INTERNET-DRAFT
Intended Status: Standard Track

Ali Sajassi
Samer Salam
Cisco

Nick Del Regno
Verizon

Jorge Rabadan
Alcatel-Lucent

Expires: August 15, 2018

February 15, 2018

(PBB-)EVPN Seamless Integration with (PBB-)VPLS
draft-ietf-bess-evpn-vpls-seamless-integ-01

Abstract

This draft discusses the backward compatibility of the (PBB-)EVPN solution with (PBB-)VPLS and provides mechanisms for seamless integration of the two technologies in the same MPLS/IP network on a per-VPN-instance basis.

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/1id-abstracts.html

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

Copyright and License Notice

Table of Contents

1 Introduction 3
1.1 Terminology 3
2. Requirements 4
3 PBB-VPLS Integration with PBB-EVPN 4
3.1 Capability Discovery 4
3.2 Forwarding Setup and Unicast Operation 5
3.3 Multicast Operation 6
3.3.1 Ingress Replication 6
3.3.2 LSM 7
4 VPLS Integration with EVPN 7
4.1 Capability Discovery 7
4.2 Forwarding Setup and Unicast Operation 7
4.3 Multicast Operation 7
4.3.1 Ingress Replication 7
4.3.2 LSM 7
5 VPLS Integration with PBB-EVPN 7
5.1 Capability Discovery 7
5.2 Forwarding Setup and Unicast Operation 7
5.3 Multicast Operation 8
5.3.1 Ingress Replication 8
5.3.2 LSM 8
6 Solution Advantages 8
7 Security Considerations 8
8 IANA Considerations 8
9 References 8
9.1 Normative References 8
9.2 Informative References 9
Authors' Addresses 9

1 Introduction

VPLS and PBB-VPLS are widely-deployed L2VPN technologies. Many SPs who are looking at adopting EVPN and PBB-EVPN want to preserve their investment in the (PBB-)VPLS networks. Hence, it is required to provide mechanisms by which (PBB-)EVPN technology can be introduced into existing L2VPN networks without requiring a fork-lift upgrade. This document discusses mechanisms for the seamless integration of the two technologies in the same MPLS/IP network.

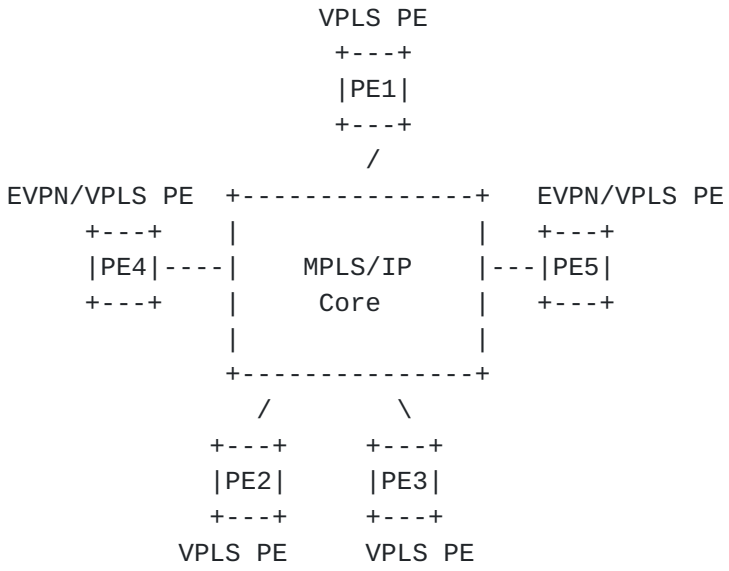

Figure 1: Seamless Integration of (PBB-)EVPN PEs & (PBB-)VPLS

Section 2 provides the details of the requirements. Section 3 discusses PBB-VPLS integration with PBB-EVPN. Section 4 discusses the integration of VPLS and EVPN. Section 5 discusses the integration of VPLS and PBB-EVPN, and finally Section 6 discusses the solution advantages.

It is worth noting that the scenario where PBB-VPLS is integrated with EVPN, is for future study and upon market validation. The reason for that is that deployments which employ PBB-VPLS typically require PBB encapsulation for various reasons. Hence, it is expected that for those deployments the evolution path would be from PBB-VPLS towards PBB-EVPN, rather than EVPN.

1.1 Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [KEYWORDS].

