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

V. Manral
Ionos Corp.
D. Eastlake 3rd
Huawei R&D USA
D. Ward
Cisco Systems
A. Banerjee
Cumulus Networks

May 2014

Transparent Interconnection of Lots of Links (TRILL): Bidirectional Forwarding Detection (BFD) Support

Abstract

This document specifies use of the Bidirectional Forwarding Detection (BFD) protocol in Routing Bridge (RBridge) campuses based on the RBridge Channel extension to the Transparent Interconnection of Lots of Links (TRILL) protocol.

BFD is a widely deployed Operations, Administration, and Maintenance (OAM) mechanism in IP and MPLS networks, using UDP and Associated Channel Header (ACH) encapsulation respectively. This document specifies the BFD encapsulation over TRILL.

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 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7175.

Copyright 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

Faster convergence is a critical feature of Transparent Interconnection of Lots of Links (TRILL) [RFC6325] networks. The TRILL IS-IS Hellos [RFC7177] [IS-IS] used between RBridges provide a basic neighbor and continuity check for TRILL links. However, failure detection by non-receipt of such Hellos is based on the Holding Time parameter that is commonly set to a value of tens of seconds and, in any case, has a minimum expressible value of one second.

Some applications, including Voice over IP, may wish, with high probability, to detect interruptions in continuity within a much shorter time period. In some cases, physical-layer failures can be detected very rapidly, but this is not always possible, such as when there is a failure between two bridges that are in turn between two RBridges. There are also many subtle failures possible at higher levels. For example, some forms of failure could affect unicast frames while still letting multicast frames through; since all TRILL IS-IS Hellos are multicast, such a failure cannot be detected with Hellos. Thus, a low-overhead method for frequently testing continuity for the TRILL Data between neighbor RBridges is necessary for some applications. The BFD protocol [RFC5880] provides a low-overhead method for the rapid detection of connectivity failures.

BFD is a widely deployed OAM [RFC6291] mechanism in IP and MPLS networks, using UDP and ACH encapsulation, respectively. This document describes a TRILL encapsulation for BFD packets for networks that forward based on the TRILL Header.

1.1. Terminology

This document uses the acronyms defined in [RFC6325] along with the following:

BFD: Bidirectional Forwarding Detection

IP: Internet Protocol

IS-IS: Intermediate System to Intermediate System

MH: Multi-Hop

PPP: Point-to-Point Protocol

OAM: Operations, Administration, and Maintenance

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].

2. BFD over TRILL

TRILL supports unicast neighbor BFD Echo and one-hop and multi-hop BFD Control, as specified below, over the RBridge Channel facility [RFC7178]. (Multi-destination BFD is a work in progress [MultiBFD].) BFD-over-TRILL support is similar to BFD-over-IP support [RFC5881], except where differences are explicitly mentioned.

Asynchronous and demand modes MUST be supported [RFC5880]. BFD over TRILL supports the Echo function; however, implementation of TRILL BFD Echo is optional, and it can only be used for single-hop sessions.

The TRILL Header hop count in the BFD packets is sent out with the maximum value of 0x3F. To prevent spoofing attacks, the TRILL hop count of a received session is checked [RFC5082]. For a single-hop session, if the hop count is less than 0x3F and the RBridge Channel Header MH flag is zero, the packet is discarded. For multi-hop sessions, the hop count check can be disabled if the MH flag is one.

As in BFD for IP, the format of the Echo Packet content is not defined.

New RBridge Channel code points for BFD TRILL Control and BFD Echo packets are specified.

Authentication mechanisms as supported in BFD are also supported for BFD running over TRILL.

2.1. Sessions and Initialization

Within an RBridge campus, there will be no more than one TRILL BFD Control session from any RBridge RB1 to RBridge RB2 for each RB1 TRILL port. This BFD session must be bound to this interface. As such, both sides of a session MUST take the "Active" role (sending initial BFD Control packets with a zero value of Your Discriminator), and any BFD packet from the remote machine with a zero value of Your Discriminator MUST be associated with the session bound to the remote system and interface.

