

INTERNET-DRAFT
Updates: RFCchannel
Intended status: Proposed Standard
Expires: April 20, 2014

Donald Eastlake
Yizhou Li
Huawei
October 21, 2013

TRILL: RBridge Channel Tunnel Protocol
<draft-eastlake-trill-channel-tunnel-00.txt>

Abstract

The IETF TRILL (Transparent Interconnection of Lots of Links) protocol includes an optional mechanism, called RBridge Channel, for the transmission of typed messages between TRILL switches in the same campus and between TRILL switches and end stations on the same link. This document specifies optional extensions to RBridge Channel that provides three facilities: (1) A mechanism to send such messages between a TRILL switch and an end station in either direction, or between two end stations, when the two devices are in the same campus but not on the same link; (2) A method to support security facilities for RBridge Channel messages; and (3) A method to tunnel a variety of payload types by encapsulating them in an RBridge Channel message.

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 authors or the TRILL working group mailing list:
trill@ietf.org

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
1.2. Terminology and Acronyms.....	3
2. Channel Tunnel Packet Format.....	5
3. Tunnel Payload Types.....	8
3.1 Null Payload.....	8
3.2 RBridge Channel Message Payload.....	8
3.3 TRILL Data Packet.....	9
3.4 TRILL IS-IS Packet.....	10
3.5 Ethernet Frame.....	11
3. Channel Tunnel Scopes.....	13
3.1 End Station to RBridge(s).....	14
3.2 RBridge to End Station.....	15
3.3 End Station to End Station.....	16
5. Security, Keying, and Algorithms.....	18
5.1 SType None.....	18
5.2 RFC 5310 Based Authentication.....	18
5.3 DTLS Based Security.....	19
6. Channel Tunnel Errors.....	20
6.1 SubERRs under ERR 6.....	20
7. IANA Considerations.....	21
8. Security Considerations.....	21
Normative References.....	22
Informative References.....	22
Acknowledgements.....	23
Authors' Addresses.....	24

1. Introduction

The IETF TRILL protocol [RFC6325] provides efficient least cost transparent frame routing in multi-hop networks with arbitrary topologies and link technologies, using link-state routing and a header with a hop count. End stations are attached to TRILL switches by Ethernet but links between TRILL switches can be arbitrary technology. In general, the TRILL way to address or specify a TRILL switch (RBridge) in a TRILL campus is by the switch's TRILL provided nickname [RFC6325] [ClearCorrect].

The TRILL protocol includes an optional RBridge Channel facility [RFCchannel] to support typed message transmission between two RBridges (for example BFD [RFCbfd]) in the same campus and between RBridges and end stations on the same link.

This document specifies optional extensions to RBridge Channel that provides three facilities:

- (1) A mechanism to send RBridge Channel messages between a TRILL switch and an end station in either direction, or between two end stations, when the two devices are in the same campus but not on the same link. This mechanism requires the cooperation of an RBridge that is on the same link as the end station or stations involved.
- (2) A method to support security facilities for RBridge Channel messages.
- (3) A method to tunnel a variety of payload types by encapsulating them in an RBridge Channel message.

Any one, two, or all three of these facilities can be use in the same message.

There is no mechanism to stop end stations on the same link, from sending native RBridge Channel messages to each other; however, such use is outside the scope of this document.

1.2. 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 acronyms defined in [RFC6325] and [RFCchannel] supplemented by the following additional acronym:

Data Label - VLAN or Fine Grained Label [RFCfgl].

Primary Nickname - If a TRILL switch holds two or more nicknames, the one it holds with the highest priority is the primary nickname. If two or more are held with the same priority, the one with the lowest value, considered as a 16-bit unsigned integer in network byte order, is the primary nickname.

TRILL switch - An alternative term for an RBridge.

2. Channel Tunnel Packet Format

The general structure of an RBridge Channel message on a link between TRILL switches (RBridges) is shown in Figure 1 below. When a native RBridge Channel message is sent between an RBridge and an end station on the same link, in either direction, the TRILL Header (including the inner Ethernet addresses and Data Label) is omitted as shown in Figure 2. The type of RBridge Channel message is given by a Protocol field in the RBridge Channel Header which indicates how to interpret the Channel Protocol Specific Payload [RFCchannel].

