TRILL Working Group
INTERNET-DRAFT
Intended status: Proposed Standard
Updates: RFCtrill

Donald Eastlake 3rd
Stellar Switches
Vishwas Manral
IP Infusion
Dave Ward
Juniper
Ayan Banerjee
Cisco

Expires: April 16, 2011
October 17, 2010

RBridges: OAM and BFD Support for TRILL
<draft-eastlake-trill-rbridge-bfd-00.txt>

Abstract

This document specifies a general channel for sending OAM (Operations, Administration, and Maintenance) messages between RBridges in a campus through an extension to the TRILL (Transparent Interconnection of Lots of Links) protocol. It further specifies use of this channel for the BFD (Bidirectional Forwarding Detection) protocol.

Status of This Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Distribution of this document is unlimited. Comments should be sent to the TRILL working group mailing list.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/1id-abstracts.html

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html

Table of Contents

1. Introduction

The TRILL IS-IS Hellos used between RBridges provide a basic neighbor and continuity check for TRILL links [RFCtrill]. However, failure detection by non-receipt of such Hellos is based on the holding time parameter which is typically set to a value over ten seconds and, in any case, has a minimum expressible value of one second.

Many applications, including voice over IP, may wish, with very 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 devices that are in turn between two RBridges, and there are many subtle failures possible at higher levels. For example, some forms of failure could affect unicast frames while still letting multicast frames through and all TRILL IS-IS frames, including Hellos, are multicast. Thus, a method of frequently testing continuity for the TRILL Data between neighbor RBridges is necessary for some applications.

Such continuity testing is one example of TRILL data plane Operations, Administration, and Maintenance (OAM) requirements. Various of such requirements can be met by a variety of protocols such as the Bidirectional Forwarding Detection (BFD) [RFC5880] [RFC5882] and [Y.1731].

This document specifies, in Section 2, a general channel for sending OAM messages between RBridges in a campus using extensions to the TRILL protocol and further specifies, in Section 3, use of this channel for the BFD protocol. TRILL BFD can be used to provide rapid detection of link continuity failure for TRILL Data frames.

1.1 Terminology

The terminology and acronyms of [RFCtrill] are used in this document.

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].

1.2 Additional Acronyms

The following acronyms are used in this document in addition to those defined in [RFCtrill]:

BFD - Bidirectional Forwarding Detection

MH - Multi-Hop

OAM - Operations, Administration, and Maintenance

OV - OAM (Message Channel) Version

SL - Silent

1.3 Acknowledgements

The authors would like to particularly thank David Katz, co-author of [RFC5880] and [RFC5882]. Some of the text is this document was adapted from those RFCs.

2. The TRILL OAM Message Channel

TRILL OAM messages are transmitted as TRILL Data frames. They are primarily identified as OAM messages by their Inner.MacDA and Inner Ethertype. This Inner Ethertype is followed by a 32-bit TRILL OAM Header used to indicate the OAM protocol of the following OAM protocol specific data. A TRILL Header bit option is provided that may optionally be used to guarantee that frames sent over the TRILL OAM Message Channel cannot accidentally be forwarded to end stations, even by RBridges that are ignorant of the TRILL OAM Message Channel mechanism.

The diagram below shows the overall structure of a TRILL OAM Message Channel frame on a link between two RBridges:

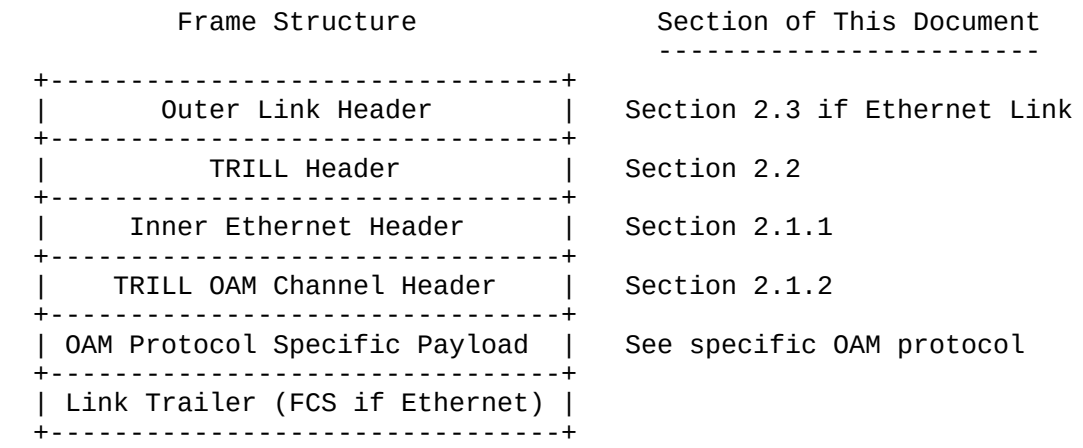

The Sections 2.1 and 2.2 below describe the Inner frame and TRILL Header for frames sent in the TRILL OAM Message Channel. As always, the Outer link header is whatever is needed to get a TRILL Data frame from one RBridge to the next, depends on link technology, and can change with each hop for multi-hop OAM messages. Section 2.3 describes the Outer link header for Ethernet. Section 2.4 goes into further detail on the OAM-Channel Bit Option. And Section 2.5 describes some details of TRILL OAM Message processing.

2.1 The OAM Message Inner Frame

The encapsulated Inner frame within A TRILL OAM Message Channel frame is as shown below.

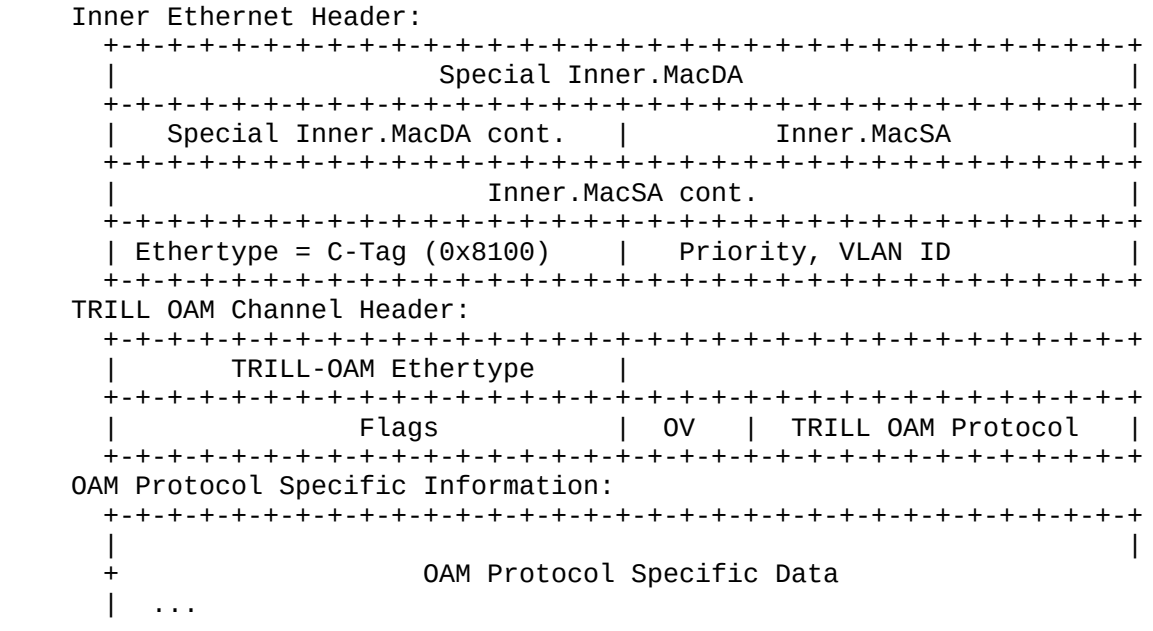

2.1.1 Inner Ethernet Header

The special Inner.MacDA is one of two values: OAM-RBridge-MAC if the OAM message is unicast or All-OAM-RBridges if the OAM message is multi-destination (see Section 6).

