

Internet Engineering Task Force (IETF)
Request for Comments: 7734
Category: Standards Track
ISSN: 2070-1721

D. Allan, Ed.
J. Tantsura
Ericsson
D. Fedyk
HPE
A. Sajassi
Cisco
January 2016

Support for Shortest Path Bridging MAC Mode over Ethernet VPN (EVPN)

Abstract

This document describes how Ethernet Shortest Path Bridging MAC mode (SPBM) can be combined with Ethernet VPN (EVPN) to interwork with Provider Backbone Bridging Provider Edges (PBB PEs) as described in the PBB-EVPN solution (RFC 7623). This is achieved via operational isolation of each Ethernet network attached to an EVPN core while supporting full interworking between the different variations of Ethernet networks.

Status of This Memo

This is an Internet Standards Track document.

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). Further information on Internet Standards is available in 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/rfc7734>.

Copyright Notice

Copyright (c) 2016 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
1.1. Requirements Language	3
2. Conventions Used in This Document	3
2.1. Terminology	3
3. Solution Overview	4
4. Elements of Procedure	5
4.1. PE Configuration	5
4.2. DF Election	6
4.3. Control-Plane Interworking ISIS-SPB to EVPN	6
4.4. Control-Plane Interworking EVPN to ISIS-SPB	7
4.5. Data-Plane Interworking SPBM Island or PBB PE to EVPN	8
4.6. Data-Plane Interworking EVPN to SPBM Island	8
4.7. Data-Plane Interworking EVPN to PBB PE	8
4.8. Multicast Support	8
5. Other Aspects	8
5.1. Transit	8
6. Security Considerations	9
7. References	10
7.1. Normative References	10
7.2. Informative References	10
Acknowledgments	11
Authors' Addresses	11

1. Introduction

This document describes how Ethernet Shortest Path Bridging MAC mode (SPBM) [IEEE.802.1Q] along with Provider Backbone Bridging Provider Edges (PBB PEs) and Provider Backbone Bridged Networks (PBBNs) can be supported by Ethernet VPNs (EVPNs) such that each SPBM island is operationally isolated while providing full L2 connectivity between the different types of PBBNs where desired. Each SPBM island uses its own control-plane instance and multipathing design, be it multiple equal-cost tree sets or multiple spanning trees.

The intention is to permit past, current, and emerging future versions of Ethernet to be seamlessly interconnected to permit large-scale, geographically diverse numbers of Ethernet end systems to be fully supported with EVPN as the unifying system.

1.1. Requirements Language

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 [RFC2119].

2. Conventions Used in This Document

2.1. Terminology

Terms used in this document are used as specified in IEEE 802.1Q-2014, which incorporates earlier IEEE 802.1 projects.

BEB: Backbone Edge Bridge

BGP: Border Gateway Protocol

B-MAC: Backbone MAC

B-VID: Backbone VLAN ID

CE: Customer Edge

DA: Destination Address

DF: Designated Forwarder

ESI: Ethernet Segment Identifier

EVPN: Ethernet VPN

IB-BEB: A BEB that has both an I-component (customer-layer VLAN-aware bridge) and a B-component (backbone-layer VLAN-aware bridge)

ISIS-SPB: IS-IS as extended for SPB

I-SID: Backbone Service Instance Identifier

NLRI: Network Layer Reachability Information

PBB: Provider Backbone Bridging as in Clauses 25 and 26 of [IEEE.802.1Q]

PBBN: Provider Backbone Bridged Network

PBB PE: Co-located BEB and EVPN PE

PE: Provider Edge

SPB: Shortest Path Bridging
 SPBM: Shortest Path Bridging MAC mode as in Clauses 27 and 28 of [IEEE.802.1Q]
 SPBM-PE: Co-located SPBM<->EVPN interworking function and EVPN PE

3. Solution Overview

The EVPN solution for SPBM, as specified in [IEEE.802.1Q], incorporates control-plane interworking in the PE to map ISIS-SPB [RFC6329] information elements into the EVPN Next Layer Reachability Information (NLRI) and vice versa. This requires each PE to act both as an EVPN BGP speaker and as an ISIS-SPB edge node. Associated with this are procedures for configuring the forwarding operations of the PE such that an arbitrary number of EVPN-attached SPBM islands can be interconnected without any topological or multipathing dependencies. This model also permits PBB PEs as defined in [RFC7623] to seamlessly communicate with the SPBM islands.




