

BESS Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 13, 2022

H. Wang
Huawei Technologies
D. Eastlake
Futurewei Technologies
T. Zhu
Huawei Technologies
July 12, 2021

EVPN ELAN FRR Loop Prevent label
draft-wang-bess-evpn-frr-label-01

Abstract

This document describes how to use Fast Re-Route (FRR) labels avoid traffic loop in CE failures when deploying FRR protection in EVPN scenarios.

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](#)].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on January 13, 2022.

Copyright Notice

Copyright (c) 2021 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

(<https://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
2.	FRR Label Extended Community	3
3.	The Control Plane Process	4
4.	The Data Plane Process	4
5.	Other considerations	5
6.	IANA Considerations	5
7.	Security Considerations	5
8.	Acknowledgements	5
9.	Normative References	5
	Authors' Addresses	5

[1.](#) Introduction

In the EVPN active-active scenario, to solve the failure of a CE access channel to one PE, we can deploy Fast Re-Route (FRR) protection mode to achieve fast convergence. All active PEs can be deployed with FRR. When a link failure occurs on the CE connection to the PE, traffic can be rapidly FRR to another PE to improve the switching performance. However, if the CE device fails, both the two PEs sense that their CE link is faulty at the same time. They will each perform fast switching according to the FRR. Then the traffic will loop between the dual PEs.

If one PE detects a failure and withdraw the ES-AD route, the other PE, after receiving the withdrawal of the ES-AD route, deletes the FRR path to the PE, and the loop is eliminated. The time until the loop is eliminated may be short, but during this time, the loop will cause traffic congestion between the dual-homing PEs.

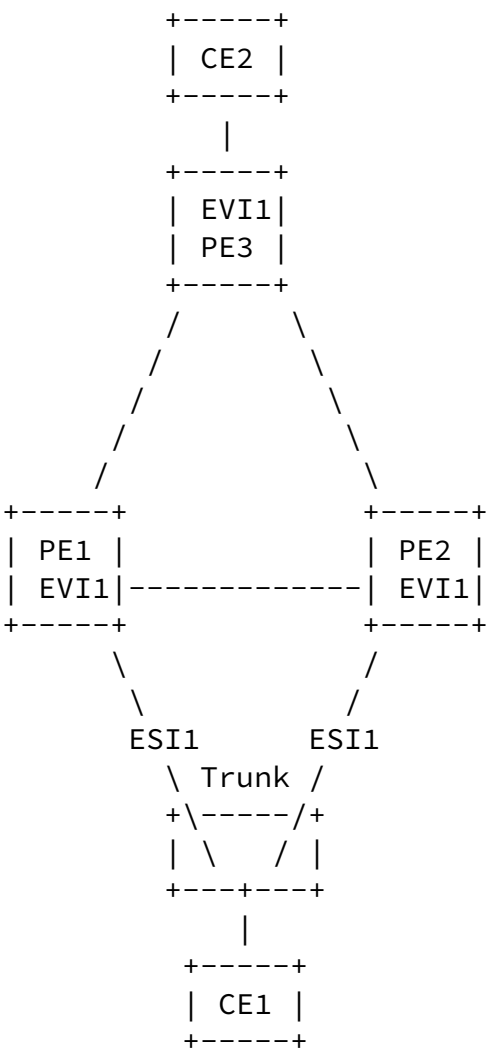


Figure 1: Basic networking of the EVPN all-active scenario

2. FRR Label Extended Community

The FRR Label Extended Community is a new transitive Extended Community having a Type field value of 0x06 and the Sub-Type TBD. It may be advertised along with MAC/IP Advertisement routes and Ethernet A-D per EVI routes.

The FRR Label Extended Community is encoded as an 8-octet value, as follows:

```

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Type=0x06      | Sub-Type=TBD  | Flags(1 octet)| Reserved=0    |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Reserved=0     |                FRR Label                |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

3. The Control Plane Process

If we enable the FRR LABEL function for an instance, then when the PE advertises the MAC-IP route or Ethernet A-D per EVI route of the instance, it also carries an FRR Label Extended Community attribute.

When another PE on the dual-homed side receives the MAC-IP route or the EVI-AD route, and uses the Label in the FRR Label Extended Community attribute as the label for the Edge FRR path. The single-homing-side PE receives the MAC-IP route or Ethernet A-D per EVI route advertised by the PE will ignores the FRR Label Extended Community attribute.

Because the FRR Label Extended Community attribute is an optional transitive attribute, if there are RR devices or ASBR devices in the network, the attributes can be transparently transmitted and processed by the final PE device.

Taking Figure 1 as an example, the EVI1 of PE2 enables the FRR LABEL function, and PE2 applies for a new label. PE2 advertises the MAC-IP route and the Ethernet A-D per EVI route carries the label through the FRR Label Extended Community attribute. Because CE1 is dual-homed to PE1 and PE2, PE1 learns the MAC address of CE1 from the data plane. Therefore, when PE1 receives the MAC-IP route or Ethernet A-D per EVI route from PE2, it can generate the MAC address learned from CE1 to form an edge FRR entry and the label filled in the FRR entry is the FRR label.

For PE3, even if CE1 is dual-homed to PE1 and PE2 in single-active

mode, PE3 form FRR does not use the FRR label.

The feature is available for EVPN ELAN service , EVPN VPWS and EVPN L3VPN service.

[4.](#) The Data Plane Process

The PE receives the traffic from the network side and finds the corresponding bridge-domain according to the Label. If the Label is a normal EVI label, the MAC address is normally queried. If the local outbound interface of the MAC fails, the FRR of the MAC is further protected. If the Label is an FRR label, the MAC address continues to be queried normally. If the local outbound interface of the MAC fails, the FRR of the MAC is no longer protected.

Taking Figure 1 as an example. CE2 send traffic to CE1, traffic arrived PE3 will encapsulate PE2 's EVI label and send to PE1. Traffic arrived at PE1 and use the EVI label to lookup MAC table, but find the out-interface is failed. PE1 will do MAC FRR for that and

encapsulate the FRR label, which is advertised by PE2. Traffic arrived at PE2 and use the FRR label to lookup MAC table. The traffic will send to CE1 according to the MAC information. When the MAC's out-interface is failed, the traffic will dropped to avoid traffic loop.

[5.](#) Other considerations

The solution of this document is not only applicable to the EVPN scenario. The traditional L3VPN can also use this solution to achieve rapid loop breaking.

[6.](#) IANA Considerations

IANA is requested to assign a new type of FRR label extended community with value TBD.

[7.](#) Security Considerations

TBD

[8.](#) Acknowledgements

9. 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, <<https://www.rfc-editor.org/info/rfc2119>>.

Authors' Addresses

Haibo Wang
Huawei Technologies
Huawei Campus, No. 156 Beiqing Road
Beijing 100095
China

Email: rainsword.wang@huawei.com

Wang, et al.

Expires January 13, 2022

[Page 5]

Internet-Draft

Abbreviated-Title

July 2021

Donald E. Eastlake
Futurewei Technologies
2386 Panoramic Circle
Apopka FL 32703
USA

Phone: +1-508-333-2270
Email: d3e3e3@gmail.com

Tong Zhu
Huawei Technologies
101 software Avenue, Yuhua District
Nanjing 210012
China

Email: zhu.tong@huawei.com