The Inner.MacSA is selected by the RBridge originating the OAM message. If it is a unicast MAC address, on decapsulation it will be learned as being attached to the ingress RBridge. If that learning is not desired, the Inner.MacSA MAY be set to All-OAM-RBridges. Address learning on decapsulation does not occur if the source MAC has the group bit on.

As with all TRILL encapsulated frames, a VLAN tag MUST be present. Use of a VLAN tag Ethertype other than 0x8100 is beyond the scope of this document. Recommendations for the frame priority are as follows:

- For one-hop known unicast OAM messages critical to network connectivity, such as one-hop BFD for rapid link failure detection in support of TRILL IS-IS, the RECOMMENDED priority is 7.

- For multi-hop known unicast OAM messages, the RECOMMENDED priority is 6.

- For multi-destination OAM messages, it is RECOMMENDED that the priority be no higher than 5.

Multi-destination TRILL OAM messages are VLAN scoped so the Inner.VLAN ID MUST be set to the VLAN of interest. To the extent that distribution tree pruning is in effect, such OAM messages will only reach RBridges advertising that they have appointed forwarder connectivity to that VLAN.

For known unicast OAM messages, if the message is one-hop it is RECOMMENDED that the Inner.VLAN ID be the Designated VLAN on that hop. For multi-hop unicast OAM messages, it is RECOMMENDED that the Inner.VLAN ID be the default VLAN 1.

2.1.2 TRILL OAM Header

After the TRILL OAM Ethertype (see Section 6) is a four-byte quantity with three sub-fields. The first, Flags, provides 16 bits of flags which, except as specified below, MUST be sent as zero, transparently copied by transit RBridges, and ignored on receipt. The next field, OV, gives the OAM Header version and MUST be zero. Lastly, a 12-bit field specified the particular TRILL OAM protocol to which the message applies. See Section 6 for IANA Considerations.

The flag bits are numbered from 0 to 15 as shown below.

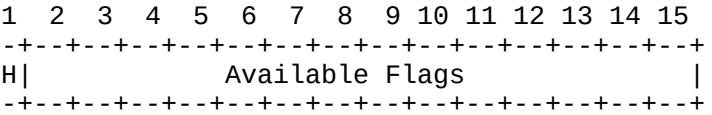

Bit 0, which is the high order bit in network order, is defined as the SL or Silent bit. If it is a one, it suppresses OAM Channel Error messages due to the use of an unknown version or OAM protocol (see Section 2.5.1). Bit 1 is the MH or Multi-Hop bit. It is used to inform the destination OAM protocol that the message was intended to be multi-hop (MH=1) or one-hop (MH=0).

The TRILL OAM Protocol field specifies the OAM protocol that the OAM Channel message relates to. Initial defined values are as listed below. See Section 6 for IANA Considerations.

Protocol	Name - Section of this Document
0x0001	OAM Channel Error - Section 2.5
0x0002	TRILL BFD Control - Section 3.2
0x0003	TRILL BFD Echo - Section 3.3

2.2 The TRILL Header for OAM Messages

After the Outer link header (which for Ethernet ends with the TRILL Ethertype) and before the encapsulated frame, the OAM message's TRILL Header appears as follows:

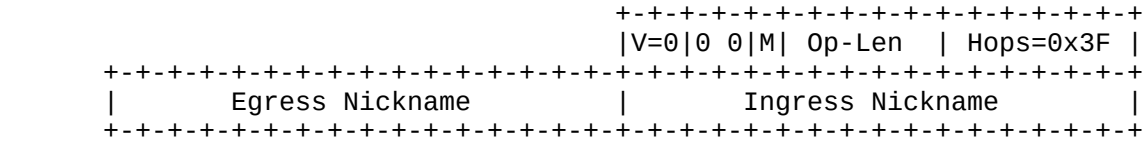

The TRILL Header version V, MUST be zero, the M bit is set appropriately as the OAM message is known unicast (M=0) or multi-destination (M=1), and Op-Len is set appropriately for the length of the options area, if any, all as specified in [RFCtrill].