Figure 1: PBB and SPBM EVPN Network

Figure 1 illustrates the generalized space addressed by this memo. SPBM networks may be multihomed onto an IP/MPLS network that implements EVPN for the purpose of interconnecting with other SPBM networks and/or PBB PEs. The multipathing configuration of each SPBM network can be unique as the backbone VLAN ID (B-VID) configuration (how multipathing is performed in SPBM) is not propagated across the IP/MPLS network implementing EVPN. As with PBB networking, the B-VID is local to the SPBM network, so in SPBM a B-MAC associated with the B-VID is advertised with the supported I-SIDs at the PBB gateway.

Each EVPN is identified by a route target. I-SID-based load-balancing as specified in [RFC7623] allows multiple gateways per B-VID (each with different I-SIDs) across the EVPN; it is supported by the interworking between PBBNs and SPBM networks. However, SPBM

only allows a single active designated forwarder (DF) per B-VID as described below. The route target identifies the set of SPBM islands and PBB PEs that are allowed to communicate. Each SPBM island is administered to have an Ethernet Segment ID (ESI) Label extended community associated with it.

BGP acts as a common repository of the I-Component Service ID (I-SID) attachment points for the set of attached PEs / SPBM islands. This is in the form of {B-MAC address, I-SID, Tx-Rx-attribute} tuples. BGP distributes I-SID information into each SPBM island on the basis of locally registered interest. If an SPBM island has no Backbone Edge Bridges (BEBs) registering interest in a particular I-SID, information about that I-SID from other SPBM islands, PBB PEs, or PBBNs MUST NOT be leaked into the local ISIS-SPB routing system. For each B-VID in an SPBM island, a single SPBM-PE MUST be elected the DF for the B-VID. An SPBM-PE can be a DF for more than one B-VID. This is described further in Section 4.2. The SPBM-PE originates IS-IS advertisements as if it were an IB-BEB that proxies for the other SPBM islands and PBB PEs in the EVPN defined by the route target, but the PE typically will not actually host any I-components.

An SPBM-PE that is a DF for a B-VID MUST strip the B-VID tag information from frames relayed towards the EVPN. The DF MUST also insert the appropriate B-VID tag information into frames relayed towards the SPBM island on the basis of the local I-SID/B-VID bindings advertised in ISIS-SPB.

4. Elements of Procedure

A PE MUST implement and perform the following procedures.

4.1. PE Configuration

At SPBM island commissioning a PE is configured with:

- 1) The route target for the service instance. Where a route target is defined as identifying the set of SPBM islands, PBBNs and PBB PEs are to be interconnected by the EVPN.
- 2) The unique ESI for the SPBM island. Mechanisms for deriving a unique ESI for the SPBM island are out of scope.

The following is configured as part of commissioning an ISIS-SPB node:

- 1) A Shortest Path Source ID (SPSourceID) used for algorithmic construction of multicast addresses. Note this is required for SPBM BEB operation independent of the EVPN operation.

- 2) The set of B-VIDs used in the SPBM island and multipathing algorithm IDs to use for each. The set of B-VIDs and multipathing algorithms used can be different in different domains. Therefore, the B-VID is local to an SPBM domain and is removed for frames carried over the IP/MPLS network.

A Type 1 Route Distinguisher for the node can be auto-derived. The actual procedure is out of scope for this document.

4.2. DF Election

PEs self-appoint themselves for the role of DF for a B-VID for a given SPBM island. The procedure used is as per Section 8.5 (Designated Forwarder Election) of [RFC7432].

