

TRILL
Internet-Draft
Intended status: Standards Track
Expires: August 7, 2015

M. Zhang
Huawei Technologies
S. Pallagatti
Juniper Networks
V. Govindan
Cisco Systems
February 03, 2015

Point to Multipoint BFD for TRILL
draft-zhang-trill-p2mp-bfd-00

Abstract

Point to multipoint (P2MP) BFD is designed to verify multipoint connectivity. This document specifies the support of P2MP BFD in TRILL. Similar as TRILL point to point BFD, BFD Control packets in TRILL P2MP BFD are also transmitted using an extended RBridge Channel.

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 August 7, 2015.

Copyright 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	2
2. Acronyms and Terminology	2
2.1. Acronyms	2
2.2. Terminology	3
3. Bootstrapping	3
4. A New RBridge Channel for P2MP BFD	3
5. Discriminators and Packet Demultiplexing	4
6. Tracking Active Tails	4
7. Security Considerations	4
8. IANA Considerations	5
9. References	5
9.1. Normative References	5
9.2. Informative References	6
Authors' Addresses	6

1. Introduction

TRILL supports multicast forwarding. Applications based on TRILL multicast wish to achieve quick detection of multicast failures using P2MP BFD. This document specifies the use of P2MP BFD in TRILL.

To use P2MP BFD, the head need to periodically transmit BFD Control packets to all tails using TRILL multicast. A new RBridge Channel 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 need to track the active status of tails [spallagatti-bfd-multipoint-active-tail]. When the tail loses connectivity from the head, it should notify the head of the lack of multipoint connectivity with unicast BFD Control packets. These packets are transmitted using the existing RBridge Channel assigned to BFD Control [RFC7175].

2. Acronyms and Terminology

2.1. Acronyms

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

BFD: Bidirectional Forwarding Detection

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][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 for P2MP BFD

RBridge Channel 0x002 is defined for TRILL point to point BFD Control packets in [RFC7175]. If the M bit of the TRILL Header of the channeled packet containing the BFD Control packet is non-zero, the packet MUST be dropped [RFC7175]. While for 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, whose code point is TBD, is specified in this document. An RBridge that supports P2MP BFD MUST support the new RBridge Channel for P2MP BFD. The capability to support the RBridge Channel for P2MP BFD, and therefore support performing P2MP BFD, is announced within the "RBridge Channel Protocols Sub-TLV" in LSPs [RFC7176].

An alternative option is to define a new RBridge Channel Tunnel protocol for P2MP BFD Control packets [I-D.ietf-trill-channel-tunnel] so that P2MP BFD Control Packets can be adapted as the payload of this Tunnel protocol.

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

5. Discriminators and Packet Demultiplexing

In [I-D.ietf-bfd-multipoint], the tail demultiplexes incoming BFD packets based on a combination of the source address and My Discriminator. In addition to this combination, TRILL P2MP BFD requires the tail to use the Data Label, which is either the inner VLAN or the Fine Grained Label [RFC7172], for demultiplexing. If the tail need 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 [spallagatti-bfd-multipoint-active-tail]. See [spallagatti-bfd-multipoint-active-tail] for detail operations of tracking active tails.

7. Security Considerations

P2MP BFD control packets can be encapsulated as the payload of the RBridge Channel Tunnel [I-D.ietf-trill-channel-tunnel]. 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.

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. References

9.1. Normative References

- [I-D.ietf-bfd-multipoint]
Katz, D., Ward, D., and J. Networks, "BFD for Multipoint Networks", draft-ietf-bfd-multipoint-05 (work in progress), January 2015.
- [I-D.ietf-trill-channel-tunnel]
Eastlake, D. and L. Yizhou, "TRILL: RBridge Channel Tunnel Protocol", draft-ietf-trill-channel-tunnel-02 (work in progress), December 2014.
- [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-02 (work in progress), December 2014.
- [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.
- [spallagatti-bfd-multipoint-active-tail]
Katz, D., Ward, D., and S. Pallagatti , "BFD Multipoint Active Tails", January 2015.

9.2. 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