

Internet Engineering Task Force (IETF)
Request for Comments: 8564
Updates: 7175, 7177
Category: Standards Track
ISSN: 2070-1721

M. Zhang
Huawei Technologies
S. Pallagatti
Vmware
V. Govindan
Cisco Systems
April 2019

Support of Point-to-Multipoint Bidirectional Forwarding Detection (BFD)
in Transparent Interconnection of Lots of Links (TRILL)

Abstract

Point-to-multipoint (P2MP) Bidirectional Forwarding Detection (BFD) is designed to verify multipoint connectivity. This document specifies the support of P2MP BFD in Transparent Interconnection of Lots of Links (TRILL). Similar to TRILL point-to-point BFD, BFD Control packets in TRILL P2MP BFD are transmitted using RBridge Channel messages. This document updates RFCs 7175 and 7177.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8564>.

Copyright Notice

Copyright (c) 2019 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. Acronyms and Terminology 3
 - 2.1. Acronyms 3
 - 2.2. Terminology 3
- 3. Bootstrapping 3
- 4. A New RBridge Channel Message for P2MP BFD 4
- 5. Discriminators and Packet Demultiplexing 4
- 6. Tracking Active Tails 5
- 7. Security Considerations 5
- 8. IANA Considerations 5
- 9. References 6
 - 9.1. Normative References 6
 - 9.2. Informative References 7
- Acknowledgements 7
- Authors' Addresses 8

1. Introduction

TRILL supports multicast forwarding. Applications based on TRILL multicast may need quick detection of multicast failures using P2MP BFD [RFC8562]. 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 [TRILL-TREES], the head needs to track the active status of tails [RFC8563]. If the tail loses connectivity as detected by not receiving 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 Sections 4 and 5.

2. Acronyms and Terminology

2.1. Acronyms

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

BFD: Bidirectional Forwarding Detection

P2MP: Point to Multipoint

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", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

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

3. Bootstrapping

The TRILL adjacency mechanism bootstraps the establishment of one-hop TRILL BFD sessions [RFC7177]. Multi-hop sessions are expected to be configured by the network manager. A slight wording update to the second sentence in Section 6 of [RFC7177] is required.

It currently reads:

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 (see [RFC7175] and [RFC8564]), 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.

In addition, a normative reference to this document is added to RFC 7177 as a result of this update.

4. A New RBridge Channel Message for P2MP BFD

RBridge Channel protocol number 0x002 is defined for TRILL point-to-point BFD Control packets in [RFC7175]. That RFC states that if the M bit of the TRILL Header of the RBridge Channel packet containing a BFD Control packet is nonzero, the packet is generally dropped. 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 (P2MP BFD Control), whose protocol code point is 0x007, 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 Link State PDUs (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] generally applies except that the test on the M bit in the TRILL Header is reversed. If the M bit is zero, the packet MUST be discarded. If the M bit is one, it is processed.

After the processing per Section 3.2 of [RFC7175], the tail demultiplexes incoming BFD packets based on a combination of the source address and My Discriminator as specified in [RFC8562]. In addition to this combination, TRILL P2MP BFD requires that 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 [RFC8562], 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.

An example use is when a multicast tree root needs to keep track of all the receivers as in [TRILL-TREES]; this can be done by maintaining a session of type MultipointClient per receiver that is of interest, as described in [RFC8563]. See [RFC8563] for detailed operations of tracking active tails.

7. Security Considerations

Multipoint BFD provides its own authentication but does not provide encryption (see the Security Considerations in [RFC8562]). As specified in this document, the point-to-multipoint BFD payloads are encapsulated in RBridge Channel messages that have been extended by [RFC7978] to provide security. [RFC7978] provides encryption only for point-to-point extended RBridge Channel messages, so its encryption facilities are not applicable to this document. However, [RFC7978] provides stronger authentication than that currently provided in BFD. Thus, there is little reason to use the BFD security mechanisms if authentication per [RFC7978] is in use. It is expected that a future TRILL document will provide for group keying; when that occurs, the use of RBridge Channel security [RFC7978] will be able to provide both encryption and authentication.