A PE that assumes the role of DF for a given B-VID is responsible for originating specific information into BGP from ISIS-SPB and vice versa. A PE that ceases to perform the role of DF for a given B-VID is responsible for withdrawing the associated information from BGP and ISIS-SPB, respectively. The actual information exchanged is outlined in the following sections.

4.3. Control-Plane Interworking ISIS-SPB to EVPN

When a PE receives an SPBM service identifier and unicast address sub-TLV as part of an ISIS-SPB MT capability TLV, the PE checks if it is the DF for the B-VID in the sub-TLV.

If it is the DF, and there is new or changed information, then a MAC/IP advertisement route NLRI is created for each new I-SID in the sub-TLV. Changed information that results in modification to existing NLRI is processed accordingly. NLRI creation/modification will ensure:

- the Route Distinguisher is set to that of the PE.
- the ESI is that of the SPBM island.
- the Ethernet Tag ID contains the I-SID (including the Tx/Rx attributes) copied from the SPBM service identifier and unicast address sub-TLV. The encoding of I-SID information is as per Figure 2. (See [RFC6329] for details on the T bit and R bit.)




Figure 2: I-SID Encoding in the Ethernet Tag ID Field

- the MAC address is copied from the SPBM service identifier and unicast address sub-TLV
- a locally assigned MPLS label (which may be common with other NLRI originated by the PE and is used to map EVPN traffic to the SPBM network)

Similarly, in the scenario where a PE became elected DF for a B-VID in an operating network, the IS-IS database would be processed in order to construct the NLRI associated with the new role of the PE.

If the BGP database has NLRI for the I-SID, and this is the first instance of registration of interest in the I-SID from the SPBM island, the NLRI for the I-SID is processed to construct an updated set of SPBM service identifier and unicast address sub-TLVs to be advertised by the PE.

The ISIS-SPB information is also used to keep current a local table indexed by I-SID to indicate the associated B-VID for processing of frames received from the EVPN. When an I-SID is associated with more than one B-VID, only one entry is allowed in the table. Rules for preventing this are out of scope for this memo.

4.4. Control-Plane Interworking EVPN to ISIS-SPB

When a PE receives a BGP NLRI that has new information, the PE checks if it is the elected DF to communicate this information into ISIS-SPB by checking if the I-SID in the Ethernet Tag ID locally maps to the B-VID for which it is an elected DF. Note that if no BEBs in the SPB island have advertised any interest in the I-SID, it will not be associated with any B-VID locally, and therefore will not be of interest. If the I-SID is of local interest to the SPBM island and the PE is the DF for the B-VID to which the I-SID is locally mapped, a SPBM service identifier and unicast address sub-TLV are constructed/updated for advertisement into ISIS-SPB.

The NLRI advertised into ISIS-SPB is also used to locally populate a forwarding table indexed by B-MAC + I-SID that points to the label stack to impose on the SPBM frame. The bottom label in the stack is that obtained from the NLRI.

4.5. Data-Plane Interworking SPBM Island or PBB PE to EVPN

When a PE receives a frame from the SPBM island in a B-VID for which it is a DF, it looks up the B-MAC/I-SID information to determine the label stack to be added to the frame for forwarding in the EVPN. The PE strips the B-VID information from the frame, adds the label information to the frame, and forwards the resulting MPLS packet.

4.6. Data-Plane Interworking EVPN to SPBM Island

When a PE receives a packet from the EVPN, it can infer the B-VID to overwrite in the SPBM frame from the I-SID or by other means (such as via the bottom label in the MPLS stack).

If the frame has a local multicast destination address (DA), it overwrites the SPSourceID in the frame with the local SPSourceID.

4.7. Data-Plane Interworking EVPN to PBB PE

A PBB PE actually has no attached PBBN nor concept of B-VID, so no frame processing is required.

