

INTERNET-DRAFT

Updates: [7178](#)

Intended status: Proposed Standard

Expires: June 7, 2015

Donald Eastlake

Yizhou Li

Huawei

December 8, 2014

TRILL: RBridge Channel Tunnel Protocol
<[draft-ietf-trill-channel-tunnel-02.txt](#)>

Abstract

The IETF TRILL (Transparent Interconnection of Lots of Links) protocol includes an optional mechanism, called RBridge Channel and specified in [RFC 7178](#), 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 two optional extensions to the RBridge Channel protocol: (1) A standard method to tunnel a variety of payload types by encapsulating them in an RBridge Channel message; and (2) A method to support security facilities for RBridge Channel messages. This document updates [RFC 7178](#).

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>.

INTERNET-DRAFT

TRILL: RBridge Channel Tunnel

Table of Contents

1. Introduction.....	3
1.1 Terminology and Acronyms.....	3
2. Channel Tunnel Packet Format.....	4
3. Tunnel Payload Types.....	7
3.1 Null Payload.....	7
3.2 RBridge Channel Message Payload.....	7
3.3 TRILL Data Packet.....	8
3.4 TRILL IS-IS Packet.....	9
3.5 Ethernet Frame.....	10
4. Security, Keying, and Algorithms.....	12
4.1 Authentication Coverage.....	12
4.2 SType None.....	13
4.3 RFC 5310 Based Authentication.....	13
4.4 xxx Based Security.....	13
5. Channel Tunnel Errors.....	14
5.1 SubERRs under ERR 6.....	14
5.2 Nested RBridge Channel Errors.....	14
6. IANA Considerations.....	15
7. Security Considerations.....	16
Normative References.....	17
Informative References.....	18
Appendix Z: Change History.....	19
Acknowledgements.....	20
Authors' Addresses.....	21

INTERNET-DRAFT

TRILL: RBridge Channel Tunnel

1. Introduction

The IETF TRILL base protocol [[RFC6325](#)] has been extended with an optional RBridge Channel [[RFC7178](#)] facility to support transmission of typed messages (for example BFD [[RFC7175](#)]) between two TRILL switches (RBridges) in the same campus and between RBridges and end stations on the same link. When sent between RBridges in the same campus, a TRILL Data packet with a TRILL header is used and the destination RBridge is indicated by nickname. When sent between a RBridge and an end station on the same link in either direction a native RBridge Channel messages [[RFC7178](#)] is used with no TRILL header and the destination port or ports is indicated by a MAC address. (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.)

This document updates [[RFC7178](#)] and specifies extensions to RBridge Channel that provides two additional facilities as listed below. Implementation and use of each of these facilities is optional, except that there are two payload types that **MUST** be implemented. Both of these facilities can be used in the same packet.

- (1) A standard method to tunnel a variety of payload types by encapsulating them in an RBridge Channel message.
- (2) A method to provide security facilities for RBridge Channel messages.

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 acronyms defined in [[RFC6325](#)] and [[RFC7178](#)] supplemented by the following additional acronym:

Data Label - VLAN or FGL.


FGL - Fine Grained Label [[RFC7172](#)].

RBridge - An alternative term for a TRILL switch.

TRILL switch - A device that implements the TRILL protocol [[RFC6325](#)], sometimes referred to as 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, inner Ethernet addresses, and Data Label are omitted as shown in Figure 2. The type of RBridge Channel message 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




Figure 2. Native RBridge Channel Frame

The RBridge Channel Header looks like this:

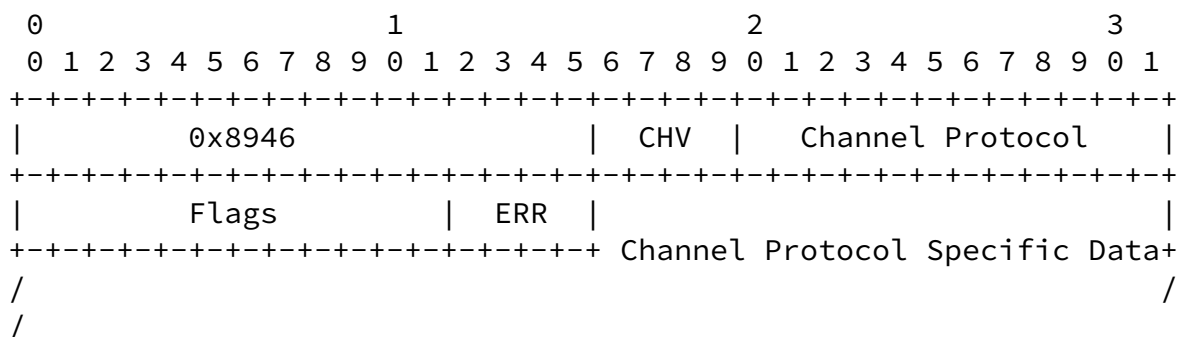


Figure 3. RBridge Channel Header

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.

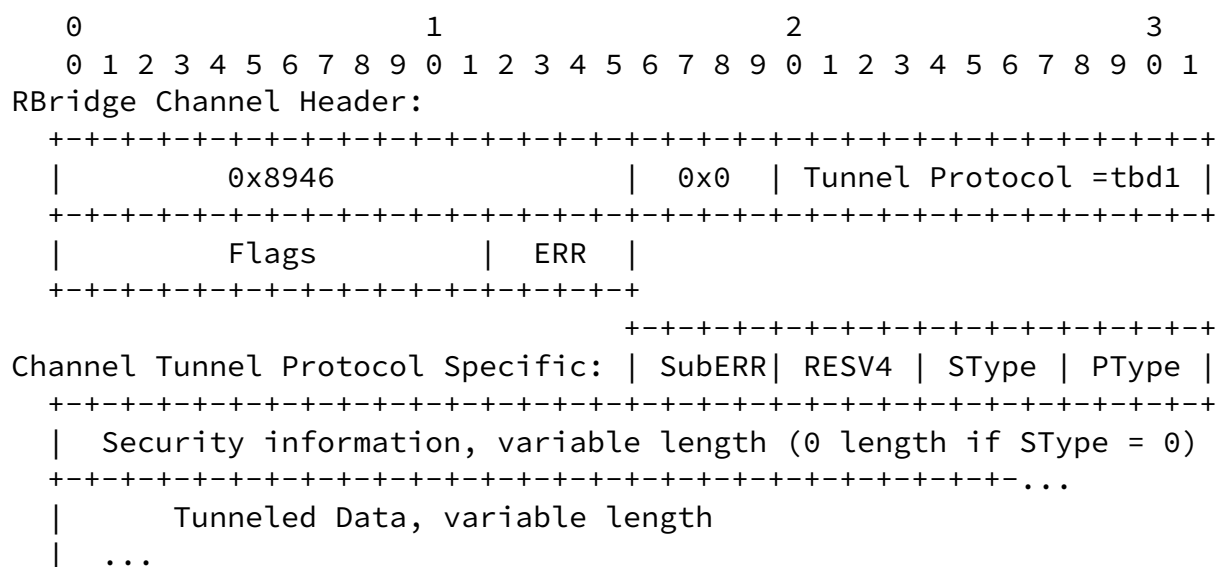


Figure 4. Channel Tunnel Header Structure

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 6](#)).

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 5](#).

RESV4: This field MUST be sent as zero. If non-zero when received,

this is an error condition (see [Section 4](#)).

SType: This field describes the type of security information and features, including keying material, being provided. See [Section 4](#).

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 (see [Section 6](#)).