2. Requirements

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 PEs.

2. The solution MUST require no changes to existing VPLS or PBB-VPLS PEs, not even a software upgrade.

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

4. The solution MUST support single-active redundancy of multi-homed networks and multi-homed devices for (PBB-)EVPN PEs.

5. In case of single-active redundancy, the participant VPN instances MAY span across both (PBB-)EVPN PEs and (PBB-)VPLS PEs as long as single-active redundancy is employed by (PBB-)EVPN PEs. In case of an ES link failure, the (PBB-)EVPN PEs will send a BGP mass-withdraw to the EVPN peers OR MAC advertisement with MAC Mobility extended community for PBB-EVPN AND an LDP MAC withdrawal to the VPLS peers.

6. The solution SHOULD support all-active redundancy of multi-homed networks and multi-homed devices for (PBB-)EVPN PEs.

7. In case of all-active redundancy, the participant VPN instances SHOULD be confined to (PBB-)EVPN PEs only.

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

3 PBB-VPLS Integration with PBB-EVPN

In order to support seamless integration with (PBB-)VPLS, the (PBB-)EVPN PEs MUST support EVPN BGP routes (EVPN SAFI) and SHOULD support VPLS AD route (VPLS SAFI). All the logic for the integration will reside on the (PBB-)EVPN PEs side. However, if a VPLS instance is setup without the use of BGP auto-discovery, it is still possible (but cumbersome) for (PBB-)EVPN PEs to integrate into that VPLS instance.

3.1 Capability Discovery

The (PBB-)EVPN PEs must advertise both the BGP VPLS auto-discovery (AD) route as well as the BGP EVPN Inclusive Multicast route for a given VPN instance. The (PBB-)VPLS PEs only advertise the BGP VPLS AD route, per current standard procedures specified in [RFC4761] and [RFC6074]. The operator may decide to use the same BGP RT for both (PBB-)EVPN and (PBB-)VPLS. In this case, when a (PBB-)VPLS PE receives the EVPN Inclusive Multicast route, it will ignore it on the basis that it belongs to an unknown SAFI. However, the operator may use two RTs (one for (PBB-)VPLS and another for (PBB-)EVPN) and employ RT-constraint in order to prevent EVPN BGP routes from reaching the (PBB-)VPLS PEs. This provides an optimization in case required by the scale of the network.

When a (PBB-)EVPN PE receives both a VPLS AD route as well as an EVPN Inclusive Multicast 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 (PBB-)EVPN capable PEs discover other (PBB-)EVPN capable PEs as well as the (PBB-)VPLS-only PEs for that VPN instance. Furthermore, all the (PBB-)VPLS-only PEs would discover the (PBB-)EVPN PEs as if they were standard (PBB-)VPLS nodes. In other words, when the discovery phase is complete, the (PBB-)EVPN PEs would have discovered all the PEs in the VPN instance, and their associated capability: (PBB-)EVPN or VPLS-only. Whereas the (PBB-)VPLS PEs would have discovered all the PEs in the VPN instance, as if they were all VPLS-only nodes.

3.2 Forwarding Setup and Unicast Operation

The procedures for forwarding 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 pseudowire to a remote PE from which it has received only a VPLS AD route, for the VPN instance in question, and set up the label stack corresponding to the pseudowire FEC. This PW is between 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 (PBB-)VPN unicast FEC to any remote PE that has advertised EVPN AD route.

- If a (PBB-)EVPN PE receives a VPLS AD route followed by an EVPN AD route from the same PE and a pseudowire is setup to that PE, then the

(PBB-)EVPN MUST bring that pseudowire operationally down.

- If a (PBB-)EVPN PE receives an EVPN AD route followed by a VPLS AD route from the same PE, then the (PBB-)EVPN PE will setup the pseudowire but MUST keep it operationally down.

When the (PBB-)EVPN PE receives traffic over the pseudowires, it learns the associated MAC addresses in the data-plane. This is analogous to dynamic learning in IEEE bridges. If the PW belongs to the same split-horizon group as the EVPN mesh, then the MAC addresses learnt and associated to the PW will NOT be advertised in the control plane to any remote (PBB-)EVPN PE. The (PBB-)EVPN PE learns MAC addresses in the control plane, via the EVPN MAC Advertisement routes sent by remote (PBB-)EVPN PEs, and updates its MAC forwarding table accordingly. This is analogous to static learning in IEEE bridges. In PBB-EVPN, a given B-MAC address can be learnt either over the BGP control-plane from a remote PBB-EVPN PE, or in the data-plane over a pseudowire 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 there is an anomaly in the network.

