

L2VPN Working Group
Internet Draft
Intended status: Standards Track
Expires: January 2015

Dave Allan, Jeff Tantsura
Ericsson
Sam Aldrin
Huawei

July 2014

**mLDP extensions for integrating EVPN and multicast
draft-allan-l2vpn-mlbp-evpn-03**

Abstract

This document describes how mLDP FECs can be encoded to support both service specific and shared multicast trees and describes the associated procedures for EVPN PEs. Thus, mLDP can implement multicast for EVPN.

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 July 2014.

Copyright and License Notice

Copyright (c) 2014 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.....2
- 1.2. Requirements Language.....3
- 2. Changes since last version.....3
- 3. Conventions used in this document.....3
- 3.1. Terminology.....3
- 4. Solution Overview.....4
- 5. Elements of Procedure.....4
- 6. FEC Encoding.....5
- 6.1. VLAN tagged FEC.....5
- 6.2. I-SID tagged FEC.....6
- 6.3. Shared FEC.....6
- 7. Acknowledgements.....7
- 8. Security Considerations.....7
- 9. IANA Considerations.....7
- 10. References.....7
- 10.1. Normative References.....7
- 10.2. Informative References.....8
- 11. Authors' Addresses.....8

1. Introduction

This document describes how mLDP FECs can be encoded to permit mLDP to implement multicast for EVPN. Such support can be applied to interconnecting 802.1ad, 802.1ah, 802.1aq, and 802.1Qbp based networks.

1.1. Authors

David Allan, Jeff Tantsura, Sam Aldrin

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. Changes since last version

- 1) Added co-author.

3. Conventions used in this document

3.1. 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
FEC: Forwarding Equivalence Class
ISIS-SPB: IS-IS as extended for SPB
I-SID: Backbone Service Instance ID
mLDP: Multicast Label Distribution Protocol
MP2MP: Multipoint to Multipoint
MVPN: Multicast VPN
NLRI: Network layer reachability information
PBBN: Provider Backbone Bridged Network
BEB-PE: Co located BEB and PE
PE: provider edge
P2MP: Point to Multipoint
P2P: Point to Point

RD: Route Distinguisher
SPB: Shortest path bridging
SPBM: Shortest path bridging MAC mode
VID: VLAN ID
VLAN: Virtual LAN

4. Solution Overview

mLDP[6] permits arbitrary FEC encodings for the naming of multicast trees to be defined. This property is leveraged to permit both service specific trees and shared trees to be utilized to augment EVPN unicast connectivity with network based multicast and avoid the inefficiencies of edge replication.

The flooding of EVPN BGP NLRI and ISIS-SPB [7] provides each PE with sufficient information to self elect as a DF, have knowledge of peer DFs, and from that construct the identifiers for the required set of multicast trees to support the current service set, which can then be encoded as mLDP FECs, and used to originate label mapping and label withdraw messages.

Both p2mp and mp2mp trees are supported with different FEC encodings for each. Service specific tree FECs encode the VID or I-SID associated with the service instance in the subtending network. Shared tree FECs encode a sorted list of the IP addresses of the leaf DFs.

5. Elements of Procedure

A PE advertises whether or not it supports shared tree (actual mechanism is TBD). Support of both shared and service specific trees is mandatory. Whether a PE supports shared trees is a network design decision.

A PE is expected to maintain a list of current multicast memberships.

A PE, upon receipt of new information from BGP or ISIS-SPB:

- 1) Evaluates it's DF roles (as described in [5]).
- 2) On the basis of the PE's DF role, determines the set of services it needs to support.
- 3) Determines the set of peer DFs for each service.

4) On the basis of requisite tree types and ESI multicast registrations (p2mp or mp2mp/service specific or shared), determines the name of the multicast tree needed for the service.

For example an ESI may only have source interest in an ISIS-SPB I-SID in which case it would:

- require a p2mp tree to the set of DFs registering receive interest in the I-SID for p2mp trees

- require an upstream label mapping to the set of DFs registering receive interest in the I_SID for mp2mp trees

5) Upon completion of evaluating the set of services, de-duplicates the required tree membership list.

