

TRILL Working Group
INTERNET-DRAFT
Intended status: Proposed Standard
Expires: April 17, 2015

Weiguo Hao
Donald Eastlake
Huawei
October 18, 2015

TRILL: Address Flush Protocol
<draft-hao-trill-address-flush-00.txt>

Abstract

The TRILL (TRAnsparent Interconnection of Lots of Links) protocol, by default, learns end station addresses from observing the data plane. This document specifies an optional message by which an originating TRILL switch can explicitly flush addresses learned by other TRILL switches through the egress of data ingress by that originating TRILL switch. This is a supplement to the TRILL automatic address forgetting and can assist in achieving more rapid convergence.

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

Table of Contents

1. Introduction.....	3
1.1 Terminology and Acronyms.....	3
2. Address Flush Message Details.....	5
3. IANA Considerations.....	9
4. Security Considerations.....	9
Normative References.....	10
Informative References.....	10
Acknowledgements.....	10
Authors' Addresses.....	11

1. Introduction

Edge TRILL (Transparent Interconnection of Lots of Links [RFC6325]) switches, also called RBridges, by default learn end station MAC addresses from observing the data plane. On receipt of a native frame from an end station, they would learn the local MAC address attachment of the source end station. And on egressing (decapsulating) a remotely originated TRILL Data frame, they learn the remote MAC address and remote attachment TRILL switch. Such learning is all appropriately scoped by data label (VLAN or Fine Grained Label [RFC7172]).

TRILL has mechanisms for timing out such learning and appropriately clearing it based on some network connectivity changes; however, there are circumstances under which it would be helpful for a TRILL switch to be able to explicitly flush (clear) learned end station reachability information to achieve more rapid convergence (see, for example, Section 6.2 of [RFC4762]). Obviously a TRILL switch R1 can easily flush any locally learned addresses it wants. This document specifies an optional message to request flushing such learned address information at remote TRILL switches. This Address Flush message makes use of the RBridge Channel facility [RFC7178], which supports typed message transmission between RBridges.

1.1 Terminology and Acronyms

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].

This document uses the terms and acronyms defined in [RFC6325] and [RFCchannel] as well as the following:

AFN - Address Family Number ([RFC4760] where it is called Address Family Identifier (AFI)).

FGL - Fine Grained Label [RFC7172].

Management VLAN - A VLAN in which all TRILL switches in a campus indicate interest so that multi-destination TRILL Data packets, including RBridge Channel messages [RFCchannel], sent with that VLAN as the Inner.VLAN will be delivered to all TRILL switches in the campus. Usually no end station service is offered in the Management VLAN.

RBridge - A alternative name for a TRILL switch.

TRILL switch - A device implementing the TRILL protocol.

Edge TRILL switch - A TRILL switch attached to one or more links that provide end station service.

2. Address Flush Message Details

The Address Flush message makes use of the RBridge Channel protocol [RFC7178].

Although initial use is expected to be to flush 48-bit MAC addresses [RFC7042], the protocol accommodates flushing other types of end station addresses; there have been suggestion for TRILL switches to learn IP addresses from the data plane [INFOCOM], TRILL might be extended to accommodate 64-bit MAC addresses, or similar future extensions might benefit from the ability to flush other types of learned addresses.

The general structure of an RBridge Channel packet on a link between TRILL switches is shown in Figure 1 below. The type of RBridge Channel packet is given by a Protocol field in the RBridge Channel Header that indicates how to interpret the Channel Protocol Specific Payload [RFC7178].




Figure 1. RBridge Channel Packet Structure

An Address Flush RBridge Channel message normally applies to addresses within the VLAN or FGL [RFC7178] Data Label in the TRILL Header. Address Flush protocol messages are usually sent as multi-destination packets (TRILL Header M bit equal to one) so as to reach all TRILL switches offering end station service in the VLAN or FGL specified by the Data Label. However, an address flush protocol message can be sent unicast, if it is desired to clear addresses at one TRILL switch only. And there are provisions for indicating the Data Label with the address(es) to be flushed for cases where the address flush protocol message is sent over a Management VLAN or the like.

Figure 2 below expands the RBridge Channel Header and Channel

Protocol Specific Payload from Figure 1 for the case of the Address Flush message.

