Internet Engineering Task Force (IETF)
Request for Comments: 7978
Updates: 7178
Category: Standards Track
ISSN: 2070-1721

D. Eastlake 3rd
Huawei
M. Umair
IPinfusion
Y. Li
Huawei

September 2016

# Transparent Interconnection of Lots of Links (TRILL): RBridge Channel Header Extension

## Abstract

The IETF TRILL (Transparent Interconnection of Lots of Links) protocol includes an optional mechanism (specified in RFC 7178) called RBridge Channel for the transmission of typed messages between TRILL switches in the same campus and the transmission of such messages between TRILL switches and end stations on the same link. This document specifies extensions to the RBridge Channel protocol header to support two features as follows: (1) a standard method to tunnel payloads whose type can be indicated by Ethertype through encapsulation in RBridge Channel messages; and (2) a method to support security facilities for RBridge Channel messages. This document updates RFC 7178.

## Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7978.

Table of Contents

## 1. Introduction

The IETF TRILL base protocol [RFC6325] [RFC7780] has been extended with the RBridge Channel [RFC7178] facility to support transmission of typed messages (for example, Bidirectional Forwarding Detection (BFD) [RFC7175]) between two TRILL switches (RBridges) in the same campus and the transmission of such messages 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 message [RFC7178] is used with no TRILL Header, and the destination port or ports are indicated by a Media Access Control (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 the RBridge Channel header that provide two additional facilities as follows:

(1) A standard method to tunnel payloads, whose type may be indicated by Ethertype, through encapsulation in RBridge Channel messages.

(2) A method to provide security facilities for RBridge Channel messages. Example uses requiring such facilities are the security of Pull Directory messages [RFC7067], address flush messages [AddrFlush], and port shutdown messages [TRILL-AF].

Use of each of these facilities is optional, except that, as specified below, if this header extension is implemented, there are two payload types that MUST be implemented. Both of the above facilities can be used in the same packet. In case of conflict between this document and [RFC7178], this document takes precedence.

### 1.1. Terminology and Acronyms

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

This document uses terminology and abbreviations defined in [RFC6325] and [RFC7178]. Some of these are listed below for convenience along with new terms and abbreviations.

application_data - A DTLS [RFC6347] message type.

Data Label - VLAN or FGL.

DTLS - Datagram Transport Layer Security [RFC6347].

FCS - Frame Check Sequence.

FGL - Fine-Grained Label [RFC7172].

HKDF - HMAC-based Key Derivation Function [RFC5869].

IS-IS - Intermediate System to Intermediate System [IS-IS].

PDU - Protocol Data Unit.

MTU - Maximum Transmission Unit.

RBridge - An alternative term for a TRILL switch.

SHA - Secure Hash Algorithm [RFC6234].

Sz - Campus-wide minimum link MTU [RFC6325] [RFC7780].

TRILL - Transparent Interconnection of Lots of Links or Tunneled Routing in the Link Layer.

TRILL switch - A device that implements the TRILL protocol [RFC6325] [RFC7780], sometimes referred to as an RBridge.

## 2. RBridge Channel Header Extension Format

The general structure of an RBridge Channel message between two TRILL switches (RBridges) in the same campus is shown in Figure 1 below. The structure of a native RBridge Channel message sent between an RBridge and an end station on the same link, in either direction, is shown in Figure 2 and, compared with the first case, omits the TRILL Header, inner Ethernet addresses, and Data Label. A Protocol field in the RBridge Channel Header gives the type of RBridge Channel message and indicates how to interpret the Channel-Protocol-Specific Payload [RFC7178].

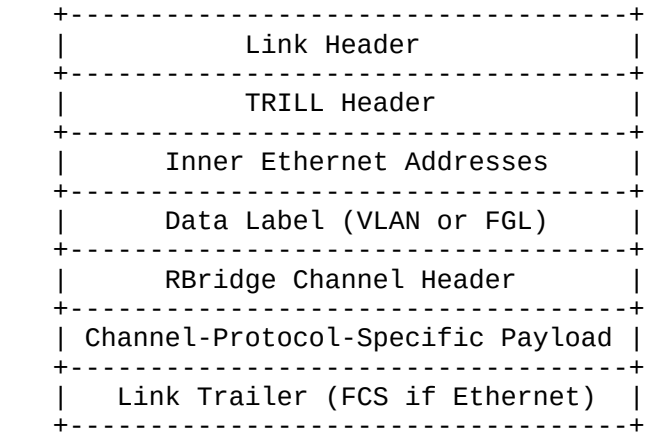


Figure 1: RBridge Channel Packet Structure

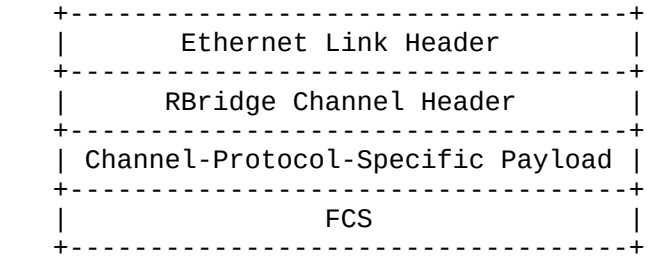


Figure 2: Native RBridge Channel Frame

The RBridge Channel Header looks like this:

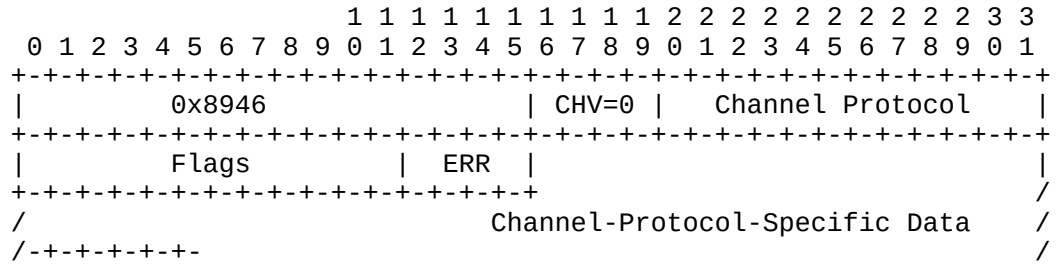


Figure 3: RBridge Channel Header

where 0x8946 is the RBridge-Channel Ethertype and CHV is the Channel Header Version. This document is based on RBridge Channel version zero.

The header extensions specified herein are in the form of an RBridge Channel protocol, the Extended RBridge Channel Protocol. Figure 4 below expands the RBridge Channel Header and Protocol-Specific Payload above for the case where the header extension is present.

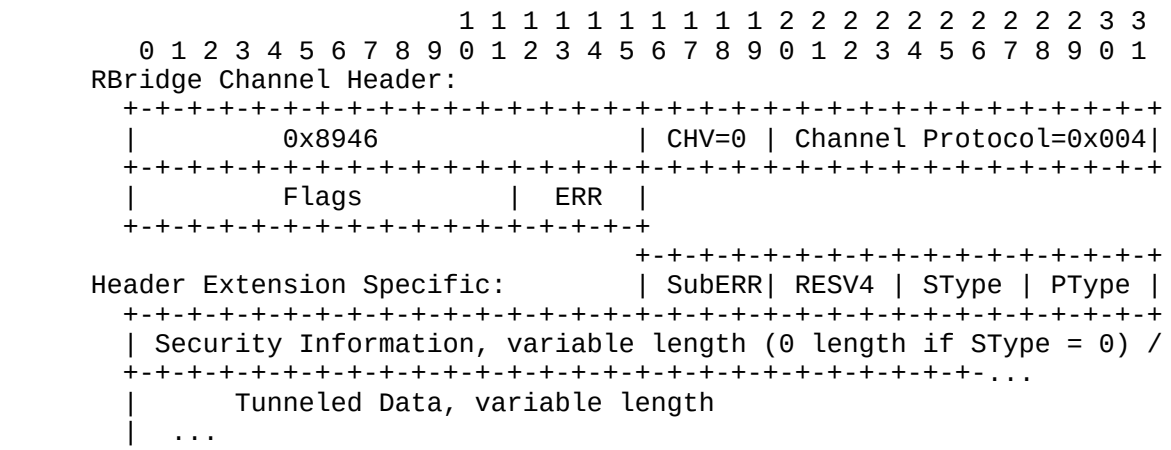


Figure 4: RBridge Channel Header Extension Structure

The RBridge Channel Header Protocol field is used to indicate that the header extension is present. Its contents MUST be the value allocated for this purpose (see Section 6). The use of an RBridge Channel protocol to indicate extensions makes it easy to determine if a remote RBridge in the campus supports extensions since RBridges advertise in their LSP which such protocols they support.

The Extended RBridge Channel-Protocol-Specific Data fields are as follows:

SubERR: This field provides further details when an 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 5.

SType: This field describes the type of security information and features, including keying material, being used or provided by the extended RBridge Channel message. See Section 4.

PType: Payload Type. This describes the tunneled data. See Section 3.

Security Information: Variable-length information. Length is zero if SType is zero. See Section 4.

The RBridge Channel Header Extension is integrated with the RBridge Channel facility. Extension 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.