When a TRILL OAM message is originated, the hop count field is always set to the maximum value, 0x3F. For messages sent a known number of hops, particularly one-hop messages or neighbor echo messages, checking the Hops (Hop Count) field provides an additional validity check as discussed in [RFC5082].

The RBridge originating a TRILL OAM message places a nickname that it holds into the ingress nickname field.

There are several cases for the egress nickname field. If the OAM message is multi-destination, then the egress nickname designates the distribution tree to use. If the OAM message is a multi-hop unicast message, then the egress nickname is a nickname of the target RBridge; this includes the special case of an "echo" OAM message where the originator places its own nickname in both the ingress and egress nickname fields. If the OAM message is a one-hop unicast message, there are two possibilities for the egress nickname.

- The egress nickname can bet set to a nickname of the target neighbor RBridge. This will usually work well but there is a small chance that, due to a nickname transient, the frame will actually be delivered to some other RBridge in the campus. Due to this possibility, both here and in the multi-hop unicast case, if a TRILL OAM message is intended for a specific RBridge in the campus topology, it is RECOMMENDED that the OAM protocol specific data include the IS-IS SystemID of the target RBridge for an added check.

- The special nickname Any-RBridge may be used. This will guarantee decapsulation at the immediate neighbor RBridge regardless of the state of nickname assignments. RBridges supporting the TRILL OAM Channel facility MUST recognize the Any-RBridge special nickname and accept TRILL Data frames having that value in the egress

nickname field as being sent to them as the egress.

2.3 OAM Message Ethernet Link Header

If the link on which a TRILL OAM frame is transmitted between neighbor RBridges is Ethernet, the link header follows the usual rules for a TRILL Data frame over Ethernet [RFCtrill]. In particular, the Outer.MacSA is the MAC address of the port from which the frame is sent. The Outer.MacDA is the MAC address of the next-hop RBridge port for unicast TRILL OAM messages or the All-RBridges multicast address for multi-destination TRILL OAM messages. If an Outer.VLAN tag is present, it must specify the Designated VLAN for that hop and the priority must be the same as in the Inner.VLAN tag.

2.4 The TRILL OAM-Channel Bit Option

A critical ingress-to-egress TRILL Header bit option, OAM-Channel, is specified associated with the TRILL OAM Channel facility. This option is NOT REQUIRED to appear in the TRILL Header in TRILL OAM message frames. It serves two functions, as follows:

- An RBridge indicates that it supports the TRILL OAM Channel facility by advertising, in the link state database, its support for this bit option.

- If this bit option is present in a TRILL OAM message frame, it guarantees that, if the inner frame is decapsulated by an RBridge that does not implement the TRILL OAM Channel it will be discarded rather than being locally flooded as a native frame out all ports for which that RBridge is appointed forwarder for the Inner.VLAN. However, if it is certain that all RBridges in the campus implement the TRILL OAM Channel or if the possible local flooding of the inner frame as specified above is acceptable, there is NO REQUIREMENT to include an options area or to set this particular option bit in the TRILL Header options area even if an options area is included.

As with any other critical ingress-to-egress option, if the bit options area is present and this bit option is set, then the summary CItE bit MUST be set at the top of the options area.

2.5 Processing TRILL OAM Messages

TRILL OAM messages are designed to look like and, to the extent practical, be processed as regular TRILL Data frames. On receiving a TRILL OAM frame, the initial tests on the Outer.MacDA, Outer Ethertype, TRILL Header V and Hop Count fields and the RPF check if the frame is multi-destination, are all performed as usual. The forwarding and/or decapsulation decisions are the same as for a regular TRILL Data frame with the exception that a RBridge implementing the TRILL OAM Channel MUST recognize the Any-RBridge egress nickname in unicast TRILL Data frames, decapsulating and not forwarding such frames if they meet other checks.

