TRILL
Internet-Draft
Intended status: Standards Track
Updates: 7175, 7177

Expires: August 11, 2017

M. Zhang
Huawei Technologies
S. Pallagatti
Juniper Networks
V. Govindan
Cisco Systems

February 7, 2017

TRILL Support of Point to Multipoint BFD
draft-ietf-trill-p2mp-bfd-04

Abstract

Point to multipoint (P2MP) BFD is designed to verify multipoint connectivity. This document specifies the support of P2MP BFD in TRILL. Similar to TRILL point-to-point BFD, BFD Control packets in TRILL P2MP BFD are transmitted using RBridge Channel message. This document updates RFC 7175 and RFC 7177.

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 http://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 December 11, 2015.

Table of Contents

1. Introduction

TRILL supports multicast forwarding. Applications based on TRILL multicast may need quick detection of multicast failures using P2MP BFD. This document specifies TRILL support of P2MP BFD.

To use P2MP BFD, the head end needs to periodically transmit BFD Control packets to all tails using TRILL multicast. A new RBridge Channel message is allocated for this purpose.

In order to execute the global protection of distribution used for multicast forwarding [I-D.ietf-trill-resilient-trees], the head needs to track the active status of tails [I-D.ietf-bfd-multipoint-active-tail]. If the tail loses connectivity of the new RBridge Channel message from the head, the tail should notify the head of the lack of multipoint connectivity with unicast BFD Control packets. These unicast BFD Control packets are transmitted using the existing RBridge Channel message assigned to BFD Control [RFC7175].

This document updates [RFC7177] as specified in Section 3 and updates [RFC7175] as specified in Section 4.

2. Acronyms and Terminology

2.1. Acronyms

Data Label: VLAN or Fine Grained Label [RFC7172].

BFD: Bidirectional Forwarding Detection

P2MP: Point to Multi-Point

TRILL: Transparent Interconnection of Lots of Links or Tunneled Routing in the Link Layer

2.2. Terminology

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].

Familiarity with [RFC6325], [RFC7175], and [RFC7178] is assumed in this document.

3. Bootstrapping

The TRILL adjacency mechanism bootstraps the establishment of the BFD session [RFC7177]. A slight wording update to the second sentence in Section 6 of [RFC7177] is required.

It currently read:

> If an RBridge supports BFD [RFC7175], it will have learned whether the other RBridge has BFD enabled by whether or not a BFD-Enabled TLV [RFC6213] was included in its Hellos.

Now it should read:

> If an RBridge supports BFD [RFC7175] [this document], it will have learned whether the other RBridge has BFD enabled by whether or not a BFD-Enabled TLV [RFC6213] was included in its Hellos.

4. A New RBridge Channel Message for P2MP BFD

RBridge Channel message protocol 0x002 is defined for TRILL point-to-point BFD Control packets in [RFC7175]. If the M bit of the TRILL Header of the RBridge channel packet containing a BFD Control packet is non-zero, the packet MUST be dropped [RFC7175]. In P2MP BFD, the head is required to probe tails using multicast. This means the M bit will be set to 1. For this reason, a new RBridge Channel message, whose protocol code point is TBD, is specified in this document. An RBridge that supports P2MP BFD MUST support the new

RBridge Channel message for P2MP BFD. The capability to support the RBridge Channel message for P2MP BFD, and therefore support performing P2MP BFD, is announced within the "RBridge Channel Protocols Sub-TLV" in LSPs [RFC7176].

As specified in [RFC7178], when the tail receives TRILL Data packets sent as BFD RBridge channel messages, it will absorb the packets itself rather than deliver these packets to its attached end-stations.

5. Discriminators and Packet Demultiplexing

The processing in Section 3.2 of [RFC7175] applies except that the test on the M bit in the TRILL Header is reversed. If the M bit is zero, the packet is discarded. If the M bit is one, it is processed.

After the Section 3.2 of [RFC7175] processing, the tail demultiplexes incoming BFD packets based on a combination of the source address and My Discriminator as specified in [I-D.ietf-bfd-multipoint]. In addition to this combination, TRILL P2MP BFD that requires the tail use the Data Label, which is either the inner VLAN or the Fine Grained Label [RFC7172], for demultiplexing. If the tail needs to notify the head about the failure of a multipath, the tail is required to send unicast BFD Control packets using the same Data Label as used by the head.