[3. Tunnel Payload Types](#)

The RBridge Channel Tunnel Protocol can carry a variety of payloads as indicated by the PType field. Values 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 testing or messages such as key negotiation or the like. It indicates that there is no payload. Any data after the 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 payload type looks like the following:

INTERNET-DRAFT

TRILL: RBridge Channel Tunnel

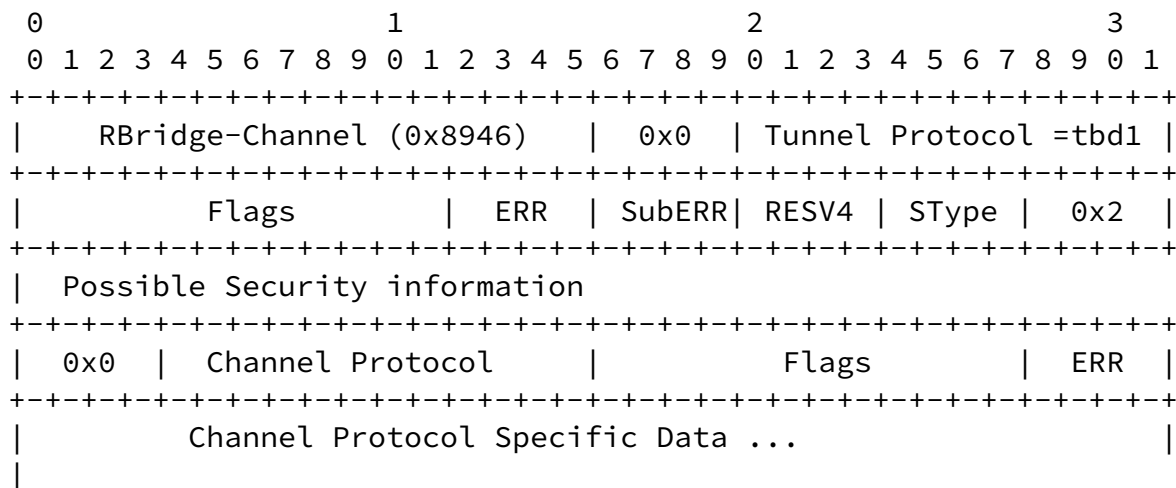


Figure 5. Tunneled Channel Message Channel Tunnel Structure

3.3 TRILL Data Packet

A PType of 3 indicates that the payload of the Tunnel protocol message is an encapsulated TRILL Data packet as shown in the figure below. (There is no TRILL Ethertype before the inner TRILL Data packet because that is just part of the Ethernet link header for a TRILL Data packet, not part of the TRILL header itself.) If this PType is implemented and the message meets local policy for acceptance, 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.