If the OAM-Channel critical ingress-to-egress bit option is present and the egressing RBridge does not implement the TRILL OAM Channel, the frame is discarded. If other options are present, they may affect processing or cause the frame to be discarded.

On decapsulation, the special Inner.MacDA values of OAM-RBridge-MAC (unicast) and All-OAM-RBridges (multicast) and/or the Inner Ethertype of TRILL-OAM MUST be recognized to trigger processing as a TRILL OAM message. If the decapsulating RBridge does not implement the TRILL OAM Channel, it will treat the frame as a regular TRILL Data frame and locally flood the decapsulated native frame out all ports where it is appointed forwarder for the Inner.VLAN.

2.5.1 Processing the TRILL OAM Channel Header

Knowing that it has a TRILL OAM Channel message, the egress RBridge looks at the OV (OAM Message Header version) and OAM Protocol fields.

If the OV field is non-zero or if the OAM Protocol field is a reserved value or a value unknown to the egress RBridge, the egress RBridge returns an OAM Channel Error frame unless the "SL" (Silent) flag is a one in the OAM message. An OAM Channel Error frame is a multi-hop unicast TRILL OAM Channel message with the ingress nickname set to the nickname of the RBridge detecting the error, and the egress nickname set to the value of the ingress nickname in the OAM message for which the error was detected. For the protocol specific data area, an OAM Channel Message Error frame has at least the first 256 bytes (or less if less are available) of the erroneous decapsulated OAM message starting with the Inner.MacDA. All RBridges implementing the TRILL OAM Message Channel MUST recognize the OAM Message Channel Error protocol value (0x001) and MUST NOT generate an OAM Message Channel Error message in response to a received OAM Message Channel Error frame, even if they always set the "SL" flag is all TRILL OAM messages they send so they would not normally expect to receive an OAM Channel Message Error frame.

If the OV field is zero and the processing RBridge recognizes the OAM Protocol value, it processes the message in accordance with that OAM protocol.

Errors within a recognized OAM Protocol are handled within that protocol and do not produce OAM Message Channel Error frames.

2.5.2 Native TRILL-OAM Frames

A TRILL OAM Message Channel frame MAY be generated, if provided for by the OAM protocol involved, as the result of the receipt by an RBridge of a native frame with the TRILL-OAM Ethertype. Such a native frame must met the usual VLAN restrictions to be accepted by the ingress RBridge generating the TRILL OAM Message Channel frame. If the native frame’s destination MAC address is not one of the special MAC destination addresses All-OAM-RBridges or OAM-RBridge-MAC, it MUST be changed to one of those two addresses before the frame is encapsulated.

The decapsulation and processing of a TRILL OAM Message Channel frame MAY, if provided for by the OAM protocol involved, result in the sending of a native frame with the TRILL-OAM Ethertype out one or more ports of the egress RBridge. The VLAN, and the MAC destination addres, of the frame MAY be set to appropriate values before it is transmitted.

3. TRILL BFD

Using the TRILL OAM Message Channel facility, described in Section 2, TRILL supports one-hop and multi-hop BFD Control and neighbor BFD Echo as detailed below. Multi-destination BFD is beyond the scope of this document.

3.1 Sessions and Initialization

Within an RBridge campus, there will be only a single TRILL BFD Control session between two RBridges over a given interface visible to TRILL. 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 PPP [TrillPPP]. Link technology specific OAM protocols may be used on a link between neighbor RBridges, for example Continuity Fault Management [802.1ag] if the link is Ethernet. But such link layer OAM and coordination between it and TRILL data plan layer OAM, such as TRILL BFD, is beyond the scope of this document.

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

3.2 TRILL BFD Control Protocol

TRILL BFD Control frames are unicast TRILL OAM Message Channel frames as described in Section 2 above supplemented by the specifications below.