6. Tracking Active Tails

According to[I-D.ietf-bfd-multipoint], the head has a session of type MultipointHead that is bound to a multipoint path. Multipoint BFD Control packets are sent by this session over the multipoint path, and no BFD Control packets are received by it. Each tail dynamically creates a MultipointTail per a multipoint path. MultipointTail sessions receive BFD Control packets from the head over multipoint paths.

If the head is keeping track of some or all of the tails [I-D.ietf-trill-resilient-trees], it has a session of type MultipointClient per tail that it cares about [I-D.ietf-bfd-multipoint-active-tail]. See [I-D.ietf-bfd-multipoint-active-tail] for detail operations of tacking active tails.

7. Security Considerations

P2MP BFD control packets can be encapsulated as the payload of the RBridge Channel Tunnel [RFC7978]. In that case, the security option of RBridge Channel Tunnel can secure the transmission of BFD control

packets.

The demultiplexing of TRILL P2MP BFD at the tail is Data Label aware. This enhances the security of the dynamic creation of MultipointTail sessions at tails. In order to forge BFD Control packets, the attacker has to acquire the right Data Label that the head uses for P2MP BFD.

For general multipoint BFD security considerations, see [I-D.ietf-bfd-multipoint].

For general RBridge Channel security considerations, see [RFC7178].

8. IANA Considerations

IANA is required to allocate one RBridge Channel protocol number from the Standards Action range, as follows:

Protocol	Number
P2MP BFD Control	TBD

9. Acknowledgements

Authors would like to thank the comments and suggestions from Gayle Noble and Donald Eastlake.

10. References

10.1. Normative References

[I-D.ietf-bfd-multipoint]
Katz, D., Ward, D., and J. Networks, "BFD for Multipoint Networks", draft-ietf-bfd-multipoint (work in progress).

[I-D.ietf-bfd-multipoint-active-tail]
Katz, D., Ward, D., and J. Networks, "BFD Multipoint Active Tails.", draft-ietf-bfd-multipoint-active-tail (work in progress).

[I-D.ietf-trill-resilient-trees]
Zhang, M., Senevirathne, T., Pathangi, J., Banerjee, A., and A. Ghanwani, "TRILL Resilient Distribution Trees", draft-ietf-trill-resilient-trees (work in progress).

[RFC7978] Eastlake 3rd, D., Umair, M., and Y. Li, "Transparent Interconnection of Lots of Links (TRILL): RBridge Channel Header Extension", RFC 7978, DOI 10.17487/RFC7978,

September 2016, <http://www.rfc-editor.org/info/rfc7978>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC6325] Perlman, R., Eastlake, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification", RFC 6325, July 2011.

[RFC7172] Eastlake, D., Zhang, M., Agarwal, P., Perlman, R., and D. Dutt, "Transparent Interconnection of Lots of Links (TRILL): Fine-Grained Labeling", RFC 7172, May 2014.

[RFC7175] Manral, V., Eastlake, D., Ward, D., and A. Banerjee, "Transparent Interconnection of Lots of Links (TRILL): Bidirectional Forwarding Detection (BFD) Support", RFC 7175, May 2014.

[RFC7176] Eastlake, D., Senevirathne, T., Ghanwani, A., Dutt, D., and A. Banerjee, "Transparent Interconnection of Lots of Links (TRILL) Use of IS-IS", RFC 7176, May 2014.

[RFC7177] Eastlake, D., Perlman, R., Ghanwani, A., Yang, H., and V. Manral, "Transparent Interconnection of Lots of Links (TRILL): Adjacency", RFC 7177, May 2014.

[RFC7178] Eastlake, D., Manral, V., Li, Y., Aldrin, S., and D. Ward, "Transparent Interconnection of Lots of Links (TRILL): RBridge Channel Support", RFC 7178, May 2014.

10.2. Informative References

[RFC6213] Hopps, C. and L. Ginsberg, "IS-IS BFD-Enabled TLV", RFC 6213, April 2011.

Authors' Addresses

Mingui Zhang
Huawei Technologies
No.156 Beiqing Rd. Haidian District
Beijing 100095
P.R. China

Email: zhangmingui@huawei.com

Santosh Pallagatti
Juniper Networks
Embassy Business Park
Bangalore KA 560093
India

Email: santoshpk@juniper.net

Vengada Prasad Govindan
Cisco Systems

Email: venggovi@cisco.com