Internet Engineering Task Force (IETF)
Request for Comments: 8560
Category: Standards Track
ISSN: 2070-1721
May 2019

A. Sajassi, Ed.
S. Salam
Cisco
N. Del Regno
Verizon
J. Rabadan
Nokia

Seamless Integration of Ethernet VPN (EVPN) with Virtual Private LAN Service (VPLS) and Their Provider Backbone Bridge (PBB) Equivalents

Abstract

This document specifies mechanisms for backward compatibility of Ethernet VPN (EVPN) and Provider Backbone Bridge Ethernet VPN (PBB-EVPN) solutions with Virtual Private LAN Service (VPLS) and Provider Backbone Bridge VPLS (PBB-VPLS) solutions. It also provides mechanisms for the seamless integration of these two technologies in the same MPLS/IP network on a per-VPN-instance basis. Implementation of this document enables service providers to introduce EVPN/PBB-EVPN Provider Edges (PEs) in their brownfield deployments of VPLS/PBB-VPLS networks. This document specifies the control-plane and forwarding behavior needed for the auto-discovery of the following: 1) a VPN instance, 2) multicast and unicast operation, and 3) a Media Access Control (MAC) mobility operation. This enables seamless integration between EVPN and VPLS PEs as well as between PBB-VPLS and PBB-EVPN PEs.

Copyright Notice

Table of Contents

1. Introduction

Virtual Private LAN Service (VPLS) and Provider Backbone Bridging VPLS (PBB-VPLS) are widely deployed Layer 2 VPN (L2VPN) technologies. Many service providers who are looking at adopting Ethernet VPN (EVPN) and Provider Backbone Bridging EVPN (PBB-EVPN) want to preserve their investments in the VPLS and PBB-VPLS networks. Hence, they require mechanisms by which EVPN and PBB-EVPN technologies can be introduced into their brownfield VPLS and PBB-VPLS networks without requiring any upgrades (software or hardware) to these networks. This document specifies procedures for the seamless integration of the two technologies in the same MPLS/IP network. Throughout this document, we use the term "(PBB-)EVPN" to correspond to both EVPN and PBB-EVPN, and we use the term "(PBB-)VPLS" to correspond to both VPLS and PBB-VPLS. This document specifies the control-plane and forwarding behavior needed for 1) auto-discovery of a VPN instance, 2) multicast and unicast operations, and 3) a MAC mobility operation. This enables seamless integration between (PBB-)EVPN Provider Edge (PE) devices and (PBB-)VPLS PEs.

```
                        VPLS PE
                         +---+
                         |PE1|
                         +---+
                           /
   EVPN/VPLS PE  +---------------+   EVPN/VPLS PE
        +---+    |               |    +---+
        |PE4|----|    MPLS/IP    |---|PE5|
        +---+    |     Core      |    +---+
                 |               |
                 +---------------+
                   /       \
                 +---+     +---+
                 |PE2|     |PE3|
                 +---+     +---+
               VPLS PE     VPLS PE
```

Figure 1: Seamless Integration of (PBB-)EVPN and (PBB-)VPLS

Section 2 provides the details of the requirements. Section 3 specifies procedures for the seamless integration of VPLS and EVPN networks. Section 4 specifies procedures for the seamless integration of PBB-VPLS and PBB-EVPN networks.

It should be noted that the scenarios for both PBB-VPLS integration with EVPN and VPLS integration with PBB-EVPN are not covered in this document because there haven't been any requirements from service providers for these scenarios; deployments that employ PBB-VPLS typically require PBB encapsulation for various reasons. Hence, it is expected that for those deployments, the evolution path would move from PBB-VPLS towards PBB-EVPN. Furthermore, the evolution path from VPLS is expected to move towards EVPN.

The seamless integration solution described in this document has the following attributes:

- When ingress replication is used for multi-destination traffic delivery, the solution reduces the scope of MMRP (which is a soft-state protocol defined in Clause 10 of [IEEE.802.1Q]) to only that of existing VPLS PEs and uses the more robust BGP-based mechanism for multicast pruning among new EVPN PEs.

- It is completely backward compatible.