As a unicast message, the M bit in the TRILL Header is zero and the Inner.MacDA is OAM-RBridge-MAC. The TRILL OAM Protocol value is 0x002.

The protocol specific data associated with the TRILL BFD Control protocol is as shown below. See [RFC5880] for further information on the fields after the initial SystemIDs.

TRILL BFD Control Protocol Data:

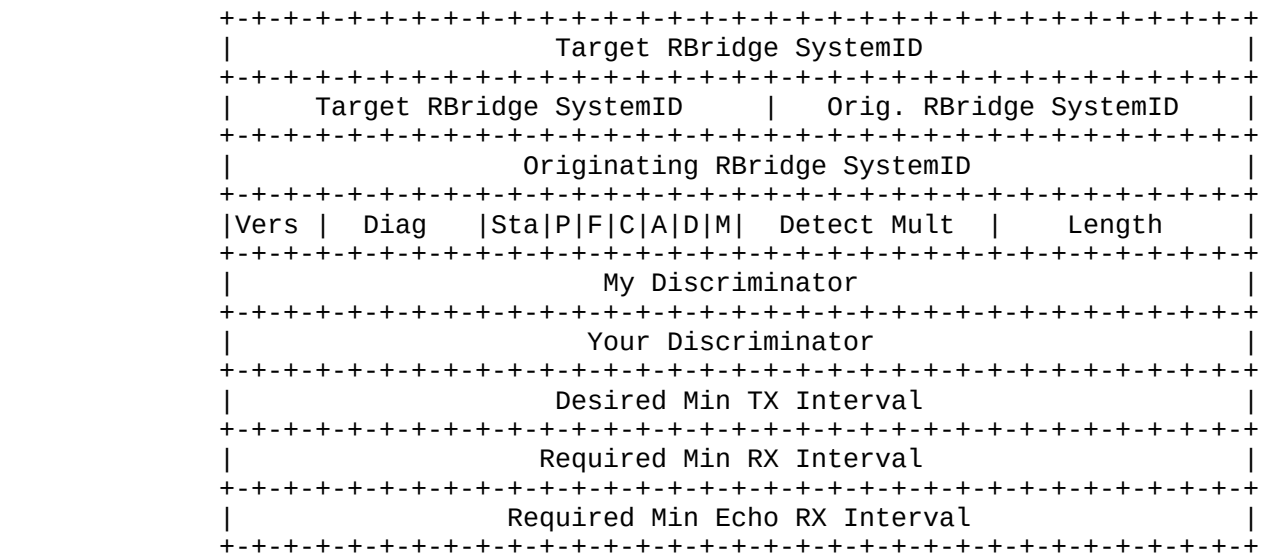

Optional Authentication Section:

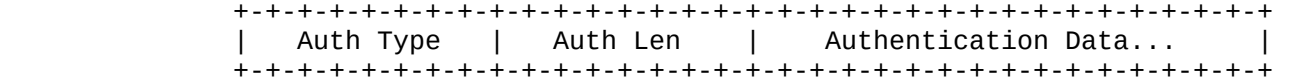

3.2.1 One-Hop TRILL BFD Control

One-hop TRILL BFD Control is typically used in support of TRILL IS-IS to rapidly detect link failure. Such TRILL BFD frames SHOULD be sent with 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 bi-directional connectivity and both RBridges indicate support of TRILL BFD is enabled. The BFD Enabled TLV is used to indicate this as specified in [RFCbfdtlv]. The indication of TRILL BFD support with the BFD Enabled TLV overrides any indication of lack of support through failure to advertise support of the OAM-Channel TRILL Header bit option in the link state database.

3.2.2 BFD Control Frame Processing

The following tests SHOULD be performed on received TRILL BFD Control frames before generic BFD processing. (In some implementations, the TRILL Header may not be available to the TRILL BFD module in which case some of these check are not possible.)

Is the M bit in the TRILL Header non-zero? If so, discard the frame. TRILL support of multi-destination BFD Control is beyond the scope of this document.