A PBB PE is required to accept SPBM-encoded multicast DAs as if they were PBB-encoded (i.e., using the Backbone Service Instance Group address) for multicast DAs. The only information of interest is that it is a multicast frame and the I-SID encoded in the lower 24 bits.

4.8. Multicast Support

Within a PBBN domain, Ethernet unicast and multicast end services are supported. PBB can tunnel multicast traffic in unicast PBB frames when using head-end replication. This is the only form of multicast traffic interworking supported by this document. Native PBB multicast forwarding over EVPN, PE replication, or optimizing PBB multicast across the EVPN is not addressed by this memo.

5. Other Aspects

5.1. Transit

Any PE that does not need to participate in the tandem calculations at the B-MAC layer can use the IS-IS overload bit to exclude SPBM tandem paths and behave as a pure interworking platform.

6. Security Considerations

Security issues associated with incorrect interconnection of customer LANs cannot be directly addressed by implementations of this document, as it requires misconfiguration in the Shortest Path Bridging domains. The identifiers so administered have global significance to the larger system. They are relayed transparently by EVPN and only policed in the SPBM domains. Therefore, care is required in synchronization of identifiers that need to be common for inter-domain operation.

There are only two identifiers unique to this solution provisioned at an SPBM-PE at service turn-up: the route target and the ESI. The ESI needs to be unique and common to all SPBM-PEs connected to a common SPBM network or PBBN, else portions of the overall network will not share reachability. (The EVPN will assume that separate networks are interconnected by SPBM.) Security issues exist when SPBM domains are incorrectly cross-connected together via EVPN; this will result in black-holing or incorrect delivery of data with associated privacy issues. This error may occur by provisioning the incorrect RT value at an SPBM-PE or associating the RT with the wrong interface. This error can be avoided by consistency-checking the route target provisioning at SPBM-PEs when performing service additions and/or changes.

The behavior that is potentially most destructive to the overall system is frequent changes to the DF elections for a given ESI. This would occur if the SPBM-PEs continuously had their links go up and down in a such a way that the SPBM-PE was being severed from and reconnected to either the IP/MPLS network or the attached SPBM network. Either of these scenarios would result in significant control-plane traffic as DF associated information was advertised and withdrawn from both the SPBM and BGP control planes. Dual-homing of SPBM-PEs on both networks would minimize the likelihood of this scenario occurring.

The issues associated with securing the BGP control plane (independent of this particular memo) are reflected in the Security Considerations section of [RFC4761].

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, <<http://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, <<http://www.rfc-editor.org/info/rfc4761>>.
- [RFC6329] Fedyk, D., Ed., Ashwood-Smith, P., Ed., Allan, D., Bragg, A., and P. Unbehagen, "IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging", RFC 6329, DOI 10.17487/RFC6329, April 2012, <<http://www.rfc-editor.org/info/rfc6329>>.
- [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, <<http://www.rfc-editor.org/info/rfc7432>>.

7.2. Informative References

- [IEEE.802.1Q] IEEE, "IEEE Standard for Local and metropolitan area networks--Bridges and Bridged Networks", IEEE 802.1Q-2014, DOI 10.1109/ieeestd.2014.6991462, December 2014.
- [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, <<http://www.rfc-editor.org/info/rfc7623>>.

Acknowledgments

The authors would like to thank Peter Ashwood-Smith, Martin Julien, and Janos Farkas for their detailed reviews of this document.

Authors' Addresses

Dave Allan (editor)
Ericsson
300 Holger Way
San Jose, CA 95134
United States

Email: david.i.allan@ericsson.com

Jeff Tantsura
Ericsson
300 Holger Way
San Jose, CA 95134
United States

Email: jeff.tantsura@ericsson.com

Don Fedyk
Hewlett-Packard Enterprise
153 Taylor Street
Littleton, MA 01460
United States

Email: don.fedyk@hpe.com

Ali Sajassi
Cisco
170 West Tasman Drive
San Jose, CA 95134
United States

Email: sajassi@cisco.com