6) Compares the required list with the existing list, and originates the necessary label mapping and label withdraw transactions to the network state up to date.


7) Configures the dataplane for the appropriate service to multicast tree bindings.

6. FEC Encoding

6.1. VLAN tagged FEC

VLAN tagged FEC uses the mLDP p2mp (0x06) type FEC and the mLDP mp2mp downstream (0x07) and upstream FECs (0x08)

The encoding of the opaque value is:




Where:

- RT is the route target for the EVPN instance
- Ethertype identifies the tag type (C 0x8100, S or B 0x88a8)
- VID is the VLAN ID tag value. If the VID=0, then this is the default MDT for the RT and how VLAN unaware RTs are encoded, else it permits MDTs to be defined for VLAN aware services.

6.2. I-SID tagged FEC

The encoding of the opaque value is:

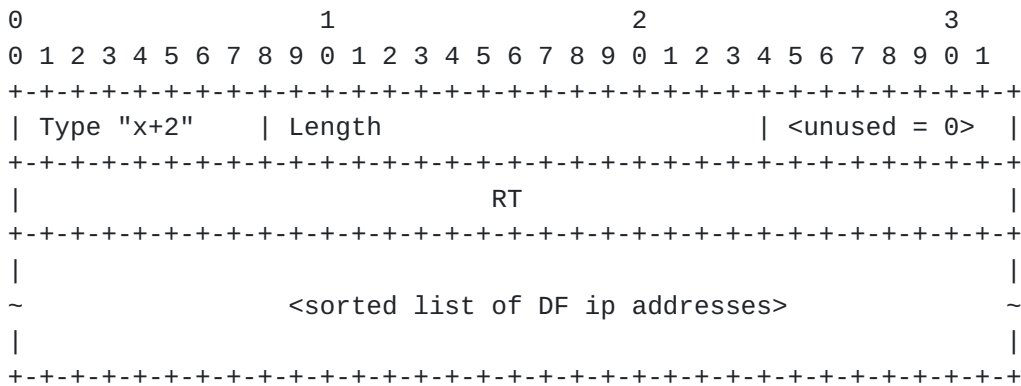


Where:

- RT is the route target for the EVPN instance
- I-SID corresponds to the I-SID that will use the tree

6.3. Shared FEC

The encoding of the opaque value is:



Where:

- RT is the route target for the EVPN instance
- Sorted list of DF addresses identifies the set of leaves that have registered interest in one or more Ethernet services (either C/S or I tagged).

7. Acknowledgements

The authors would like to thank Panagiotis Saltsidis, Jakob Heitz, Don Fedyk and Janos Farkas for their detailed review of this draft.

8. Security Considerations

For a future version of this document.

9. IANA Considerations

For a future version of this document.

10. References

10.1. 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
- [5] Allan et.al. "802.1aq and 802.1Qbp Support over EVPN", IETF work in progress, draft-allan-l2vpn-spb-evpn-03, February 2013
- [6] Wijnands et.al. "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths". IETF RFC 6388, November 2011

10.2. Informative References

- [7] IEEE 802.1aq "IEEE Standard for Local and Metropolitan Area Networks: Bridges and Virtual Bridged Local Area Networks - Amendment 9: Shortest Path Bridging", June 2012
- [8] IEEE 802.1Qbp "Draft IEEE Standard for Local and Metropolitan Area Networks---Virtual Bridged Local Area Networks - Amendment: Equal Cost Multiple Paths (ECMP), 802.1Qbp", draft 1.3, February 2013
- [9] Sajassi et.al. "PBB E-VPN", IETF work in progress, draft-ietf-l2vpn-pbb-evpn-03, June 2012
- [10] IEEE 802.1Q-2011 "IEEE Standard for Local and metropolitan area networks--Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", August 2011

11. 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

Sam Aldrin
Huawei Technologies
2330 Central Expressway
Santa Clara, CA 95051
EMail: aldrin.ietf@gmail.com