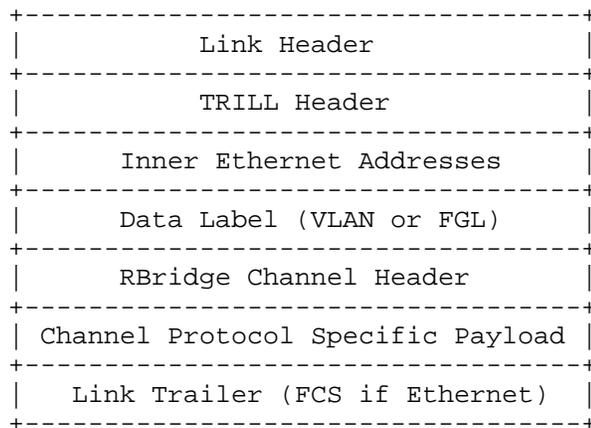


Figure 1. RBridge Channel Packet Structure

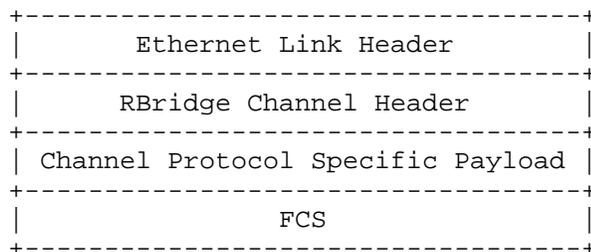


Figure 2. Native RBridge Channel Frame

The RBridge Channel Header looks like this:

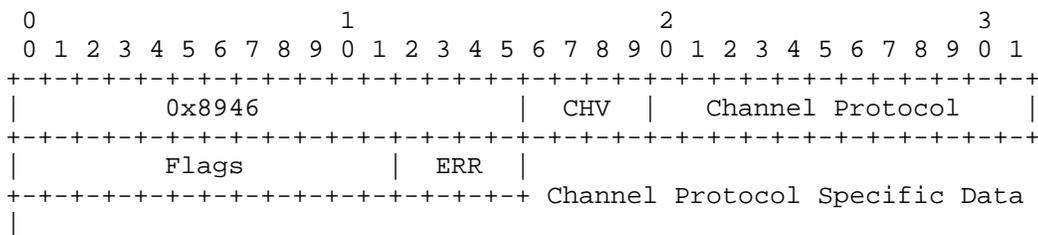


Figure 3. RBridge Channel Header

where 0x8946 is the RBridge Channel Ethertype and CHV is the Channel Header Version, currently zero.

The extensions specified herein are in the form of an RBridge Channel protocol, the Channel Tunnel Protocol. Figure 4 below expands the RBridge Channel Header and Protocol Specific Payload above for the case of the Channel Tunnel Protocol.

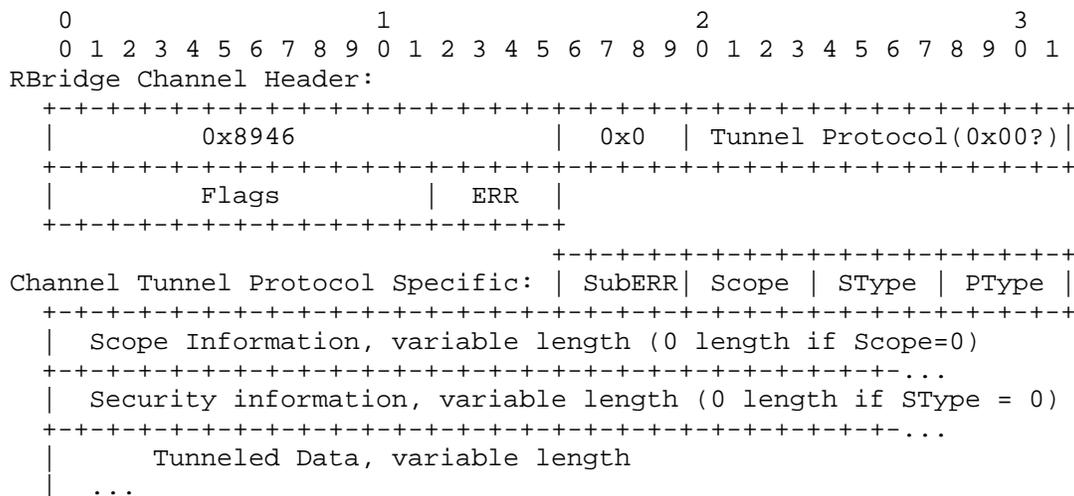


Figure 4. Channel Tunnel Header Structure

The RBridge Channel Header field specific to the RBridge Channel Tunnel Protocol is the Protocol field. Its contents MUST be the value allocated for this purpose (see Section 7).