INTERNET-DRAFT

TRILL: RBridge Channel Tunnel

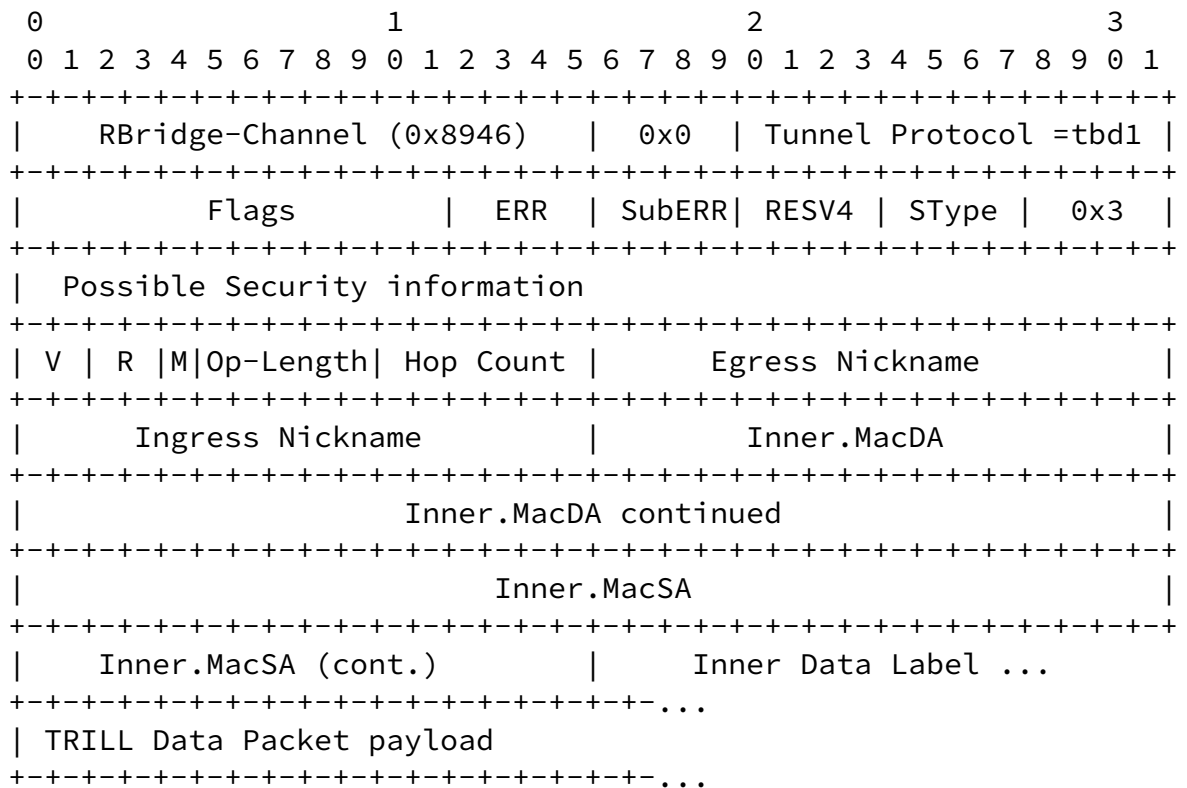


Figure 6. Nested TRILL Data Packet Channel Tunnel Structure

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 PDU 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. One possible use is to expedite the receipt of a link state PDU by some TRILL switch or

switches with an immediate requirement for the enclosed link state PDU. Any link local IS-IS PDU (Hello, CSNP, or PSNP [IS-IS]; MTU-probe, MTU-ack [RFC7176]; or circuit scoped FS-LSP, FS-CSNP or FS-PSNP [RFC7356]) received via this channel tunnel payload type MUST be discarded.

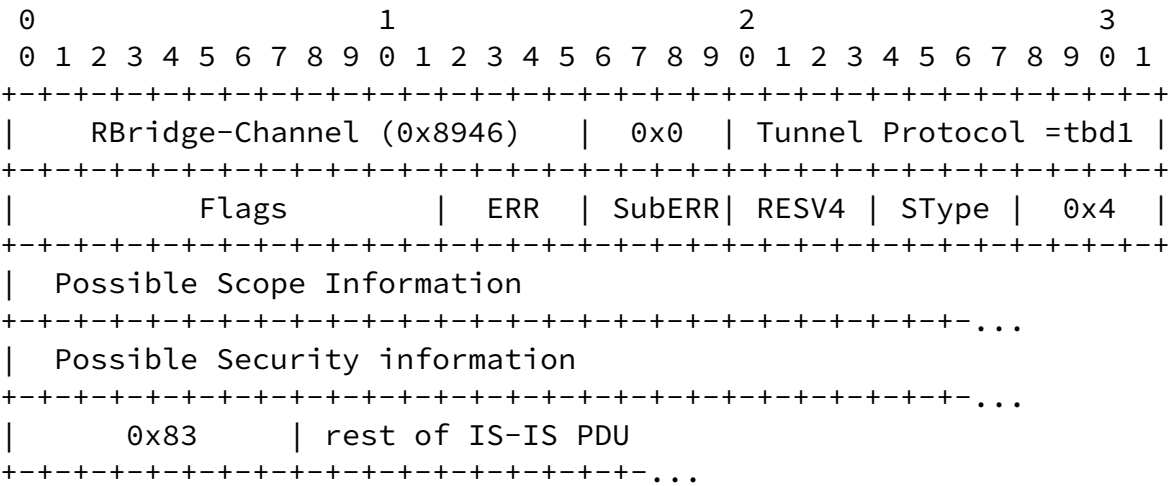


Figure 7. Tunneled TRILL IS-IS Packet Structure

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 and the message meets local policy, the tunneled frame is handled as if it

had been received on the port on which the Tunnel Protocol message was received.