For general multipoint BFD security considerations, see [RFC8562].

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

8. IANA Considerations

IANA has allocated the following from the Standards Action [RFC8126] range of the "RBridge Channel Protocols" registry, which is part of the "Transparent Interconnection of Lots of Links (TRILL) Parameters" registry.

Protocol	Number
-----	-----
P2MP BFD Control	0x007

9. References

9.1. 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>>.
- [RFC6325] Perlman, R., Eastlake 3rd, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification", RFC 6325, DOI 10.17487/RFC6325, July 2011, <<https://www.rfc-editor.org/info/rfc6325>>.
- [RFC7172] Eastlake 3rd, D., Zhang, M., Agarwal, P., Perlman, R., and D. Dutt, "Transparent Interconnection of Lots of Links (TRILL): Fine-Grained Labeling", RFC 7172, DOI 10.17487/RFC7172, May 2014, <<https://www.rfc-editor.org/info/rfc7172>>.
- [RFC7175] Manral, V., Eastlake 3rd, D., Ward, D., and A. Banerjee, "Transparent Interconnection of Lots of Links (TRILL): Bidirectional Forwarding Detection (BFD) Support", RFC 7175, DOI 10.17487/RFC7175, May 2014, <<https://www.rfc-editor.org/info/rfc7175>>.
- [RFC7176] Eastlake 3rd, D., Senevirathne, T., Ghanwani, A., Dutt, D., and A. Banerjee, "Transparent Interconnection of Lots of Links (TRILL) Use of IS-IS", RFC 7176, DOI 10.17487/RFC7176, May 2014, <<https://www.rfc-editor.org/info/rfc7176>>.
- [RFC7177] Eastlake 3rd, D., Perlman, R., Ghanwani, A., Yang, H., and V. Manral, "Transparent Interconnection of Lots of Links (TRILL): Adjacency", RFC 7177, DOI 10.17487/RFC7177, May 2014, <<https://www.rfc-editor.org/info/rfc7177>>.
- [RFC7178] Eastlake 3rd, D., Manral, V., Li, Y., Aldrin, S., and D. Ward, "Transparent Interconnection of Lots of Links (TRILL): RBridge Channel Support", RFC 7178, DOI 10.17487/RFC7178, May 2014, <<https://www.rfc-editor.org/info/rfc7178>>.
- [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, <<https://www.rfc-editor.org/info/rfc7978>>.

- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8562] Katz, D., Ward, D., Pallagatti, S., Ed., and G. Mirsky, Ed., "Bidirectional Forwarding Detection (BFD) for Multipoint Networks", RFC 8562, DOI 10.17487/RFC8562, April 2019, <<https://www.rfc-editor.org/info/rfc8562>>.
- [RFC8563] Katz, D., Ward, D., Pallagatti, S., Ed., and G. Mirsky, Ed., "Bidirectional Forwarding Detection (BFD) Multipoint Active Tails", RFC 8563, DOI 10.17487/RFC8563, April 2019, <<https://www.rfc-editor.org/info/rfc8563>>.

9.2. Informative References

- [RFC6213] Hopps, C. and L. Ginsberg, "IS-IS BFD-Enabled TLV", RFC 6213, DOI 10.17487/RFC6213, April 2011, <<https://www.rfc-editor.org/info/rfc6213>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [TRILL-TREES]
Zhang, M., Senevirathne, T., Pathangi, J., Banerjee, A., and A. Ghanwani, "TRILL: Resilient Distribution Trees", Work in Progress, draft-ietf-trill-resilient-trees-09, January 2018.

Acknowledgements

The authors would like to thank Gayle Noble and Donald Eastlake 3rd for their comments and suggestions.

Authors' Addresses

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

Email: zhangmingui@huawei.com

Santosh Pallagatti
Vmware

Email: santosh.pallagatti@gmail.com

Vengada Prasad Govindan
Cisco Systems

Email: venggovi@cisco.com