The RBridge Tunnel Channel Protocol Specific fields are as follows:

SubERR: This field provides further details when a Tunnel Channel error is indicated in the RBridge Channel ERR field. If ERR is zero, then SubERR MUST be sent as zero and ignored on receipt. See Section 6.

Scope: This field describes the transport scope of the instance of Channel Tunnel. See Section 4.

SType: This field describes the type of security information and features, including keying material, being provided. See Section 5.

PType: Payload type. The describes the tunneled data. See Section 3 below.

The Channel Tunnel protocol is integrated with the RBridge Channel facility. Channel Tunnel errors are reported as if they were RBridge Channel errors, using newly allocated code points in the ERR field of the RBridge Channel Header supplemented by the SubErr field. Additional RBridge Channel Header flags are specified and used by Channel Tunnel.

3. Tunnel Payload Types

The RBridge Channel Tunnel Protocol can carry a variety of payloads as indicated by the PType field. Value are shown in the table below with further explanation after the table.

PType	Section	Description
----	-----	-----
0		Reserved
1	3.1	Null
2	3.2	RBridge Channel message
3	3.3	TRILL Data packet
4	3.4	TRILL IS-IS packet
5	3.5	Ethernet Frame
6-14		(Available for assignment by IETF Review)
15		Reserved

Table 1. Payload Type Values

While implementation of the Channel Tunnel protocol is optional, if it is implemented PTypes 1 (Null) and 2 (RBridge Channel message) MUST be implemented. PTypes 3, 4, and 5 MAY be implemented. The processing of any particular Channel Protocol message and its payload depends on meeting local security and other policy at the destination TRILL switch or end station.

3.1 Null Payload

The Null payload type is intended to be used for messages such as key negotiation or the like. It indicates that there is no payload. Any data after the possible Scope Information and Security Information fields is ignored.

3.2 RBridge Channel Message Payload

A PType of 2 indicates that the payload of the Channel Tunnel message is an encapsulated RBridge Channel message without the initial RBridge Channel Ethertype. Typical reasons for sending an RBridge Channel message inside a Channel Tunnel message are to provide security services, such as authentication or encryption, or to forward it through a cooperating border TRILL switch in either direction between an end station and a TRILL switch not on the same link.

This looks like the following:

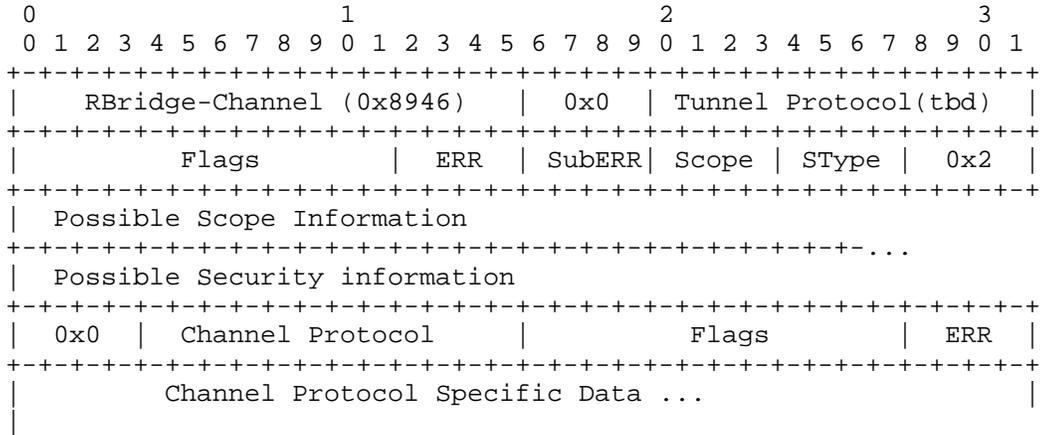


Figure 5. Tunneled Channel Message Channel Tunnel Structure

3.3 TRILL Data Packet

A PType of 3 indicates that the payload of the Tunnel protocol message is an encapsulated TRILL Data packet without the initial TRILL Ethertype as shown in the figure below. If this PType is implemented, the tunneled TRILL Data packet is handled as if it had been received by the destination TRILL switch on the port where the Channel Tunnel message was received.