Note that TRILL BFD provides OAM facilities for the TRILL data plane. This is above whatever protocol is in use on a particular link, such as a pseudowire [RFC7173], Ethernet [RFC6325], or PPP link [RFC6361]. Link-technology-specific OAM protocols may be used on a link between neighbor RBridges, for example, Continuity Fault Management [802.1Q] if the link is Ethernet. But such link-layer OAM (and coordination between it and OAM in the TRILL data-plane layer, such as TRILL BFD) is beyond the scope of this document.

If lower-level mechanisms are in use, such as link aggregation [802.1AX], that present a single logical interface to TRILL IS-IS, then only a single TRILL BFD session can be established to any other RBridge over this logical interface. However, lower-layer OAM could be aware of and/or run separately on each of the components of an aggregation.

3. TRILL BFD Control Protocol

TRILL BFD Control frames are unicast TRILL RBridge Channel frames [RFC7178]. The RBridge Channel Protocol value is given in Section 8. The protocol-specific data associated with the TRILL BFD Control protocol is as shown in Section 4.1 of [RFC5880].

3.1. One-Hop TRILL BFD Control

One-hop TRILL BFD Control is typically used to rapidly detect link and RBridge failures. TRILL BFD frames over one hop for such purposes SHOULD be sent with high priority; that is, the Inner.VLAN tag priority should be 7, they should be queued for transmission as maximum priority frames, and, if they are being sent on an Ethernet link where the output port is configured to include an Outer.VLAN tag, that tag should specify priority 7.

For neighbor RBridges RB1 and RB2, each RBridge sends one-hop TRILL BFD Control frames to the other only if TRILL IS-IS has detected bidirectional connectivity; that is, the adjacency is in the 2-Way or Report state [RFC7177], and both RBridges indicate support of TRILL BFD is enabled. The BFD-Enabled TLV is used to indicate this as specified in [RFC6213].

3.2. BFD Control Frame Processing

The following tests SHOULD be performed on received TRILL BFD Control frames before generic BFD processing.

- Is the M bit in the TRILL Header non-zero? If so, discard the frame. (Multi-destination BFD is a work in progress [MultiBFD].) Failure to perform this test would make a denial-of-service attack using bogus multi-destination BFD Control frames easier.

- If the Channel Header MH flag is zero, indicating one hop, test that the TRILL Header hop count received was 0x3F (i.e., is 0x3E if it has already been decremented); if it is any other value, discard the frame. If the Channel Header MH flag is one, indicating multi-hop, test that the TRILL Header hop count received was not less than a configurable value that defaults to 0x30. If it is less, discard the frame. Failure to perform these tests would make it easier to spoof BFD Control frames. However, if forged BFD Control frames are a concern, then BFD Authentication [RFC5880] should be used.

4. TRILL BFD Echo Protocol

A TRILL BFD Echo frame is a unicast RBridge Channel frame, as specified in [RFC7178], which should be forwarded back by an immediate neighbor because both the ingress and egress nicknames are set to a nickname of the originating RBridge. Normal TRILL Data frame forwarding will cause the frame to be returned unless micro-loop suppression logic in the neighbor RBridge prohibits sending a frame back out the port on which it was received or the like. RBridges with such prohibitions cannot support BFD Echo. The TRILL OAM protocol number for BFD Echo is given in Section 8.

TRILL BFD Echo frames SHOULD be sent on a link only if the following conditions are met. An Echo originating under other circumstances will consume bandwidth and CPU resources but is unlikely to be returned.

- A TRILL BFD Control session has been established,
- TRILL BFD Echo support is indicated by the RBridge that would potentially respond to the BFD Echo,
- The adjacency is in the Report state [RFC7177], and
- The TRILL BFD Echo originating RBridge wishes to make use of this optional feature.

Since the originating RBridge is the RBridge that will be processing a returned Echo frame, the entire TRILL BFD Echo protocol-specific data area is considered opaque and left to the discretion of the originating RBridge. Nevertheless, it is suggested that this data include information by which the originating RBridge can authenticate the returned BFD Echo frame and confirm the neighbor that echoed the frame back. For example, it could include its own System ID, the neighbor’s System ID, a session identifier, and a sequence count as well as a Message Authentication Code.