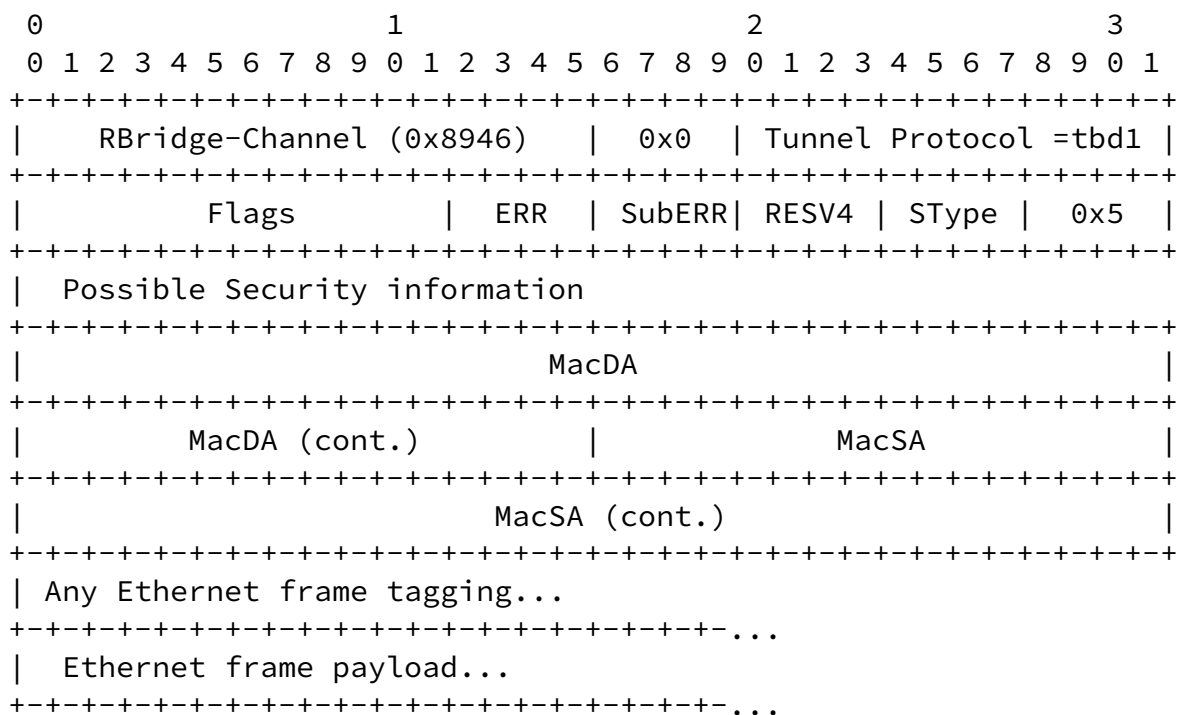


Figure 8. Ethernet Frame Channel Tunnel Structure

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. This constructed address MAY be used as the MacSA and, if the RBridge Channel message is link local, the source TRILL switch will have the information to construct such a MAC address for the destination TRILL switch port and that MAC address MAY be used as the MacDA.

These MAC addresses are constructed as follows: 0xFEFF, the nickname of the TRILL switch used in TRILL Hellos on that port, and the Port ID that the TRILL switch has assigned to that port, as shown in Figure 9. Both the nickname and Port ID appear in the Special VLANs and Flags sub-TLV [[RFC7176](#)]. The resulting MAC address has the Local bit on and the Group bit off [[RFC7042](#)]. Since end stations are connected to TRILL switches over Ethernet, there will 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 9. Synthetic MAC Address

4. Security, Keying, and Algorithms

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

SType	Section	Meaning
0	4.2	None
1	4.3	[RFC5310] Based Authentication
2	4.4	xxx Based Security

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:




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 bytes, as an unsigned integer, of Info in the Security Information after the Size byte itself.

