example of the PBB packet for regular Ethernet PW is depicted in Figure 3 on the right hand side. The MPLS packet from MPLS core network is received by the PBB-VPLS PE. The PW Forwarder function of the PE uses 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 is performed according to the [IEEE 802.1ah] where bridging based on B-MAC DA is performed which result in one of the three outcomes:

1. The packet is forwarded to a physical interface on the B-component. In this case, the 802.1ah 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 scenario, the I-component removes the B-MAC header according to [IEEE 802.1ah] and bridges the packet using C-MAC DA.
4. If the destination B-MAC is an unknown 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 re-used in a PBB-VPLS to setup the PW infrastructure in the service provider and/or customer bridging space. This allows porting the existing control plane procedures (e.g. BGP-AD, PW setup, VPLS MAC Flush, 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. HVPLS 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 to control the distribution of the Multicast traffic to non-member sites.

[IEEE802.1ah] specifies also the use of MMRP protocol [[IEEE802.1ak](#)] 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 draft.

## **7. PBB VPLS OAM**

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

PBB OAM procedures make use of the IEEE 802.1ag and Y.1731 tools in both I-component and B-component.

Both set 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 that are described in [[RFC4761](#)] and [[RFC4762](#)].

## **9. IANA Considerations**

IANA does not need to take any action for this draft.

## **10. References**

### **10.1. Normative References**

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

[RFC4762] Lasserre, M. and Kompella, V. (Editors), "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", [RFC 4762](#), January 2007.

[RFC6074] E. Rosen, et Al. "Provisioning, Autodiscovery and Signaling in L2VPNs", [RFC 6074](#), January 2011

## **[10.2. Informative References](#)**

[RFC3985] Bryant, S. and Pate, P. (Editors), "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", [RFC 3985](#), May 2005.

[RFC4664] Andersson, L. and Rosen, E. (Editors), "Framework for Layer 2 Virtual Private Networks (L2VPNs)", [RFC 4664](#), Sept 2006

[IEEE802.1ah] IEEE 802.1ah "Virtual Bridged Local Area Networks, Amendment 6: Provider Backbone Bridges", Approved Standard June 12th, 2008

[IEEE802.1ak] IEEE Draft P802.1ak/D8.0 "Virtual Bridged Local Area Networks, Amendment 7: Multiple Registration Protocol", Work in Progress, November 29, 2006

[RFC4026] Andersson, L. et Al., "Provider Provisioned Virtual Private Network (VPN) Terminology", [RFC 4026](#), May 2005.

## **[11. Contributors](#)**

The following authors contributed to this document: John Hoffmans (KPN), Geraldine Calvignac (France Telecom), Olen Stokes (Extreme Networks), Raymond Zhang and Matthew Bocci (Alcatel-Lucent).

## **[12. Acknowledgments](#)**

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

Authors' Addresses

Ali Sajassi  
Cisco  
170 West Tasman Drive  
San Jose, CA 95134, U.S.  
Email: sajassi@cisco.com

Nabil Bitar  
Verizon  
40 Sylvan Road  
Waltham, MA 02145  
Email: nabil.bitar@verizon.com

Florin Balus  
Alcatel-Lucent  
701 E. Middlefield Road  
Mountain View, CA, USA 94043  
Email: florin.balus@alcatel-lucent.com

Matthew Bocci  
Alcatel-Lucent,  
Voyager Place  
Shoppenhangers Road  
Maidenhead  
Berks, UK  
e-mail: matthew.bocci@alcatel-lucent.co.uk

Raymond Zhang  
BT  
2160 E. Grand Ave.  
El Segundo, CA 900245 USA  
EMail: raymond.zhang@bt.com

Geraldine Calvignac  
France Telecom  
2, avenue Pierre-Marzin  
22307 Lannion Cedex  
France  
Email: geraldine.calvignac@orange-ftgroup.com

John Hoffmans  
KPN  
Regulusweg 1  
2516 AC Den Haag  
Nederland  
Email: john.hoffmans@kpn.com

Olen Stokes  
Extreme Networks  
PO Box 14129  
RTP, NC 27709  
USA  
Email: [ostokes@extremenetworks.com](mailto:ostokes@extremenetworks.com)