3.3 Multicast Operation

3.3.1 Ingress Replication The procedures for multicast operation on the (PBB-)VPLS PE, using ingress replication, are per [RFC4761], [RFC4762], 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 Inclusive multicast routes, as described in [RFC7623]. This will be referred to as sub-list A. It comprises MP2P tunnels used for delivering PBB-EVPN BUM traffic [RFC7432].

- The PBB-EVPN PE builds a replication sub-list per VPN instance to all the remote PBB-VPLS PEs , as a result of the exchange of the VPLS AD routes. This will be referred to as sub-list B. It comprises pseudowires 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, if [MMRP] is run over the PBB-VPLS network. This will be referred to as sub-list C. This list comprises a pruned set of the pseudowires 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 tunnels.

3.3.2 LSM Will be covered in a future revision of this document.

4 VPLS Integration with EVPN

4.1 Capability Discovery

The procedures for capability discovery are per Section 3.1 above.

4.2 Forwarding Setup and Unicast Operation

The operation here is largely similar to that of PBB-EVPN integration with PBB-VPLS, with the exception of the need to handle MAC mobility, the details of which will be covered in a future revision of this document.

4.3 Multicast Operation

4.3.1 Ingress Replication

The operation is per the procedures of Section 3.3.1 above for the scenario WITHOUT [MMRP]. The replication list is maintained per VPN instance, rather than per I-SID.

4.3.2 LSM Will be covered in a future revision of this document.

5 VPLS Integration with PBB-EVPN

5.1 Capability Discovery

The procedures for capability discovery are per Section 3.1 above.

5.2 Forwarding Setup and Unicast Operation

The operation here is largely similar to that of PBB-EVPN integration with PBB-VPLS, with a few exceptions listed below:

- When a PW is setup between a PBB-EVPN PE and a VPLS PE, it gets setup between the I-component of PBB-EVPN PE and the bridge component of VPLS PE.

- The MAC mobility needs to be handled. The details of which will be

covered in a future revision of this document.

5.3 Multicast Operation

5.3.1 Ingress Replication

The operation is per the procedures of Section 3.3.1 above for the scenario WITHOUT [MMRP]. The replication list is maintained per I-SID on the PBB-EVPN PEs and per VPN instance on the VPLS PEs.

5.3.2 LSM Will be covered in a future revision of this document.

6 Solution Advantages

The solution for seamless integration of (PBB-)EVPN with (PBB-)VPLS has the following advantages:

- When ingress replication is used for multi-destination traffic delivery, the solution reduces the scope of [MMRP] (which is a soft-state protocol) 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 multi-homing mechanisms and provisioning simplifications of PBB-EVPN:
   1. Auto-sensing of MHN / MHD
   2. Auto-discovery of redundancy group
   3.Auto-provisioning of DF election and VLAN carving

7 Security Considerations

No new security considerations beyond those for VPLS and EVPN.

8 IANA Considerations

This document has no actions for IANA.

9 References

9.1 Normative References

[KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate

Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC8077] Martini, et al., "Pseudowire Setup and Maintenance using the Label Distribution Protocol", RFC 8077, February 2017.

[RFC7432] Sajassi et al., "BGP MPLS Based Ethernet VPN", RFC 7432, February, 2015.

[RFC7623] Sajassi et al., "Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)", RFC 7623, September, 2015.

[RFC4761] Kompella, K., Ed., and Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", RFC 4761, January 2007, <http://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, January 2007, <http://www.rfc-editor.org/info/rfc4762>.

[RFC6074] Rosen et al., "Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)", RFC 6074, January 2011.

9.2 Informative References

[MMRP] Clause 10 of "IEEE Standard for Local and metropolitan area networks - Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", IEEE Std 802.1Q, 2013.

[RFC7041] Balus et al., "Extensions to VPLS PE model for Provider Backbone Bridging", RFC 7041, November 2013.

[RFC7080] Sajassi et al., "VPLS Interoperability with Provider Backbone Bridges", RFC 7080, December, 2013.

Authors' Addresses

Ali Sajassi
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