- New PEs can leverage the extensive multihoming mechanisms and provisioning simplifications of (PBB-)EVPN:

  (a) Auto-sensing of Multihomed Networks (MHNs) / Multihomed Devices (MHDs)

  (b) Auto-discovery of redundancy groups

  (c) Auto-provisioning of Designated Forwarder election and VLAN carving

1.1. Specification of Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

1.2. Abbreviations

B-MAC: Backbone MAC, e.g., the PE's MAC address

C-MAC: Customer MAC, e.g., a host or CE's MAC address

CE: A Customer Edge device, e.g., a host, router, or switch

ES: Ethernet Segment -- refers to the set of Ethernet links that connects a customer site (device or network) to one or more PEs

FEC: Forwarding Equivalence Class

FIB: Forwarding Information Base -- an instantiation of a forwarding table on a MAC-VRF

I-SID: Service Instance Identifier

LSP: Label Switched Path

MAC: Media Access Control

MAC-VRF: A Virtual Routing and Forwarding table for Media Access Control (MAC) addresses on an EVPN PE

MHD: Multihomed Device

MHN: Multihomed Network

MP2P: Multipoint to Point -- an MP2P LSP typically refers to an LSP for unicast traffic as the result of a downstream-assigned label

P2MP: Point to Multipoint -- a P2MP LSP typically refers to an LSP for multicast traffic

PBB: Provider Backbone Bridge

(PBB-)EVPN: Both PBB-EVPN and EVPN -- this document uses this abbreviation when a given description applies to both technologies

(PBB-)VPLS: Both PBB-VPLS and VPLS -- this document uses this abbreviation when a given description applies to both technologies

PE: Provider Edge device

PW: Pseudowire

RIB: Routing Information Base -- an instantiation of a routing table on a MAC-VRF

VSI: Virtual Switch Instance

VPLS: Virtual Private LAN Service

VPLS A-D: Virtual Private LAN Service with BGP-based auto-discovery as in [RFC6074]

1.3. Terminology

All-Active redundancy mode: When all PEs attached to an Ethernet segment are allowed to forward known unicast traffic to/from that Ethernet segment for a given VLAN, then the Ethernet segment is defined as operating in All-Active redundancy mode.

Bridge table: An instantiation of a broadcast domain on a MAC-VRF (VPN Routing and Forwarding).

Broadcast domain: In a bridged network, the broadcast domain corresponds to a Virtual LAN (VLAN), where a VLAN is typically represented by a single VLAN ID (VID) but can be represented by several VIDs where Shared VLAN Learning (SVL) is used, per [IEEE.802.1Q].

Ethernet Tag: An Ethernet Tag identifies a particular broadcast domain, e.g., a VLAN. An EVPN instance consists of one or more broadcast domains.

Single-Active redundancy mode: When only a single PE, among all the PEs attached to an Ethernet segment, is allowed to forward traffic to/from that Ethernet segment for a given VLAN, then the Ethernet segment is defined as operating in Single-Active redundancy mode.

2. Requirements

The following are the key requirements for backward compatibility between (PBB-)EVPN and (PBB-)VPLS:

1. The solution must allow for staged migration towards (PBB-)EVPN on a site-by-site basis per VPN instance, e.g., new EVPN sites to be provisioned on (PBB-)EVPN Provider Edge devices (PEs).

2. The solution must not require any changes to existing VPLS or PBB-VPLS PEs, not even a software upgrade.

3. The solution must allow for the coexistence of PE devices running (PBB-)EVPN and (PBB-)VPLS for the same VPN instance and single-homed segments.

4. The solution must support single-active redundancy of multihomed networks and multihomed devices for (PBB-)EVPN PEs.

5. In cases of single-active redundancy, the participant VPN instances may span across both (PBB-)EVPN PEs and (PBB-)VPLS PEs as long as the MHD or MHN is connected to (PBB-)EVPN PEs.

6. Support of the All-Active redundancy mode across both (PBB-)EVPN PEs and (PBB-)VPLS PEs is outside the scope of this document. All-Active redundancy is not applicable to VPLS and PBB-VPLS. Therefore, when EVPN (or PBB-EVPN) PEs need to operate seamlessly with VPLS (or PBB-VPLS) PEs, they MUST use a redundancy mode that is applicable to VPLS (or PBB-VPLS). This redundancy mode is Single-Active.

These requirements collectively allow for the seamless insertion of (PBB-)EVPN technology into brownfield (PBB-)VPLS deployments.