Info: Variable length Security Information.

[4.1 Authentication Coverage](#)

Authentication is computed across relevant Channel Tunnel header information and payload with the area when their authentication value is to be stored set to zero. To be more precise, the covered area starts just after the first byte of the RBridge Channel header flags and extends to just before the link trailer. Thus RBridge Channel header flags bits 8 through 11 are protected by authentication while flag bits 0 through 7 are not.

[4.2 SType None](#)

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

4.3 [RFC 5310](#) Based Authentication

The Security Information (see Figure 6) is the flags and Size bytes specified above together with the value of the [[RFC5310](#)] Key ID and Authentication Data as shown in Figure 13.

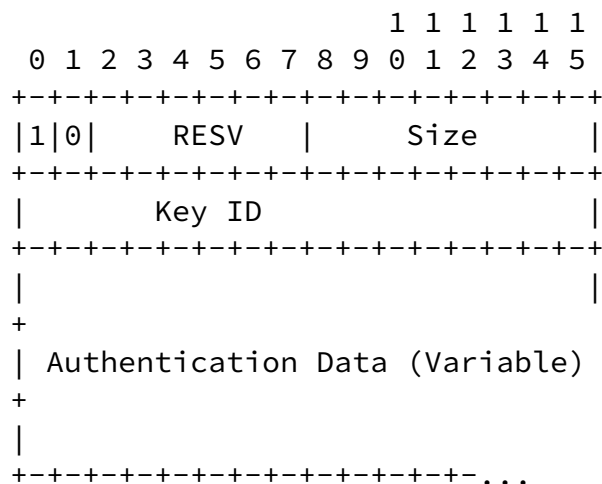


Figure 13. SType 1 Security Information

The Key ID specifies the same keying value and algorithm that Key ID specifies for core TRILL IS-IS LSP Authentication TLVs; however, to avoid using the same key for multiple purposes, the keying value actually used for Tunnel Channel messages with this SType is a value derived from that TRILL IS-IS value as follows:

HMAC-SHA256 (("Channel Tunnel" | Data Label), IS-IS-key)

where "|" indicates concatenation, HMAC-SHA256 is as in [[FIPS-180](#)] [[RFC6234](#)], "Channel Tunnel" is the 14 character [[ASCII](#)] string indicated, and Data Label is the Data Label in which Channel Tunnel message is being sent as either (1) a right justified 12-bit VLAN ID in two bytes with 4 zero high order bits or (2) three bytes of FGL.

4.4 xxx Based Security

xxx - permits key negotiation, provides both encryption and authentication ...

5. 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
8	Error in nested RBridge Channel message
(more TBD?)	

Table 4. Additional ERR Values

5.1 SubERRs under ERR 6

If the ERR field is 6, the SubERR field indicates the problematic field or value as show in the table below.

SubERR	Meaning (for ERR = 6)
-----	-----
0	Non-zero RESV4 nibble
1	Unsupported SType
2	Unsupported PType
3	Unknown or reserved Scope Egress Nickname in an ESRB scope Tunnel Channel message
4	Unsupported crypto algorithm
5	Unknown Key ID for SType 1
(more TBD)	

Table 5. SubERR values under ERR 6

5.2 Nested RBridge Channel Errors

If

a Channel Tunnel message is sent with security and with a payload type (PType) indicating a nested RBridge Channel message

and

there is an error in the processing of that nested message that results in a return RBridge Channel message with a non-zero ERR field,

then that returned message SHOULD also be nested in an Channel Tunnel message using the same type of security. In this case, the ERR field in the Channel Tunnel envelope is set to 8 indicating that there is a nested error.