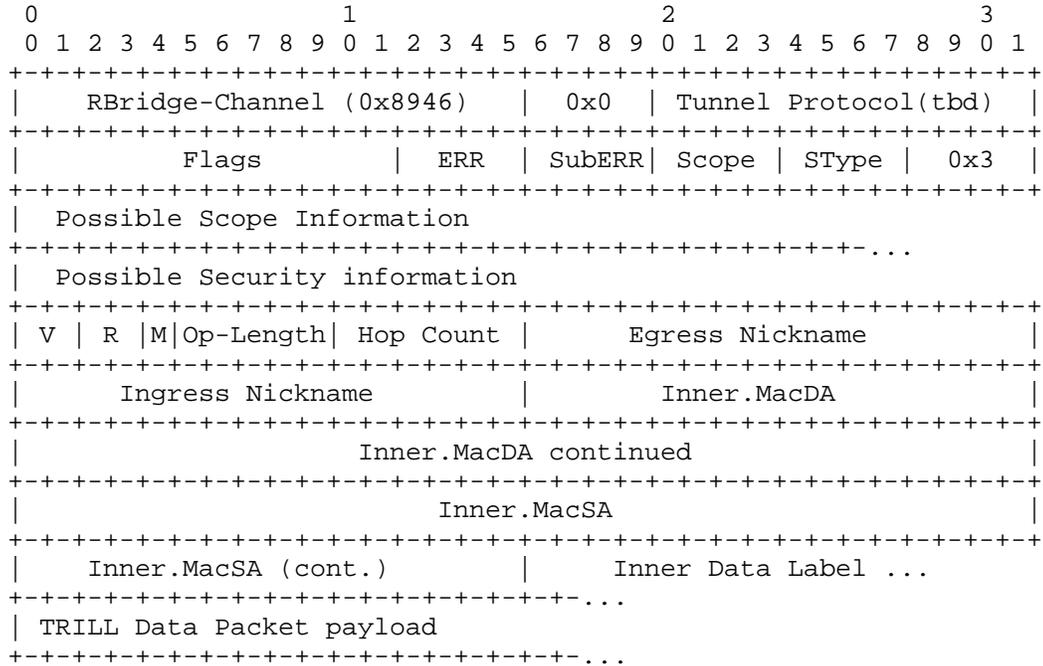


Figure 6. Nested TRILL Data Packet Channel Tunnel Structure

3.4 TRILL IS-IS Packet

A PType of 4 indicates that the payload of the Tunnel protocol message is an encapsulated TRILL IS-IS packet without the initial L2-IS-IS Ethertype as shown in the figure below. If this PType is implemented, the tunneled TRILL IS-IS packet is processed by the destination RBridge if it meets local policy. The intended use is to expedite the receipt of a link state PDU by some TRILL switch with an immediate requirement for the enclosed link state data. It is RECOMMENDED that any link local IS-IS PDU (Hello, xSNP, MTU-x) received via this channel tunnel payload type be discarded.

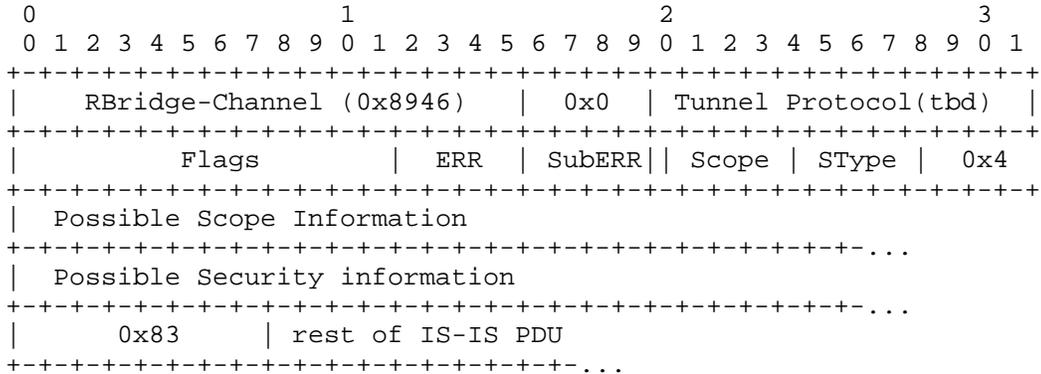


Figure 7. Tunneled TRILL IS-IS Packet Structure

3.5 Ethernet Frame

If PType is 5, the Tunnel Protocol payload is an Ethernet frame as might be received from or sent to an end station except that the tunneled Ethernet frame's FCS is omitted, as shown in Figure 8. (There is still an overall FCS if the RBridge Channel message is being sent on an Ethernet link.) If this PType is implemented, the tunneled frame is handled as if it had been received on the port on which the Tunnel Protocol message was received.

