

BESS Working Group
Internet Draft
Intended status: Standards Track
Expires: February 2016

Dave Allan, Jeff Tantsura
Ericsson
Don Fedyk
HP
Ali Sajassi
Cisco

July 2015

Shortest Path Bridging, MAC Mode Support over EVPN
draft-ietf-bess-spbm-evpn-00

Abstract

This document describes how Ethernet Shortest Path Bridging MAC mode (802.1aq) can be combined with EVPN in a way that interworks with PBB-PEs as described in the PBB-EVPN solution. This is achieved via operational isolation of each Ethernet network subtending an EVPN core while supporting full interworking between the different variations of Ethernet networks.

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 February 2016.

Copyright and License Notice

Copyright (c) 2015 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. 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. 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 802.1aq SPBM island or PBB-PE to EVPN	8
4.6. Data plane Interworking EVPN to 802.1aq SPBM island	8
4.7. Data plane interworking EVPN to 802.1ah PBB-PE	8
4.8. Multicast Support	8
5. Other Aspects	8
5.1. Flow Ordering	8
5.2. Transit	8
6. Acknowledgements	9
7. Security Considerations	9
8. IANA Considerations	9
9. References	9
9.1. Normative References	9
[RFC4761] Kompella (ed.), "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", IETF RFC 4761 , January 2007	
9.2. Informative References	10
10. Authors' Addresses	10

1. Introduction

This document describes how Ethernet Shortest Path Bridging MAC mode (802.1aq) along with PBB-PEs 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, or multiple spanning trees.

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](#) [[RFC2119](#)].

2. Conventions used in this document

2.1. Terminology

BEB: Backbone Edge Bridge
B-MAC: Backbone MAC Address
B-VID: Backbone VLAN ID
CE: Customer Edge
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: I-Component Service ID
NLRI: Network Layer Reachability Information
PBBN: Provider Backbone Bridged Network

Ethernet Segment ID (ESI) associated with it. This manifests itself as a set of Ethernet segments, where each ESI is unique within the route target.

BGP acts as a common repository of the I-SID attachment points for the set of subtending PEs/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 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 designated forwarder for the B-VID. An SPBM-PE may be a DF for more than one B-VID. This is described further in [section 5.2](#). 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 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 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.

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 addresses. Note this is required for SPBM BEBs independent of the EVPN operation.

- 2) The set of VLANs (identified by B-VIDs) used in the SPBM island and multi-pathing algorithm IDs to use. The set of B-VIDs and multi-pathing algorithms used may 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 of this document.

4.2. DF Election

PEs self appoint in the role of DF for a B-VID for a given SPBM island. The procedure used is as per [section 8.5 of \[RFC7432\]](#) "Designated Forwarder election".


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 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/IP advertisement route NLRI is created for each new I-SID in the sub-TLV. Changed information that results in modification to existing NLRI are processed accordingly.

- 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). 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 is copied from the sub-TLV
- a locally assigned MPLS label

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 information associated with the new role of the PE.

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 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 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 of this memo.

4.4. Control plane interworking EVPN to ISIS-SPB

When a PE receives a BGP NLRI that has new information, it 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 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 PE is the DF for the B-VID that the I-SID is locally mapped to, a SPBM service identifier and unicast address sub-TLV is constructed/updated for advertisement into ISIS-SPB.

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 in the stack being that obtained from the NLRI.

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

When an 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 802.1aq SPBM island

When a PE 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-PE

A PBB-PE actually has no subtending 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 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. Multicast Support

Not addressed by this memo.

5. Other Aspects

5.1. Flow Ordering

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

5.2. Transit

Any PE that does not need to participate in the tandem calculations at the B-MAC layer 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, Martin Julien and Janos Farkas for their detailed review of this draft.

7. Security Considerations

Security issues associated with incorrect interconnect 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.

Security issues exist when SPBM domains are incorrectly cross connected together via EVPN which will result in back-holing or incorrect delivery of data with associated privacy issues. This would be a consequence of conflicting administration of SPBM LAN identifiers.

Most of the issues associated with securing the BGP control plane are reflected in the security considerations section of [[RFC4761](#)].

8. IANA Considerations

No IANA assignments are required by this document.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4761] Kompella (ed.), "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", IETF [RFC 4761](#), January 2007
- [RFC6329] Fedyk et.al. "IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging", IETF [RFC 6329](#), April 2012
- [RFC7432] Aggarwal et.al. "BGP MPLS Based Ethernet VPN", IETF work in progress, IETF [RFC 7432](#), February 2015

9.2. Informative References

[802.1aq]

802.1aq(2012) IEEE Standard for Local and Metropolitan Area Networks: Bridges and Virtual Bridged Local Area Networks - Amendment 9: Shortest Path Bridging

[PBB-EVPN] Sajassi et.al. "PBB E-VPN", IETF work in progress, [draft-ietf-l2vpn-pbb-evpn-10](#), May 2015

[802.1Q]

802.1Q (2011) IEEE Standard for Local and metropolitan area networks--Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks

10. 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
Hewlett-Packard
153 Taylor Street
Littleton, MA, 01460
don.fedyk@hp.com

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

