

L2VPN Working Group
Internet Draft
Intended status: Standards Track
Expires: January 2013

Dave Allan, Jeff Tantsura
Ericsson
Don Fedyk
Alcatel-Lucent
Ali Sajassi
Cisco

July 2012

802.1aq and 802.1Qbp Support over EVPN
draft-allan-l2vpn-spbm-evpn-01

Abstract

This document describes how Ethernet Shortest Path Bridging - MAC mode (802.1aq) and (802.1Qbp) can be combined with EVPN in a way that interworks with PBB-MESs as described in the PBB-EVPN solution in a way that permits operational isolation of each Ethernet network subtending an EVPN core while supporting full interworking between the 3 variations of Ethernet operation.

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/lid-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 January 2013.

Copyright and License 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.

Table of Contents

1. Introduction.....	2
1.1. Authors.....	3
1.2. 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. MES 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 802.1aq SPBM island or PBB-MES to EVPN.....	7
4.6. Data plane Interworking EVPN to 802.1aq SPBM island.....	8
4.7. Data plane interworking EVPN to 802.1ah PBB-MES.....	8
4.8. Dataplane interworking between 802.1Qbp islands and EVPN.....	8
4.9. Multicast Stitching.....	8
5. Other Aspects.....	8
5.1. Flow Ordering.....	8
5.2. Loop Avoidance and Black Holing...Error! Bookmark not defined.	
5.3. Transit.....	8
6. Acknowledgements.....	8
7. Security Considerations.....	9
8. IANA Considerations.....	9
8.1. Normative References.....	9
8.2. Informative References.....	9
9. Authors' Addresses.....	10

1. Introduction

This document describes how Ethernet Shortest Path Bridging - MAC mode (802.1aq) and (802.1Qbp) along with PBB-MESs and PBBNs (802.1ah)

can be supported by EVPN such that each island is operationally isolated while providing full L2 connectivity between them. Each island can use its own control plane instance and multi-pathing design, be it multiple ECT sets, multiple spanning trees, or ECMP.

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

1.1. Authors

David Allan, Jeff Tantsura, Don Fedyk, Ali Sajassi

1.2. 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 RFC2119 [1].

2. Conventions used in this document

2.1. Terminology

BCB: Backbone Core Bridge
BEB: Backbone Edge Bridge
BU: Broadcast/Unknown
B-MAC: Backbone MAC Address
B-VID: Backbone VLAN ID
CE: Customer Edge
C-MAC: Customer/Client MAC Address
DF: Designated Forwarder
ESI: Ethernet segment identifier
EVPN: Ethernet VPN
ISIS-SPB: IS-IS as extended for SPB
I-SID: I-Component Service ID
MES: MPLS Edge Switch
MP2MP: Multipoint to Multipoint
MVPN: Multicast VPN
NLRI: Network layer reachability information

PBBN: Provider Backbone Bridged Network
 PBB-MES: Co located BEB and MES
 P2MP: Point to Multipoint
 P2P: Point to Point
 RD: Route Distinguisher
 SPB: Shortest path bridging
 SPBM: Shortest path bridging MAC mode

3. Solution Overview

The EVPN solution for 802.1aq SPBM incorporates control plane interworking in the MES to map ISIS-SPB [2] information elements into the EVPN NLRI information and vice versa. This requires each MES 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 MES such that an arbitrary number of EVPN subtending SPB islands may be interconnected without any topological or multipathing dependencies. This requires each MES connected to an SPBM island to act both as an EVPN BGP speaker and as an ISIS-SPB edge node. This model also permits PBB-MESs as defined in draft-l2vpn-pbb-evpn-02[6] to be seamlessly communicate with the SPB islands. The next version of this document will add support for 802.1Qbp permitting seamless interworking between 802.1ah, 802.1aq and 802.1Qbp as well as supporting subtending 802.1ad based PBNS.