In the case of a non-Ethernet link, such as a PPP link [RFC6361], the ports on the link are considered to have link local synthetic 48-bit MAC addresses constructed by concatenating three 16-bit quantities: 0xFEFF, the primary nickname of the TRILL switch (see Section 1.2), and the Port ID that the RBridge has assigned to that port, as shown in Figure 9. The resulting MAC address has the Local bit on and the Group bit off [RFC5342bis]. Since end stations are connected to TRILL switches only over Ethernet, there can be no end stations on a non-Ethernet link in a TRILL campus. Thus such synthetic MAC addresses cannot conflict on the link with an end station address.

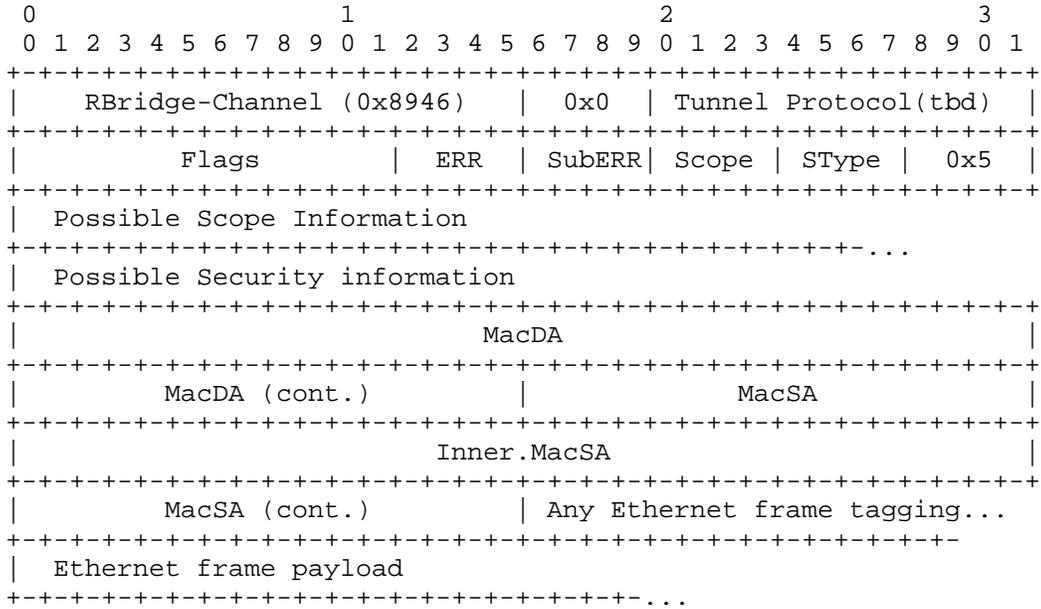


Figure 8. Ethernet Frame Channel Tunnel Structure

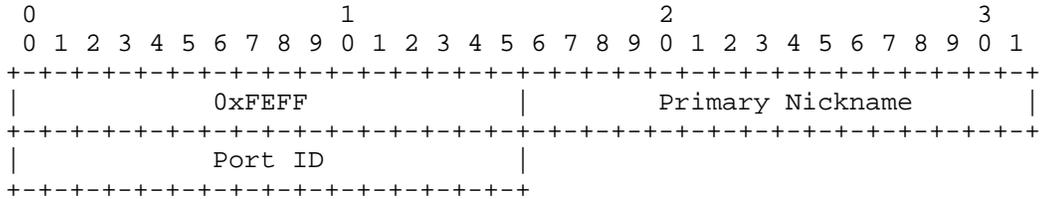


Figure 9. Synthetic MAC Address

3. Channel Tunnel Scopes

The Channel Tunnel protocol extends the RBridge Channel facility to optionally support typed messages between an end station and a TRILL switch, in either direction, or between two end stations, when these devices are part of the same TRILL campus but not on the same link. The scopes specified in this document are as follows:

Scope	Symbol	Section	From-To
----	-----	-----	-----
0	NORM		Normal
1	ESRB	3.1	End Station to RBridge
2	RBES	3.2	RBridge to End Station
3	ESES	3.3	End Station to End Station
4-14		Available for assignment by IETF Review	
15		Reserved	

Table 2. Scope Values

If the Channel Tunnel protocol is supported, then the NORM scope MUST be supported. All other scopes MAY be supported. In cases where a sequence of steps is given, other processing sequences producing the same result are, as always, allowed. The detail are given below.

NORM: This is the normal scope of an RBridge Channel message. The base RBridge Channel mechanisms apply [RFCchannel]. The scope dependent addressing information is of zero length. This scope is typically used when just the security or payload type features of the Tunnel Protocol are desired. If a TRILL switch supports the Channel Tunnel facility, it MUST support NORM scope.

ESRB: From end station to RBridge(s) not on the same link. The scope dependent address information is eight bytes long. See Section 4.1 for further details. This scope MAY be supported.

RBES: From RBridge to end station not on the same link. The scope dependent address information is eight bytes long. See Section 4.2 for further details. This scope MAY be supported.

ESES: From end station to en station not on the same link. The scope information is twelve bytes long. See Section 4.3 for further details. This scope MAY be supported.

It is an implementation option and may depend on local policy whether or not an edge TRILL switch that has been requested to forward a Channel Tunnel protocol message due to a non-NORM Scope examines the SType and, if it does examine the SType, whether it verifies any authentication.

3.1 End Station to RBridge(s)

The ESRB scope additional information is as follows:

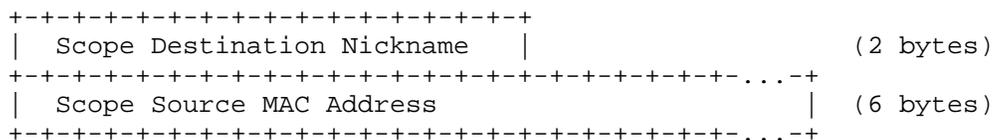


Figure 10. ESRB Scope Information

To support the case where an end station originates a multi-destination RBridge Channel message to all the TRILL switches advertising interest in a Data Label, the BR (Broadcast) bit in the RBridge Channel Header Flags field is used (see Section 7).

Steps by the source end station:

If the RBridge Channel message is intended to a single destination RBridge, the source end station sets the Scope Destination Nickname to the nickname of that RBridge and ensures that the BR bit is zero. If the message is intended to be broadcast to the RBridges indicating interest in a Data Label, the end stations sets the BR bit, uses that Data Label as part of the TRILL Header information, and the contents of the Scope Destination Nickname field is ignored.

Steps by the ingress TRILL switch on receiving the native RBridge Channel message from the end station:

0. As with any RBridge Channel message, determine, as a matter of local policy, whether the native RBridge Channel message is acceptable and discard it if it is not. This test might take into account, for example, whether the message is authenticated (see Section 5), whether or not the BR flag is set, and whether or not the original native destination MAC address is All-Edge-RBridges.
1. Store the native RBridge Channel message's source MAC address into the Scope Source MAC Address field.
2. Clear the NA bit and set the MH bit in the RBridge Channel Header flags.
3. Set the native RBridge Channel message's MAC destination address to All-Egress-RBridges.
4. Set the native RBridge Channel message's MAC source address to the MAC address that the ingress RBridge normally uses as the Inner.MacSA for RBridge Channel messages it originates.
- 5.a. If the BR flag is zero, ingress the modified native frame as a unicast TRILL RBridge Channel message with egress nickname set from the Scope Destination Nickname. If that Scope

RBridge port's MAC address, respectively.

If the original RBridge Channel message was multi-destination it might be egressed by more than one TRILL switch, each of which would perform the above transform. Whether such a multi-destination RBridge Channel Tunnel Protocol message would be accepted by any particular egress TRILL switch is a matter of local policy.