3. VPLS Integration with EVPN

In order to support seamless integration with VPLS PEs, this document requires that VPLS PEs support VPLS A-D per [RFC6074], and it requires EVPN PEs to support both BGP EVPN routes per [RFC7432] and VPLS A-D per [RFC6074]. All the logic for seamless integration shall reside on the EVPN PEs. If a VPLS instance is set up without the use of VPLS A-D, it is still possible (but cumbersome) for EVPN PEs to integrate that VPLS instance by manually configuring pseudowires (PWs) to all the VPLS PEs in that instance (i.e., the integration is no longer seamless).

3.1. Capability Discovery

The EVPN PEs MUST advertise both the BGP VPLS auto-discovery (A-D) route as well as the BGP EVPN Inclusive Multicast Ethernet Tag (IMET) route for a given VPN instance. The VPLS PEs only advertise the BGP VPLS A-D route, per the procedures specified in [RFC4761], [RFC4762] and [RFC6074]. The operator may decide to use the same Route Target (RT) to identify a VPN on both EVPN and VPLS networks. In this case, when a VPLS PE receives the EVPN IMET route, it MUST ignore it on the basis that it belongs to an unknown Subsequent Address Family Identifier (SAFI). However, the operator may choose to use two RTs -- one to identify the VPN on the VPLS network and another for the EVPN network -- and employ RT Constrain mechanisms [RFC4684] in order to prevent BGP EVPN routes from reaching the VPLS PEs.

When an EVPN PE receives both a VPLS A-D route as well as an EVPN IMET route from a given remote PE for the same VPN instance, it MUST give preference to the EVPN route for the purpose of discovery. This ensures that, at the end of the route exchanges, all EVPN-capable PEs discover other EVPN-capable PEs in addition to the VPLS-only PEs for that VPN instance. Furthermore, all the VPLS-only PEs will discover the EVPN PEs as if they were standard VPLS PEs. In other words, when the discovery phase is complete, the EVPN PEs will have discovered all the PEs in the VPN instance along with their associated capability (EVPN or VPLS-only), whereas the VPLS PEs will have discovered all the PEs in the VPN instance as if they were all VPLS-only PEs.

3.2. Forwarding Setup and Unicast Operation

The procedures for the forwarding state setup and unicast operation on the VPLS PE are per [RFC8077], [RFC4761], and [RFC4762].

The procedures for forwarding state setup and unicast operation on the EVPN PE are as follows:

- The EVPN PE MUST establish a PW to each remote PE from which it has received only a VPLS A-D route for the corresponding VPN instance and MUST set up the label stack corresponding to the PW FEC. For seamless integration between EVPN and VPLS PEs, the PW that is set up between a pair of VPLS and EVPN PEs is between the VSI of the VPLS PE and the MAC-VRF of the EVPN PE.

- The EVPN PE MUST set up the label stack corresponding to the MP2P VPN unicast FEC to any remote PE that has advertised an EVPN IMET route.

- If an EVPN PE receives a VPLS A-D route from a given PE, it sets up a PW to that PE. If it then receives an EVPN IMET route from the same PE, the EVPN PE MUST bring that PW operationally down.

- If an EVPN PE receives an EVPN IMET route followed by a VPLS A-D route from the same PE, then the EVPN PE will set up the PW but MUST keep it operationally down.

- In case VPLS A-D is not used in some VPLS PEs, the EVPN PEs need to be provisioned manually with PWs to those remote VPLS PEs for each VPN instance. In that case, if an EVPN PE receives an EVPN IMET route from a PE to which a PW exists, the EVPN PE MUST bring the PW operationally down.

When the EVPN PE receives traffic over the VPLS PWs, it learns the associated C-MAC addresses in the data plane. The C-MAC addresses learned over these PWs MUST be injected into the bridge table of the associated MAC-VRF on that EVPN PE. The learned C-MAC addresses MAY also be injected into the RIB/FIB tables of the associated MAC-VRF on that EVPN PE. For seamless integration between EVPN and VPLS PEs, because these PWs belong to the same split-horizon group (see [RFC4761] and [RFC4762]) as the MP2P EVPN service tunnels, the C-MAC addresses learned and associated with the PWs MUST NOT be advertised in the control plane to any remote EVPN PEs. This is because every

EVPN PE can send and receive traffic directly to/from every VPLS PE belonging to the same VPN instance; thus, every EVPN PE can learn the C-MAC addresses over the corresponding PWs directly.