## 3. Extended RBridge Channel Payload Types

The Extended RBridge Channel Protocol can carry a variety of payloads as indicated by the PType (Payload Type) field. Values are shown in the table below with further explanation below the table (see also Section 6.2.2).

| PType | Description | Reference |
|---|---|---|
| 0 | Reserved | |
| 1 | Null | Section 3.1 of RFC 7978 |
| 2 | Ethertyped Payload | Section 3.2 of RFC 7978 |
| 3 | Ethernet Frame | Section 3.3 of RFC 7978 |
| 4-14 | Unassigned | |
| 15 | Reserved | |

Table 1: Payload Type Values

While implementation of the RBridge Channel Header Extension is optional, if it is implemented, PType 1 (Null) MUST be implemented and PType 2 (Ethertyped Payload) with the RBridge-Channel Ethertype MUST be implemented. PType 2 for any Ethertypes other than the RBridge-Channel Ethertype MAY be implemented. PType 3 MAY be implemented.

The processing of any particular extended header RBridge Channel 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 (PType = 1) is intended to be used for testing or for messages such as key negotiation or the like where only security information is present. It indicates that there is no user data payload. Any tunneled user data after the Security Information field is ignored. If the RBridge Channel Header Extension is implemented, the Null Payload MUST be supported in the sense that an "Unsupported PType" error is not returned (see Section 5). Any particular use of the Null Payload should specify what VLAN or FGL

and what priority should be used in the inner Data Label of the RBridge Channel message (or in an outer VLAN tag for the native RBridge Channel message case) when those values are relevant.

## 3.2. Ethertyped Payload

A PType of 2 indicates that the payload (tunneled data) of the extended RBridge Channel message begins with an Ethertype. A TRILL switch supporting the RBridge Channel Header Extension MUST support a PType of 2 with a payload beginning with the RBridge-Channel Ethertype as described in Section 3.2.1. Other Ethertypes, including the TRILL and L2-IS-IS Ethertypes as described in Sections 3.2.2 and 3.2.3, MAY be supported.

### 3.2.1. RBridge Channel Message as the Payload

A PType of 2 whose payload has an initial RBridge-Channel Ethertype indicates an encapsulated RBridge Channel message. A typical reason for sending an RBridge Channel message inside an extended RBridge Channel message is to provide security services, such as authentication or encryption, for the encapsulated message.

This RBridge Channel message type looks like the following:

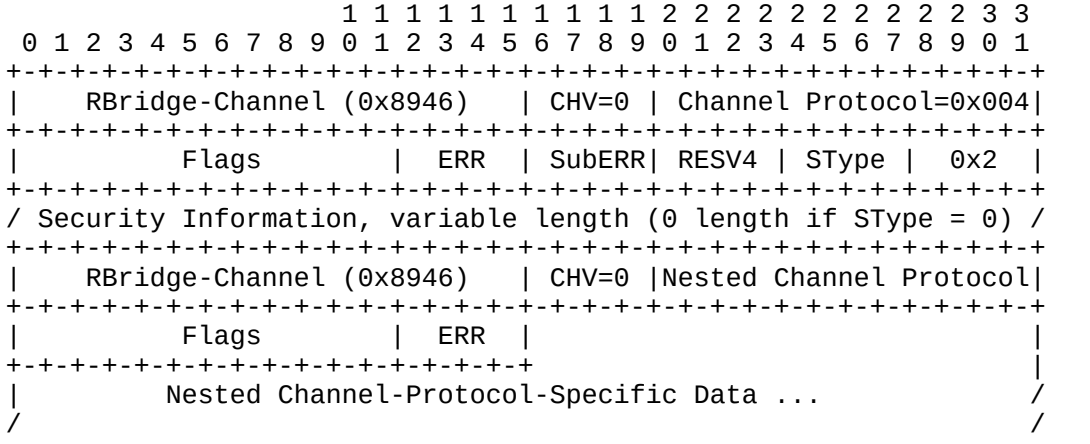


Figure 5: Message Structure with RBridge Channel Payload

### 3.2.2. TRILL Data Packet as the Payload

A PType of 2 whose payload has an initial TRILL Ethertype indicates an encapsulated TRILL Data packet as shown in Figure 6. If this Ethertype is supported for PType = 2 and the message meets local policy for acceptance, the TRILL Data packet is handled as if it had been received by the destination TRILL switch on the port where the Extended RBridge Channel message was received.