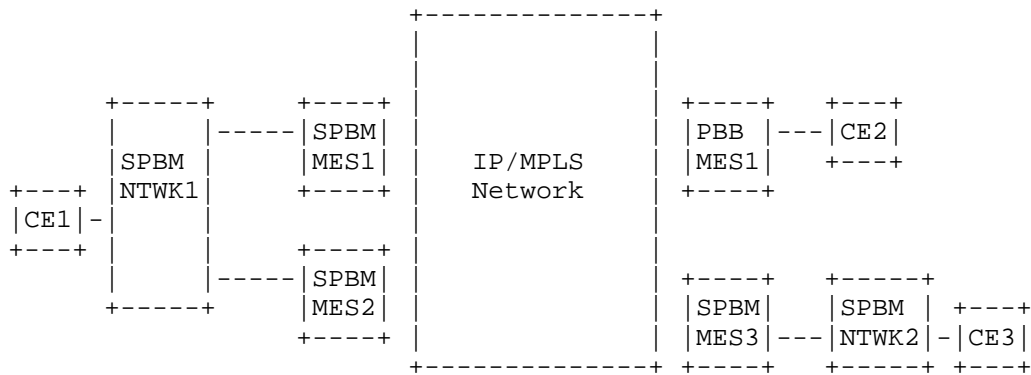


Figure 1: PBB and SPBM EVPN Network

Each EVPN is identified by a route target. The route target identifies the set of SPB islands and BEB-MESs that are allowed to communicate. This manifests itself as a set of Ethernet segments, where each Ethernet segment ID is unique within the route target.

BGP acts as a common repository of the I-SID attachment points for the set of subtending MESs/SPBM islands. This is in the form of B-MAC address/I-SID/Tx-Rx-attribute tuples. BGP filters leaking I-SID information into each SPBM ISLAND on the basis of locally registered interest. If an SPBM ISLAND has no BEBs registering interest in an I-SID, information about that I-SID from other SPBM island, PBB-MESs or PBBNs will not be leaked into the local ISIS-SPB routing system. Each SPBM island is administered to have an associated Ethernet Segment ID (ESI) associated with it.

For each B-VID in an SPBM island, a single SPBM-MES is elected the designated forwarder for the B-VID. An SPBM-MES may be a DF for more than one B-VID. This is described further in section 4.2. The SPBM-MES originates IS-IS advertisements as if it were an I-BEB or IB-BEB that proxy for the other SPBM islands and PBB MESs in the VPN defined by the route target, but the MES typically will not actually host any I-components.

An SPBM-MES that is a DF for a B-VID strips the B-VID tag information from frames relayed towards the EVPN. The DF also inserts 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

4.1. MES Configuration

At SPBM island commissioning a MES is configured with:

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

And 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 DA addresses. Note this is
- 2) The set of VLANs (identified by B-VIDs Ethernet frames) used in the SPBM island and multipathing algorithm IDs to use. The B-VID may be different in different domains and may be removed as carried over the IP/MPLS network.

A type-1 RD for the node can be auto-derived. This will be described in a future version of the document.

4.2. DF Election

MESs self appoint in the role of DF for a B-VID for a given SPBM island. The procedure used is as per section 9.5.2 of draft-ietf-l2vpn-evpn-01[4] "DF election with service carving".

4.3. Control plane interworking ISIS-SPB to EVPN

When a MES receives an SPBM service identifier and unicast address sub-TLV as part of an ISIS-SPB MT capability TLV it 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 advertisement route NLRI is created for each new I-SID in the sub-TLV.

- the Route Distinguisher (RD) is set to that of the MES
- the ESI is that of the SPBM island
- the Ethernet tag ID contains the I-SID (including the Tx/Rx attributes). The encoding of I-SID information is as per figure 2.