[6.](#) IANA Considerations

IANA has assigned tbd1 as the RBridge Channel protocol number the "Channel Tunnel" protocol from the range assigned by Standards Action.

The added RBridge Channel protocols registry entry on the TRILL Parameters web page is as follows:

Protocol	Description	Reference
-----	-----	-----
tbd1	Tunnel Channel	[this document]

INTERNET-DRAFT

TRILL: RBridge Channel Tunnel

[7](#). 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 RBridge Channel messages. Some RBridge Channel message payloads provide their own security [[RFC7175](#)] but where this is not true, consideration should be give to requiring use of the security features of the Tunnel Protocol.

On the negative side, the optional 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 taking. 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 [[RFC7178](#)] for general RBridge Channel Security Considerations.

See [[RFC6325](#)] for general TRILL Security Considerations.

INTERNET-DRAFT

TRILL: RBridge Channel Tunnel

Normative References

- [ASCII] - American National Standards Institute (formerly United States of America Standards Institute), "USA Code for Information Interchange", ANSI X3.4-1968, 1968. ANSI X3.4-1968 has been replaced by newer versions with slight modifications, but the 1968 version remains definitive for the Internet.
- [FIPS-180] - "Secure Hash Standard (SHS)", United States of American, National Institute of Science and Technology, Federal Information Processing Standard (FIPS) 180-4, March 2012, <http://csrc.nist.gov/publications/fips/fips180-4/fips-180-4.pdf>
- [IS-IS] - ISO/IEC 10589:2002, Second Edition, "Information technology -- Telecommunications and information exchange between systems -- Intermediate System to Intermediate System intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)", 2002.
- [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, "RBriges: 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](#), May 2014.
- [RFC7176] - Eastlake 3rd, D., Senevirathne, T., Ghanwani, A., Dutt, D., and A. Banerjee, "Transparent Interconnection of Lots of Links (TRILL) Use of IS-IS", [RFC 7176](#), May 2014, <<http://www.rfc-editor.org/info/rfc7176>>.
- [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.
- [RFC7356] - Ginsberg, L., Previdi, S., and Y. Yang, "IS-IS Flooding Scope Link State PDUs (LSPs)", [RFC 7356](#), September 2014, <<http://www.rfc-editor.org/info/rfc7356>>.

Informative References

- [RFC6234] - Eastlake 3rd, D. and T. Hansen, "US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF)", [RFC 6234](#), May 2011.
- [RFC6361] - Carlson, J. and D. Eastlake 3rd, "PPP Transparent Interconnection of Lots of Links (TRILL) Protocol Control Protocol", [RFC 6361](#), August 2011
- [RFC7042] - Eastlake 3rd, D. and J. Abley, "IANA Considerations and IETF Protocol and Documentation Usage for IEEE 802 Parameters", [BCP 141](#), [RFC 7042](#), October 2013.

[RFC7175] - Manral, V., Eastlake 3rd, D., Ward, D., and A. Banerjee,
"Transparent Interconnection of Lots of Links (TRILL):
Bidirectional Forwarding Detection (BFD) Support", [RFC 7175](#),
May 2014.

Appendix Z: Change History

From -00 to -01

1. Fix references for RFCs published, etc.

2. Explicitly mention in the Abstract and Introduction that this document updates [[RFC7178](#)].
3. Add this Change History Appendix.

From -01 to -02

1. Remove section on the "Scope" feature as mentioned in <http://www.ietf.org/mail-archive/web/trill/current/msg06531.html>
2. Editorial changes to IANA Considerations to correspond to [draft-leiba-cotton-iana-5226bis-11.txt](#).
- x. Other Editorial changes.

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.

INTERNET-DRAFT

TRILL: RBridge Channel Tunnel

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

INTERNET-DRAFT

TRILL: RBridge Channel Tunnel

Copyright, Disclaimer, and Additional IPR Provisions

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. 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.