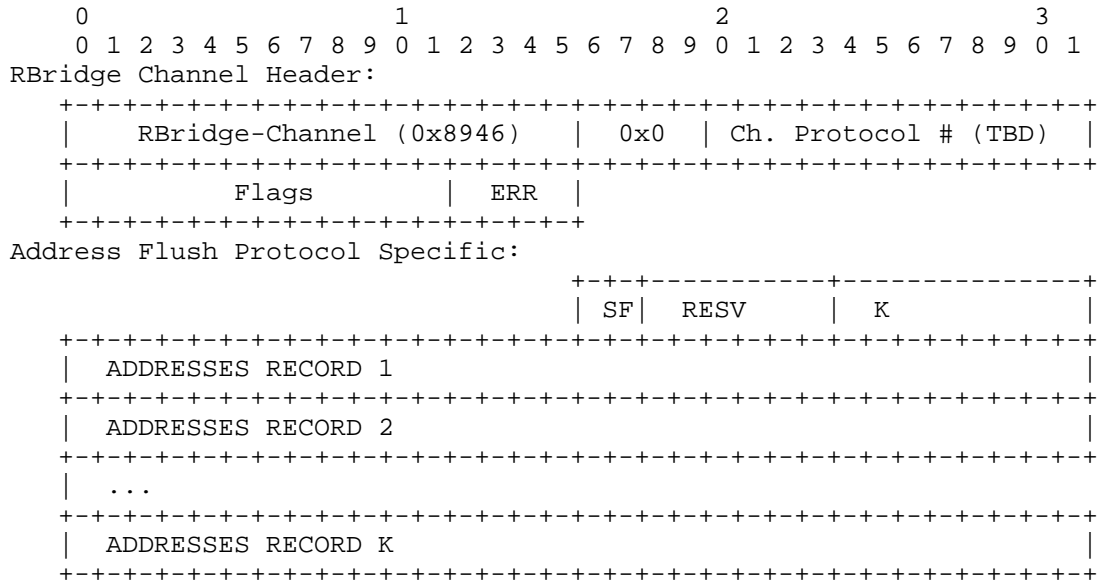


Figure 2. Address Flush Channel Message Structure

The fields in Figure 2 related to the Address Flush message are as follows:

Channel Protocol: The RBridge Channel Protocol value allocated for Address Flush (see Section 3).

SF: The 2-bit SF ("super flush") field values have the following meanings:

0: No special effect.

1: All addresses learned at the receiving TRILL switch due to egressing TRILL Data packets from the TRILL switch originating this Address Flush message are flushed for the data label in the TRILL Header. Any ADDRESS RECORDs in the rest of the message for that data label can be ignored but there may be ADDRESS RECORDs present that apply to other data labels.

2: All addresses learned at the receiving TRILL switch due to egressing TRILL Data packets from the TRILL switch originating this Address Flush message are flushed across all data labels. The remainder of the Address Flush message, including the value of K, are ignored.

3: Reserved. Ignored on receipt.

RESV: 4 reserved flag bits. Must be sent as zero and ignored on receipt.

K: The number of ADDRESS RECORDs present. See below.

The structure of the ADDRESSES RECORD is as follows:

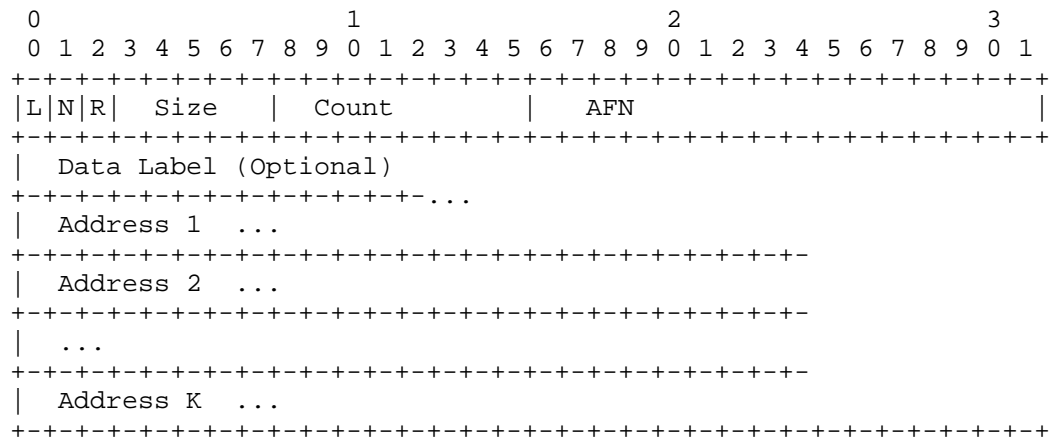


Figure 3. Structure of ADDRESSES RECORD

L: Label present. If this bit is a one, the optional Data Label shown in present. If it is zero, there is no data label and the addresses listed are within the data label given in the TRILL Header.