The C-MAC addresses learned over local Attachment Circuits (ACs) by an EVPN PE are learned in the data plane. For EVPN PEs, these C-MAC addresses MUST be injected into the corresponding MAC-VRF and advertised in the control plane using BGP EVPN routes. Furthermore, the C-MAC addresses learned in the control plane via the BGP EVPN routes sent by remote EVPN PEs are injected into the corresponding MAC-VRF table.

In case of a link failure in a single-active Ethernet segment, the EVPN PEs MUST perform both of the following tasks:

1. send a BGP mass withdraw to the EVPN peers

2. follow existing VPLS MAC Flush procedures with the VPLS peers

3.3. MAC Mobility

In EVPN, host addresses (C-MAC addresses) can move around among EVPN PEs or even between EVPN and VPLS PEs.

When a C-MAC address moves from an EVPN PE to a VPLS PE, as soon as Broadcast, Unknown Unicast, and Multicast (BUM) traffic is initiated from that MAC address, it is flooded to all other PEs (both VPLS and EVPN PEs), and the receiving PEs update their MAC tables (VSI or MAC-VRF). The EVPN PEs do not advertise the C-MAC addresses learned over the PW to each other because every EVPN PE learns them directly over its associated PW to that VPLS PE. If only known unicast traffic is initiated from the moved C-MAC address toward a known C-MAC, the result can be the black-holing of traffic destined to the C-MAC that has moved until there is BUM traffic that has been originated with the moved C-MAC address as the source MAC address (e.g., as a result of the MAC age-out timer expiring). Such black-holing happens for traffic destined to the moved C-MAC from both EVPN and VPLS PEs and is typical for VPLS PEs.

When a C-MAC address moves from a VPLS PE to an EVPN PE, then as soon as any traffic is initiated from that C-MAC address, the C-MAC is learned and advertised in the BGP to other EVPN PEs, and the MAC mobility procedure is performed among EVPN PEs. For BUM traffic, both EVPN and VPLS PEs learn the new location of the moved C-MAC address; however, if there is only known unicast traffic, then only EVPN PEs learn the new location of the C-MAC that has moved and not VPLS PEs. This can result in the black-holing of traffic sent from VPLS PEs destined to the C-MAC that has moved until there is BUM traffic originated with the moved C-MAC address as the source MAC address (e.g., as a result of the MAC age-out timer expiring). Such black-holing happens for traffic destined to the moved C-MAC for only VPLS PEs but not for EVPN PEs and is typical for VPLS PEs.

3.4. Multicast Operation

3.4.1. Ingress Replication

The procedures for multicast operation on the VPLS PE using ingress replication are per [RFC4761], [RFC4762], and [RFC7080].

The procedures for multicast operation on the EVPN PE for ingress replication are as follows:

- The EVPN PE builds a replication sub-list to all the remote EVPN PEs per EVPN instance as the result of the exchange of the EVPN IMET routes per [RFC7432]. This will be referred to as sub-list A. It comprises MP2P service tunnels (for ingress replication) used for delivering EVPN BUM traffic [RFC7432].

- The EVPN PE builds a replication sub-list per VPLS instance to all the remote VPLS PEs. This will be referred to as sub-list B. It comprises PWs from the EVPN PE in question to all the remote VPLS PEs in the same VPLS instance.

The replication list, maintained per VPN instance, on a given EVPN PE will be the union of sub-list A and sub-list B. The EVPN PE MUST enable split horizon over all the entries in the replication list across both PWs and MP2P service tunnels.

3.4.2. P2MP Tunnel

The procedures for multicast operation on the EVPN PEs using P2MP tunnels are outside of the scope of this document.

4. PBB-VPLS Integration with PBB-EVPN

In order to support seamless integration between PBB-VPLS and PBB-EVPN PEs, this document requires that PBB-VPLS PEs support VPLS A-D per [RFC6074] and PBB-EVPN PEs support both BGP EVPN routes per [RFC7432] and VPLS A-D per [RFC6074]. All the logic for this seamless integration shall reside on the PBB-EVPN PEs.