4.1. BFD Echo Frame Processing

The following tests MUST be performed on returned TRILL BFD Echo frames before other processing. The RBridge Channel document [RFC7178] requires that the information in the TRILL Header be given to the BFD protocol.

- o Is the M bit in the TRILL Header non-zero? If so, discard the frame. (Multi-destination BFD is a work in progress [MultiBFD].)

- o The TRILL BFD Echo frame should have gone exactly two hops, so test that the TRILL Header hop count as received was 0x3E (i.e., 0x3D if it has already been decremented), and if it is any other value, discard the frame. The RBridge Channel Header in the frame MUST have the MH bit equal to one, and if it is zero, discard the frame.

5. Management and Operations Considerations

The TRILL BFD parameters on an RBridge are configurable. The default values are the same as in the IP BFD case [RFC5881], except where specified in this document, such as for hop count.

It is up to the operator of an RBridge campus to configure the rates at which TRILL BFD frames are transmitted on a link to avoid congestion (e.g., link, input/output (I/O), CPU) and false failure detection. See also the discussion of congestion in Section 2 of [RFC5881].

As stated in [RFC5880]:

> It is worth noting that a single BFD session does not consume a large amount of bandwidth. An aggressive session that achieves a detection time of 50 milliseconds, by using a transmit interval of 16.7 milliseconds and a detect multiplier of 3, will generate 60 packets per second. The maximum length of each packet on the wire is on the order of 100 bytes, for a total of around 48 kilobits per second of bandwidth consumption in each direction.

6. Default Authentication

Consistent with TRILL's goal of being able to operate with minimum configuration, the default for BFD authentication between neighbor RBridges is based on the state of the IS-IS shared secret authentication for Hellos between those RBridges as detailed below. The BFD authentication algorithm and methods in this section MUST be implemented at an RBridge if TRILL IS-IS authentication and BFD are implemented at that RBridge. If such BFD authentication is configured, then its configuration is not restricted by the configuration of IS-IS security.

If IS-IS authentication is not in effect between neighbor RBridges, then, by default, TRILL BFD between those RBridges is also unsecured.

If such IS-IS authentication is in effect, then, unless configured otherwise, TRILL BFD Control frames sent between those RBridges MUST use BFD Meticulous Keyed SHA1 authentication [RFC5880]. The BFD authentication keys between neighbor RBridges by default are derived

from the IS-IS shared secret authentication keys for Hellos between those RBridges as detailed below. However, such BFD authentication keys MAY be configured to some other value.

```
HMAC-SHA256 ( ( "TRILL BFD Control" | originPortID | originSysID ),
              IS-IS-shared-key )
```

In the above, "|" indicates concatenation; HMAC-SHA256 is as described in [FIPS180] and [RFC6234]; and "TRILL BFD Control" is the 17-byte US ASCII [ASCII] string indicated that is then concatenated with the 2-byte Port ID of the originating port and the 6-byte IS-IS System ID of the originating RBridge, the last two items being in network byte order. The Port and System IDs are included to minimize exposure of the same key to improve resistance to cryptanalysis. IS-IS-shared-key is secret keying material being used for IS-IS authentication on the link.

The use of the above derived key is accomplished by associating the above default authentication type and key with the Key ID of the IS-IS-shared-key used in the derivation and then using that Key ID in the Authentication Section of the BFD Control frame OAM protocol-specific data. Also, Auth Type would be 5, and Auth Len would be 28 in the Authentication Section. RBridges MAY be configured to use other BFD security modes or keying material or configured to use no security.

Authentication for TRILL BFD Echo is a local implementation issue as BFD Echo frames are authenticated by their sender when returned by a neighbor. However, if TRILL IS-IS and BFD Control are being authenticated to a neighbor and BFD Echo is in use, BFD Echo frames to be returned by that neighbor should be authenticated, and such authentication should use different keying material from other types of authentication. For example, it could use keying material derived as follows, where "|" indicates concatenation:

```
HMAC-SHA256 ( ( "TRILL BFD Echo" | originPortID | originSysID ),
              IS-IS-shared-key )
```

7. Security Considerations

BFD over TRILL utilizes the RBridge Channel extension to the TRILL protocol and is generally analogous to BFD over IP. As such, the BFD authentication facility is available to authenticate BFD-over-TRILL packet payloads, but no encryption or other security features are provided at the BFD-over-TRILL level. See the following:

- [RFC5881] for general BFD security considerations,

- [RFC7178] for general RBridge Channel security considerations, and

- [RFC6325] for general TRILL protocol security considerations.