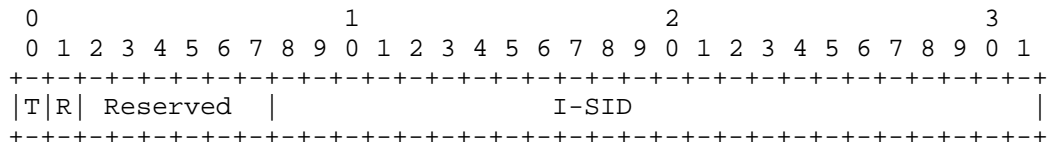


Figure 2: I-SID encoding in the Ethernet tag-ID field

- the MAC address from the sub-TLV

- an MPLS label

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

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

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 EVPN. When an I-SID is associated with more than one B-VID, only one entry is allowed in the table. Rules for this will be in a future version of the document.

4.4. Control plane interworking EVPN to ISIS-SPB

When a MES receives a BGP NLRI that is new information, it checks if the I-SID in the Ethernet Tag ID locally maps to the B-VID it is an elected DF for. 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 not of interest. If the I-SID is of local interest to the SPBM island and the MES is the DF for the B-VID that that I-SID is locally mapped to, a SPBM service identifier and unicast address sub-TLV is constructed/updated for advertisement into IS-IS.

The NLRI information 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 being that offered in the NLRI.

4.5. Data plane Interworking 802.1aq SPBM island or PBB-MES to EVPN

When an MES 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 MES 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 802.1aq SPBM island

When a MES receives a packet from the EVPN it may 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 DA, it overwrites the SPsourceID in the frame with the local SPsourceID.

4.7. Data plane interworking EVPN to 802.1ah PBB-MES

A PBB-MES actually has no subtending PBBN nor concept of B-VID so no frame processing is required.

A PBB-MES is required to accept SPBM encoded multicast DAs as if they were 802.1ah encoded multicast DAs. The only information of interest being that it is a multicast frame, and the I-SID encoded in the lower 24 bits.

4.8. Dataplane interworking between 802.1Qbp islands and EVPN

For a future version of the document

4.9. Multicast Stitching

For a future version of the document

5. Other Aspects

5.1. Flow Ordering

When per I-SID multicast is implemented via MES replication, a stable network will preserve frame ordering between known unicast and BU traffic (e.g. race conditions will not exist). This cannot be guaranteed when multicast is used in the EVPN.

5.2. Transit

Any MES that does not need to participate in the tandem calculations may use the IS-IS overload bit to exclude SPBM tandem paths and behave as pure interworking platform.

6. Acknowledgements

The authors would like to thank Peter Ashwood-Smith and Janos Farkas for their detailed review of this draft.

7. Security Considerations

For a future version of this document.

8. IANA Considerations

For a future version of this document.

8.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [2] Fedyk et.al. "IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging", IETF RFC 6329, April 2012
- [3] Rosen et.al., "BGP/MPLS IP Virtual Private Networks (VPNs)", IETF RFC 4364, February 2006
- [4] Aggarwal et.al. "BGP MPLS Based Ethernet VPN", IETF work in progress, draft-ietf-l2vpn-evpn-01, July 2012

8.2. Informative References

- [5] IEEE Standard for Local and Metropolitan Area Networks: Bridges and Virtual Bridged Local Area Networks - Amendment 9: Shortest Path Bridging
- [6] Draft IEEE Standard for Local and Metropolitan Area Networks---Virtual Bridged Local Area Networks - Amendment: Equal Cost Multiple Paths (ECMP), 802.1Qbp draft 1.0
- [7] Sajassi et.al. "PBB E-VPN", IETF work in progress, draft-ietf-l2vpn-pbb-evpn-03, June 2012
- [8] 802.1Q (2011) IEEE Standard for Local and metropolitan area networks--Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks

9. Authors' Addresses

Dave Allan (editor)
Ericsson
300 Holger Way
San Jose, CA 95134
USA
Email: david.i.allan@ericsson.com

Jeff Tantsura
Ericsson
300 Holger Way
San Jose, CA 95134
Email: jeff.tantsura@ericsson.com

Don Fedyk
Alcatel-Lucent
Groton, MA 01450
USA
Email: Donald.Fedyk@alcatel-lucent.com

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