4.3 End Station to End Station

The ESES scope additional addressing information is as follows:

```

+-----+
| Scope Destination MAC Address | (6 bytes)
+-----+
| Scope Source MAC Address     | (6 bytes)
+-----+

```

Figure 12. ESES Scope Information

Steps by the source end stations:

If the RBridge Channel message is intended for a single destination end station, the source end station sets the Scope Destination MAC address to the MAC address of that end station and ensures that the BR bit is zero. If the message is intended to be broadcast to a set of end stations via a multicast MAC address or the broadcast MAC address, the end station sets the Scope Destination MAC address to that multicast or broadcast address and sets the BR bit. All of this is within the VLAN of the native RBridge Channel message or its Fine Grained Label (FGL) if the ingress port is configured to map to an FGL.

Steps by the ingress TRILL switch:

0. As with any RBridge Channel message, determine, as a matter of local policy, whether the native RBridge Channel message is acceptable and discard it if it is not. This test might take into account, for example, whether the message is authenticated (see Section 5), whether or not the BR flag is set, and whether or not the original Outer.MacDA is All-Edge-RBridges.
1. Store the native RBridge Channel message's source MAC address into the Scope Source MAC Address.
2. Clear the NA bit and set the MH bit in the RBridge Channel Header flags.
3. Set the native RBridge Channel message's MAC destination address to All-Egress-RBridges.

4. Set the native RBridge Channel message's MAC source address to the MAC address that the ingress RBridge normally uses as the Inner.MacSA for RBridge Channel messages it originates.
- 5.a. If the BR flag is zero, lookup the Scope Destination MAC Address and ingress the modified native frame as if it were a unicast native frame with that destination MAC address. This will result in either a unicast TRILL Data packet to the Scope Destination MAC Address or in unknown MAC flooding.
- 5.b. If the BR flag is one, select a distribution tree and ingress the modified native frame as a multi-destination TRILL RBridge Channel message.
- 5.c. Regardless of the BR flag value, the Inner.VLAN is the VLAN ID reported by the ingress port or, if that port is configured for FGL, the Inner.Lable is the FGL that VLAN maps to.
6. Process the resulting RBridge Channel message. Note that if it is unicast to the ingress RBridge as egress, it is then egressed. And if it is multi-destination and the ingress RBridge qualifies, a copy is egressed as well as a copy being sent on the selected distribution tree. It is possible that the Scope Destination MAC is actually out a different or even the same port of the ingress TRILL switch as the port on which the native RBridge Channel message was received.

Steps by the egress TRILL switch(es):

The egress RBridge sets the NA bit in the RBridge Channel Header flags. It then egresses the frame as a native RBridge Channel message, setting the native frame's outer destination and source MAC addresses to the Scope Destination MAC Address and the egress RBridge port's MAC address, respectively.

If the original RBridge Channel message was multi-destination it might be egressed by more than one RBridge, each of which would perform the above transform. Whether such a multi-destination RBridge Channel Tunnel Protocol message would be accepted by egress RBridges is a matter of local policy.

5. Security, Keying, and Algorithms

The following table gives the assigned values of the SType field and their meaning.

SType	Section	Meaning
0	5.1	None
1	5.2	RFC 5310 Based Authentication
2	5.3	DTLS Based Security
3-14		Available for assignment on IETF Review
15		Reserved

Table 3. SType Values

For all SType values except zero, the Security Information starts with a byte of flag bits and a byte of remaining length as follows:

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+...
|A|E|   RESV   |   Size   |   Info
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+...

```

Figure 12. Security Information Format

The fields are as follows:

A: Zero if authentication is not being provided. One if it is.

E: Zero if encryption is not being provided. One if it is.

RESV: Six reserved bits that MUST be sent as zero and ignored on receipt.

Size: The number of byte of Info as an unsigned byte.

Info: Variable length Security Information.

5.1 SType None

No security services are being invoked. The length of the Security Information field (see Figure 6) is zero.

5.2 RFC 5310 Based Authentication

The security information is the same as the value of the Authentication TLV as specified in [RFC5310]. See figure below.

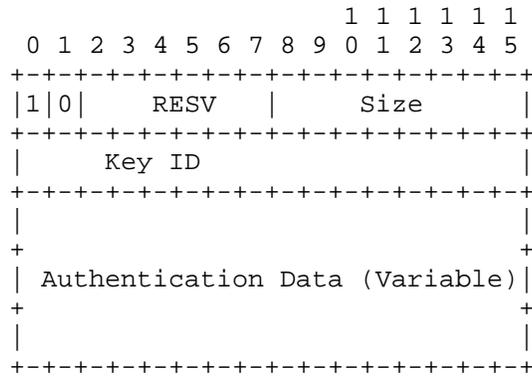


Figure 13. SType 1 Security Information

The Key ID normally specifies a keying value and algorithm.

5.3 DTLS Based Security

TBD - permits key negotiation, provides both encryption and authentication [RFC6347]...

6. Channel Tunnel Errors

RBridge Channel Tunnel Protocol errors are reported like RBridge Channel level errors. The ERR field is set to one of the following error codes:

ERR	Meaning
---	-----
6	Unknown or unsupported field value
7	Authentication failure
	(more TBD)

Table 4. Additional ERR Values

6.1 SubERRs under ERR 6

If the ERR field is 6, the SubERR field indicates

SubERR	Meaning (for ERR = 6)
-----	-----
0	Unsupported Scope
1	Unsupported SType
2	Unsupported PType
3	Unknown or reserved Scope Egress Nickname in an ESRB scope Tunnel Channel message.
4	Unsupported crypto algorithm
	(more TBD)

Table 5. SubERR values under ERR 6

7. IANA Considerations

IANA is requested to allocate a new RBridge Channel protocol number from the range based on Standards Action for the "Channel Tunnel" protocol.

IANA is requested to allocate a new RBridge Channel Header flag bit for the Broadcast (BR) flag with this document as reference.

8. Security Considerations

The RBridge Channel tunnel facility has potentially positive and negative effects on security.

On the positive side, it provides optional security that can be used to authenticate and/or encrypt channel messages. Some RBridge Channel message payloads provide their own security [RFCbfd] but where this is not true, careful consideration should be given to requiring use of the security features of the Tunnel Protocol.

On the negative side, the ability to tunnel various payload types and to tunnel them not just between TRILL switches but to and from end stations can increase risk unless precautions are taken. The processing of decapsulated Tunnel Protocol payloads is not a good place to be liberal in what you accept as the tunneling facility makes it easier for unexpected messages to pop up in unexpected places in a TRILL campus due to accidents or the actions of an adversary. Local policies should generally be strict and only process payload types required and then only with adequate authentication for the particular circumstances.

See [RFCchannel] 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.
- [RFC5310] - Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, February 2009.
- [RFC6325] - Perlman, R., D. Eastlake, D. Dutt, S. Gai, and A. Ghanwani, "RBridges: Base Protocol Specification", RFC 6325, July 2011.
- [RFC6347] - Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, January 2012.
- [ClearCorrect] - Eastlake, D., M. Zhang, A. Ghanwani, V. Manral, A. Banerjee, "TRILL: Clarifications, Corrections, and Updates", draft-ietf-trill-clear-correct, in RFC Editor's queue.
- [RFCchannel] - D. Eastlake, V. Manral, Y. Li, S. Aldrin, D. Ward, "TRILL: RBridge Channel Support", draft-ietf-trill-rbridge-channel-08.txt, in RFC Editor's queue.
- [RFCfgl] - D. Eastlake, M. Zhang, P. Agarwal, R Perlman, D. Dutt, "TRILL: Fine-Grained Labeling", draft-ietf-trill-fine-labeling, in RFC Editor's queue.

Informative References

- [RFC6361] - Carlson, J. and D. Eastlake 3rd, "PPP Transparent Interconnection of Lots of Links (TRILL) Protocol Control Protocol", RFC 6361, August 2011
- [RFC5342bis] - D. Eastlake, J. Abley, " IANA Considerations and IETF Protocol and Documentation Usage for IEEE 802 Parameters", draft-eastlake-rfc5342bis, work in progress.
- [RFCbfd] - Manral, V., D. Eastlake, D. Ward, A. Banerjee, "TRILL (Transparent Interconnection of Lots of Links): Bidirectional Forwarding Detection (BFD) Support", draft-ietf-trill-rbridge-bfd, in RFC Editor's queue.

Acknowledgements

The contributions of the following are hereby acknowledged:

TBD

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

Authors' Addresses

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

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

Yizhou Li
Huawei Technologies
101 Software Avenue,
Nanjing 210012, China

Phone: +86-25-56622310
Email: liyizhou@huawei.com

Copyright, Disclaimer, and Additional IPR Provisions

Copyright (c) 2013 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.