Section 3.2 describes security concerns with multi-hop BFD Control packets and failure to check the TRILL Header M bit in BFD Control packets.

8. IANA Considerations

IANA has allocated two RBridge Channel protocol numbers [RFC7178] from the Standards Action range, as follows:

Protocol	Number
BFD Control	0x002
BFD Echo	0x003

9. Acknowledgements

The authors would like to specially thank Dave Katz, an author of [RFC5880] and [RFC5881], from which some material herein has been reproduced.

The following individuals are thanked for their comments and suggestions: Scott Bradner, Stewart Bryant, Stephen Farrell, Eric Gray, Brian Haberman, Barry Leiba, Erik Nordmark, John Scudder, Robert Sparks, Martin Stiemerling, and Sean Turner.

10. References

10.1. Normative References

[ASCII] American National Standards Institute, "Coded Character Set - 7-bit American Standard Code for Information Interchange", ANSI X3.4, 1986.

[FIPS180] National Institute of Science and Technology, "Secure Hash Standard (SHS)", Federal Information Processing Standard (FIPS) 180-4, March 2012, <http://csrc.nist.gov/publications/fips/fips180-4/fips-180-4.pdf>.

[IS-IS] International Organization for Standardization, "Intermediate System to Intermediate System intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)", Second Edition, November 2002.

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

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, June 2010.

[RFC5881] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, June 2010.

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

[RFC6325] Perlman, R., Eastlake 3rd, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification", RFC 6325, July 2011.

[RFC7177] Eastlake 3rd, D., Perlman, R., Ghanwani, A., Yang, H., and V. Manral, "Transparent Interconnection of Lots of Links (TRILL): Adjacency", RFC 7177, May 2014.

[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, May 2014.

10.2. Informative References

[802.1AX] IEEE, "IEEE Standard for Local and metropolitan area networks -- Link Aggregation", IEEE Std 802.1AX-2008, January 2008.

[802.1Q] IEEE, "IEEE Standard for Local and metropolitan area networks -- Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", IEEE Std 802.1Q-2011, August 2011.

[MultiBFD] Katz, D. and D. Ward, "BFD for Multipoint Networks", Work in Progress, February 2014.

[RFC5082] Gill, V., Heasley, J., Meyer, D., Savola, P., Ed., and C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)", RFC 5082, October 2007.

[RFC6234] Eastlake 3rd, D. and T. Hansen, "US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF)", RFC 6234, May 2011.

[RFC6291] Andersson, L., van Helvoort, H., Bonica, R., Romascanu, D., and S. Mansfield, "Guidelines for the Use of the "OAM" Acronym in the IETF", BCP 161, RFC 6291, June 2011.

[RFC6361] Carlson, J. and D. Eastlake 3rd, "PPP Transparent Interconnection of Lots of Links (TRILL) Protocol Control Protocol", RFC 6361, August 2011.

[RFC7173] Yong, L., Eastlake 3rd, D., Aldrin, S., and J. Hudson, "Transparent Interconnection of Lots of Links (TRILL) Transport Using Pseudowires", RFC 7173, May 2014.

Authors' Addresses

Vishwas Manral
Ionos Corp.
4100 Moorpark Ave.
San Jose, CA 95117
USA

EMail: vishwas@ionosnetworks.com

Donald Eastlake 3rd
Huawei R&D USA
155 Beaver Street
Milford, MA 01757
USA

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

Dave Ward
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95138
USA

EMail: dward@cisco.com

Ayan Banerjee
Cumulus Networks
1089 West Evelyn Avenue
Sunnyvale, CA 94086
USA

EMail: ayabaner@gmail.com