4.1. Capability Discovery

The procedures for capability discovery are per Section 3.1.

4.2. Forwarding Setup and Unicast Operation

The procedures for forwarding state setup and unicast operation on the PBB-VPLS PE are per [RFC8077] and [RFC7080].

The procedures for forwarding state setup and unicast operation on the PBB-EVPN PE are as follows:

- The PBB-EVPN PE MUST establish a PW to each remote PBB-VPLS PE from which it has received only a VPLS A-D route for the corresponding VPN instance and MUST set up the label stack corresponding to the PW FEC. For seamless integration between PBB-EVPN and PBB-VPLS PEs, the PW that is set up between a pair of PBB-VPLS and PBB-EVPN PEs is between the B-components of PBB-EVPN PE and PBB-VPLS PE per Section 4 of [RFC7041].

- The PBB-EVPN PE MUST set up the label stack corresponding to the MP2P VPN unicast FEC to any remote PBB-EVPN PE that has advertised an EVPN IMET route.

- If a PBB-EVPN PE receives a VPLS A-D route from a given PE, it sets up a PW to that PE. If it then receives an EVPN IMET route from the same PE, the PBB-EVPN PE MUST bring that PW operationally down.

- If a PBB-EVPN PE receives an EVPN IMET route followed by a VPLS A-D route from the same PE, then the PBB-EVPN PE will set up the PW but MUST keep it operationally down.

- In case VPLS A-D is not used in some PBB-VPLS PEs, the PBB-EVPN PEs need to be provisioned manually with PWs to those remote PBB-VPLS PEs for each VPN instance. In that case, if a PBB-EVPN PE receives an EVPN IMET route from a PE to which a PW exists, the PBB-EVPN PE MUST bring the PW operationally down.

- When the PBB-EVPN PE receives traffic over the PBB-VPLS PWs, it learns the associated B-MAC addresses in the data plane. The B-MAC addresses learned over these PWs MUST be injected into the bridge table of the associated MAC-VRF on that PBB-EVPN PE. The learned B-MAC addresses MAY also be injected into the RIB/FIB tables of the associated MAC-VRF on that BPP-EVPN PE. For seamless integration between PBB-EVPN and PBB-VPLS PEs, since these PWs belong to the same split-horizon group as the MP2P EVPN service tunnels, the B-MAC addresses learned and associated with the PWs MUST NOT be advertised in the control plane to any remote PBB-EVPN PEs. This is because every PBB-EVPN PE can send and receive traffic directly to/from every PBB-VPLS PE belonging to the same VPN instance.

- The C-MAC addresses learned over local Attachment Circuits (ACs) by a PBB-EVPN PE are learned in the data plane. For PBB-EVPN PEs, these C-MAC addresses are learned in the I-component of PBB-EVPN PEs and are not advertised in the control plane, per [RFC7623].

- The B-MAC addresses learned in the control plane via the BGP EVPN routes sent by remote PBB-EVPN PEs are injected into the corresponding MAC-VRF table.

In case of a link failure in a single-active Ethernet segment, the PBB-EVPN PEs MUST perform both of the following tasks:

1. send a BGP B-MAC withdraw message to the PBB-EVPN peers *or* MAC advertisement with the MAC Mobility extended community

2. follow existing VPLS MAC Flush procedures with the PBB-VPLS peers

4.3. MAC Mobility

In PBB-EVPN, a given B-MAC address can be learned either over the BGP control plane from a remote PBB-EVPN PE or in the data plane over a PW from a remote PBB-VPLS PE. There is no mobility associated with B-MAC addresses in this context. Hence, when the same B-MAC address shows up behind both a remote PBB-VPLS PE as well as a PBB-EVPN PE, the local PE can deduce that it is an anomaly and SHOULD notify the operator.

4.4. Multicast Operation

4.4.1. Ingress Replication

The procedures for multicast operation on the PBB-VPLS PE using ingress replication are per [RFC7041] and [RFC7080].

The procedures for multicast operation on the PBB-EVPN PE for ingress replication are as follows:

- The PBB-EVPN PE builds a replication sub-list per I-SID to all the remote PBB-EVPN PEs in a given VPN instance as a result of the exchange of the EVPN IMET routes, as described in [RFC7623]. This will be referred to as sub-list A. It comprises MP2P service tunnels used for delivering PBB-EVPN BUM traffic.