If the MH OAM Header 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) and if it is any other value discard the frame. If the MH OAM 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.

Check that the target IS-IS SystemID in the OAM protocol data is your SystemID. If not, discard the frame.

If the MH OAM Header flag is zero, test that the originating SystemID is that of a neighbor RBridge. If not, discard the frame.

3.3 TRILL BFD Echo Protocol

A TRILL BFD Echo frame is a unicast TRILL OAM Message Channel frame, as specified in Section 2, which should be bounced 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.

TRILL BFD Echo frames SHOULD only be sent on a link if a TRILL BFD Control session has been established, TRILL BFD Echo support is indicated by the potentially echo responding RBridge, 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 RECOMMENDED 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 SystemID, the neighbor’s SystemID, a session identifier and a sequence count as well as a Message Authentication Code.

3.3.1 BFD Echo Frame Processing

The following tests SHOULD be performed on returned TRILL BFD Echo frames before other processing. (In some implementations, the TRILL

Header may not be available to the TRILL BFD Echo module in which case these check are not possible.)

Is the M bit in the TRILL Header non-zero? If so, discard the frame. TRILL support of multi-destination BFD Echo is beyond the scope of this document.

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 value of the MH flag is ignored for TRILL BFD Echo protocol.)

4. Management and Operations Considerations

The TRILL BFD parameters at an RBridge are configurable... The default values are ... TBD.

It is required that the operator of an RBridge campus configure the rates at which TRILL BFD frames are transmitted on a link to avoid congestion (e.g., link, I/O, CPU) and false failure detection.

5. Allocations Considerations

The following subsection give IANA and IEEE Registration Authority Considerations.

5.1 IANA Considerations

In this section, the allocation procedures "Standards Action", "IETF Review", and "RFC Publication" are as specified in [RFC5226].

IANA hereby allocates a previously unassigned TRILL Nickname as follows:

Any-RBridge	TBD (0xFFCO suggested)

IANA hereby allocates a previously unassigned TRILL Multicast address as follows:

All-OAM-RBridges	TBD (01-80-C2-00-00-43 suggested)

IANA hereby allocates a previously unassigned TRILL critical ingress-to-egress Bit Option as follows:

TBD	OAM-Option

IANA allocates the following block of 16 globally unique unicast MAC addresses for use with the TRILL protocol and creates a sub-registry in the TRILL Parameter Registry for these addresses:

00-00-5E-xx-xx-x0 - OAM-RBridge-MAC
00-00-5E-xx-xx-x1 to 00-00-5E-xx-xx-xF - available for allocation (suggested 00-00-5E-00-03-00 through 00-00-5E-00-03-0F)

Allocation of unicast MAC values from the above block for TRILL use is based on IETF Review.

IANA creates an additional sub-registry in the TRILL Parameter Registry for TRILL OAM Protocols, with initial contents as follows:

Protocol	Use
0x000	Reserved
0x001	OAM Channel Error
0x002	BFD Control
0x003	BFD Echo
0x004-0x0FF	Available for allocation (1)
0x100-0xFF7	Available for allocation (2)
0xFF8-0xFFE	For Experimental use, will not be allocated
0xFFF	Reserved

(1) TRILL OAM protocol code points from 0x004 to 0x0FF require an IETF Standards Action for allocation.

(2) TRILL OAM protocol code points from 0x100 to 0xFF7 require RFC Publication to allocate a single value or IETF Review to allocate multiple values.

IANA creates an additional sub-registry in the TRILL Parameter Registry for TRILL OAM Header Flags with initial contents as follows:

Flag Bit	Mnemonic	Allocation
0	SL	Silent
1	MH	Multi-hop
2-15	-	Available for allocation

Allocation of TRILL OAM Header Flags is based on IETF Standards Action [RFC5226].

5.2 IEEE Registration Authority Considerations

The Ethertype <tbd> is assigned by the IEEE Registration Authority for TRILL-OAM.