N: No Specific Addresses. If this bit is one and Count is zero and L is one, then flush all addresses learned at the receiving TRILL switch due to egressing TRILL Data packets from the TRILL switch originating this Address Flush message are flushed for the Data Label given in the ADDRESS RECORD. If this bit is zero or Count is non-zero or L is zero, then this special flush action is not performed.

R: A reserved bit that MUST be sent as zero and is ignored on receipt.

Size: The size of each Address in bytes. The presence of this field makes it possible for a receiving TRILL switch to skip an ADDRESS RECORD even if it does not understand the value in the AFN field. Size MUST NOT be zero; a zero size field indicates a corrupt Addresses Flush message and the entire message is ignored. MUST be the correct size for an Address

of the type indicated by the AFN field, for example 6 for 48-bit MAC addresses. If these conditions are violated, the Address Flush message is discarded.

Count: The number of occurrences of an Address to flush in this ADDRESS RECORD. May be zero. All Addresses MUST fit within the RBridge Channel Message. If they do not, the message is discarded.

AFN: The Address Family Number for the type of addresses present as assigned by IANA. (The AFN for 48-bit MAC addresses is 0x4005.)

Data Label: An optional Data Label (VLAN or FGL) in the same format as Data Labels that appear in the TRILL Header. Included in an ADDRESS RECORD only if the L bit is a one.

Address: An instance of an address to be flushed.

3. IANA Considerations

IANA has allocated tbd1 for the Address Flush RBridge Channel Protocol number from the range of RBridge Channel protocols allocated by Standards Action [RFC7178].

4. Security Considerations

The Address Flush RBridge Channel Protocol provides no security assurances or features. However, use of the Address Flush protocol can be nested inside the RBridge Channel Tunnel Protocol [RFCtunnel] using the RBridge Channel message payload type. The Channel Tunnel protocol can provide some security services.

See [RFC7178] for general RBridge Channel Security Considerations.

See [RFC6325] for general TRILL Security Considerations.

Normative References

- [RFC2119] - Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4760] - Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, January 2007.
- [RFC6325] - Perlman, R., D. Eastlake, D. Dutt, S. Gai, and A. Ghanwani, "RBrigdes: Base Protocol Specification", RFC 6325, July 2011.
- [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, <<http://www.rfc-editor.org/info/rfc7172>>.
- [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, <<http://www.rfc-editor.org/info/rfc7178>>.

Informative References

- [INFOCOM] - Perlman, R., "RBrigdes: Transparent Routing", Proc. Infocom 2005, March 2004.
- [RFC4762] - Lasserre, M., Ed., and V. Kompella, Ed., "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", RFC 4762, January 2007.
- [RFC7042] - Eastlake 3rd, D. and J. Abley, "IANA Considerations and IETF Protocol and Documentation Usage for IEEE 802 Parameters", BCP 141, RFC 7042, DOI 10.17487/RFC7042, October 2013, <<http://www.rfc-editor.org/info/rfc7042>>.
- [RFCtunnel] - Eastlake, D., ... "TRILL: Channel Tunnel", draft-eastlake-trill-channel-tunnel, work in progress.

Acknowledgements

The document was prepared in raw nroff. All macros used were defined within the source file.

Authors' Addresses

Weiguo Hao
Huawei Technologies
101 Software Avenue,
Nanjing 210012, China

Phone: +86-25-56623144
Email: haoweiguo@huawei.com

Donald E. Eastlake, 3rd
Huawei Technologies
155 Beaver Street
Milford, MA 01757 USA

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

Copyright, Disclaimer, and Additional IPR Provisions

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. The definitive version of an IETF Document is that published by, or under the auspices of, the IETF. Versions of IETF Documents that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of IETF Documents. The definitive version of these Legal Provisions is that published by, or under the auspices of, the IETF. Versions of these Legal Provisions that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of these Legal Provisions. For the avoidance of doubt, each Contributor to the IETF Standards Process licenses each Contribution that he or she makes as part of the IETF Standards Process to the IETF Trust pursuant to the provisions of RFC 5378. No language to the contrary, or terms, conditions or rights that differ from or are inconsistent with the rights and licenses granted under RFC 5378, shall have any effect and shall be null and void, whether published or posted by such Contributor, or included with or in such Contribution.