- The PBB-EVPN PE builds a replication sub-list per VPN instance to all the remote PBB-VPLS PEs. This will be referred to as sub-list B. It comprises PWs from the PBB-EVPN PE in question to all the remote PBB-VPLS PEs in the same VPN instance.

- The PBB-EVPN PE may further prune sub-list B on a per-I-SID basis by running MMRP (see Clause 10 of [IEEE.802.1Q]) over the PBB-VPLS network. This will be referred to as sub-list C. This list comprises a pruned set of the PWs in sub-list B.

The replication list maintained per I-SID on a given PBB-EVPN PE will be the union of sub-list A and sub-list B if MMRP is not used and the union of sub-list A and sub-list C if MMRP is used. Note that the PE MUST enable split horizon over all the entries in the replication list, across both pseudowires and MP2P service tunnels.

4.4.2. P2MP Tunnel: Inclusive Tree

The procedures for multicast operation on the PBB-EVPN PEs using P2MP tunnels are outside of the scope of this document.

5. Security Considerations

All the security considerations in [RFC4761], [RFC4762], [RFC7080], [RFC7432], and [RFC7623] apply directly to this document because it leverages the control-plane and data-plane procedures described in those RFCs.

This document does not introduce any new security considerations beyond those of the above RFCs because the advertisements and processing of MAC addresses in BGP follow [RFC7432], and the processing of MAC addresses learned over PWs follows [RFC4761], [RFC4762], and [RFC7080].

6. IANA Considerations

This document has no IANA actions.

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC4761] Kompella, K., Ed. and Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", RFC 4761, DOI 10.17487/RFC4761, January 2007, <https://www.rfc-editor.org/info/rfc4761>.

[RFC4762] Lasserre, M., Ed. and V. Kompella, Ed., "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", RFC 4762, DOI 10.17487/RFC4762, January 2007, <https://www.rfc-editor.org/info/rfc4762>.

[RFC6074] Rosen, E., Davie, B., Radoaca, V., and W. Luo, "Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)", RFC 6074, DOI 10.17487/RFC6074, January 2011, <https://www.rfc-editor.org/info/rfc6074>.

[RFC7041] Balus, F., Ed., Sajassi, A., Ed., and N. Bitar, Ed., "Extensions to the Virtual Private LAN Service (VPLS) Provider Edge (PE) Model for Provider Backbone Bridging", RFC 7041, DOI 10.17487/RFC7041, November 2013, <https://www.rfc-editor.org/info/rfc7041>.

[RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <https://www.rfc-editor.org/info/rfc7432>.

[RFC7623] Sajassi, A., Ed., Salam, S., Bitar, N., Isaac, A., and W. Henderickx, "Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)", RFC 7623, DOI 10.17487/RFC7623, September 2015, <https://www.rfc-editor.org/info/rfc7623>.

[RFC8077] Martini, L., Ed. and G. Heron, Ed., "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", STD 84, RFC 8077, DOI 10.17487/RFC8077, February 2017, <https://www.rfc-editor.org/info/rfc8077>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

7.2. Informative References

[IEEE.802.1Q] IEEE, "IEEE Standard for Local and Metropolitan Area Network -- Bridges and Bridged Networks", IEEE Standard 802.1Q, DOI 10.1109/IEEESTD.2018.8403927, July 2018.

[RFC4684] Marques, P., Bonica, R., Fang, L., Martini, L., Raszuk, R., Patel, K., and J. Guichard, "Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)", RFC 4684, DOI 10.17487/RFC4684, November 2006, <https://www.rfc-editor.org/info/rfc4684>.

[RFC7080] Sajassi, A., Salam, S., Bitar, N., and F. Balus, "Virtual Private LAN Service (VPLS) Interoperability with Provider Backbone Bridges", RFC 7080, DOI 10.17487/RFC7080, December 2013, <https://www.rfc-editor.org/info/rfc7080>.

Authors' Addresses

Ali Sajassi (editor)
Cisco

Email: sajassi@cisco.com

Samer Salam
Cisco

Email: ssalam@cisco.com

Nick Del Regno
Verizon

Email: nick.delregno@verizon.com

Jorge Rabadan
Nokia

Email: jorge.rabadan@nokia.com