6. Security Considerations

The following sections provide security considerations for the TRILL OAM Message Channel and for TRILL BFD.

See [RFCtrill] for general RBridge Security Considerations.

6.1 OAM Channel Security Considerations

-- TBD --

6.2 BFD Security Considerations

BFD Control frames can be secured by authentication mechanisms native to BFD [RFC5880].

If shared secret IS-IS authentication is not in effect for the Hellos exchanged by two neighbor RBridges then, by default, TRILL BFD between those RBridges is also unsecured.

If shared secret IS-IS authentication is in effect for the Hellos exchanged by two neighbor RBridges then, by default, TRILL BFD Control frames sent between those RBridges use BFD Keyed SHA1 authentication with keying material derived as follows:

```
HMAC-SHA1 ( ( "TRILL BFD Control" | smallerSystemID |
            largerSystemID ), IS-IS-key )
```

where HMAC-SHA1 is as specified in [RFC2104] (see also [RFC4634]), "TRILL BFD Control" is the seventeen byte US ASCII string indicated which is then concatenated with the SystemIDs of both of the neighbor RBridges sorted as unsigned 48-bit integers, and IS-IS-key is the secret keying material being used for IS-IS authentication on the link. In the Authentication Section of the BFD Control frame OAM protocol specific data, Auth Type would be 4, Auth Len would be 28, and Auth Key ID is zero. The RBridges MAY be configured to use other BFD security modes or keying material including configuration to use no security.

Authentication for TRILL BFD Echo SHOULD be provided but is a local implementation issue as BFD Echo frames are only authenticated by their sender when received in the form of Echo responses. 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 authenticate SHOULD use different keying material from other types of authentication. For example, it

could use keying material derived as follows:

HMAC-SHA1 (("TRILL BFD Echo" | smallerSystemID | largerSystemID), IS-IS-key)

7. Normative References

[RFC2104] - Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", RFC 2104, February 1997.

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

[RFC5226] - Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

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

[RFC5882] - D. Katz, D. Ward, "Generic Application of Bidirectional Forwarding Detection (BFD)", June 2010.

[RFCtrill] - R. Perlman, D. Eastlake, D. Dutt, S. Gai, and A. Ghanwani, "RBridges: Base Protocol Specification", draft-ietf-trill-rbridge-protocol-16.txt, in RFC Editor queue.

[RFCbfdtlv] - C. Hopps, L. Ginsberg, "IS-IS BFD Enabled TLV", draft-ietf-isis-bfd-tlv-02.txt, work in progress, 4 January 2010.

8. Informative References

[802.1AX] - IEEE, "IEEE Standard for Local and metropolitan area networks / Link Aggregation", 802.1AX-2008, 1 January 2008.

[802.1ag] - IEEE, "IEEE Standard for Local and metropolitan area networks / Virtual Bridged Local Area Networks / Connectivity Fault Management", 802.1ag-2007, 17 December 2007.

[RFC4634] - Eastlake 3rd, D. and T. Hansen, "US Secure Hash Algorithms (SHA and HMAC-SHA)", RFC 4634, July 2006.

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

[TrillPPP] - Carlson, J., "PPP TRILL Protocol Control Protocol", draft-ietf-pppext-trill-protocol-01.txt, work in progress, May 2010.

[Y.1731] - ITU-T Recommendation Y.1731 (02/08), "OAM functions and mechanisms for Ethernet based networks", February 2008

Authors' Addresses

Donald Eastlake 3rd
Stellar Switches
155 Beaver Street
Milford, MA 01757 USA

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

Vishwas Manral
IP Infusion Inc.
1188 E. Arques Ave.
Sunnyvale, CA 94089 USA

Tel: +1-408-400-1900
EMail: vishwas@ipinfusion.com

Dave Ward
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089-1206
USA

Phone: +1-408-745-2000
EMail: dward@juniper.net

Ayan Banerjee
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95138 USA

Phone: +1-408-525-8781
EmMail: ayabaner@cisco.com