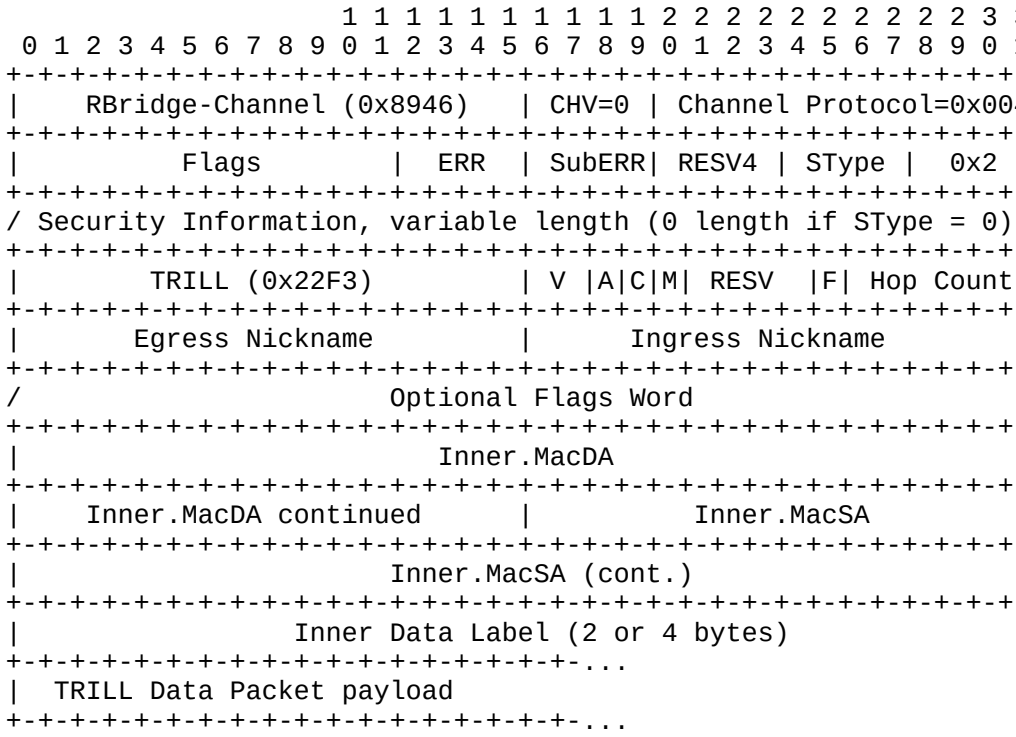


Figure 6: Message Structure with TRILL Data Packet Payload

The optional flags word is only present if the F bit in the TRILL Header is one [RFC7780].

### 3.2.3. TRILL IS-IS Packet as the Payload

A PType of 2 and an initial L2-IS-IS Ethertype indicate that the payload of the Extended RBridge Channel protocol message is an encapsulated TRILL IS-IS PDU as shown in Figure 7. If this Ethertype is supported for PType = 2, 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 (LSP) by

some TRILL switch or switches with an immediate requirement for the link state information. A link local IS-IS PDU would not normally be sent via this Extended RBridge Channel method except possibly to encrypt the PDU since such PDUs can just be transmitted on the link and do not normally need RBridge Channel handling. (Link local IS-IS PDUs are (1) Hello, CSNP, PSNP [IS-IS]; (2) MTU-probe, MTU-ack [RFC7176]; and (3) circuit scoped FS-LSP, FS-CSNP, and FS-PSNP [RFC7356].)

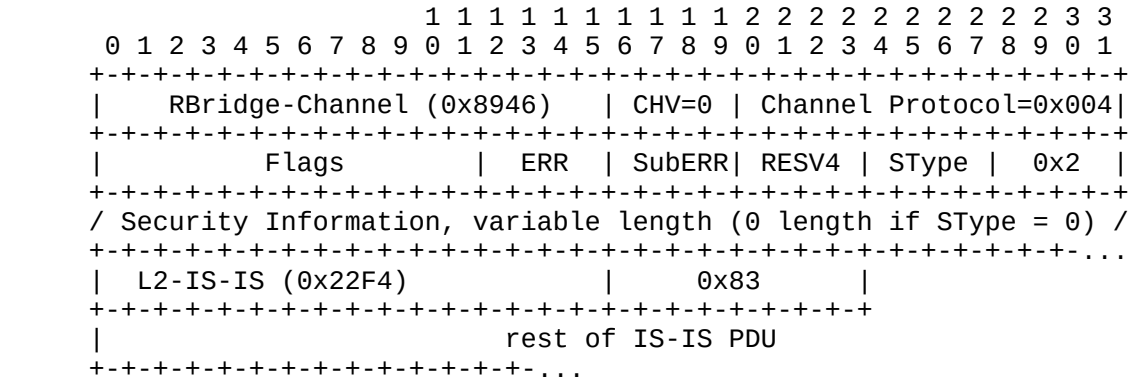


Figure 7: Message Structure with TRILL IS-IS Packet Payload

## 3.3. Ethernet Frame

If PType is 3, the extended RBridge Channel payload is an Ethernet frame as might be received from or sent to an end station except that the encapsulated Ethernet frame's FCS is omitted, as shown in Figure 8. (There is still an overall final 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 encapsulated frame is handled as if it had been received on the port on which the Extended RBridge Channel message was received.

The priority of the RBridge Channel message can be copied from the Ethernet frame VLAN tag, if one is present, except that priority 7 SHOULD only be used for messages critical to establishing or maintaining adjacency and priority 6 SHOULD only be used for other important control messages.

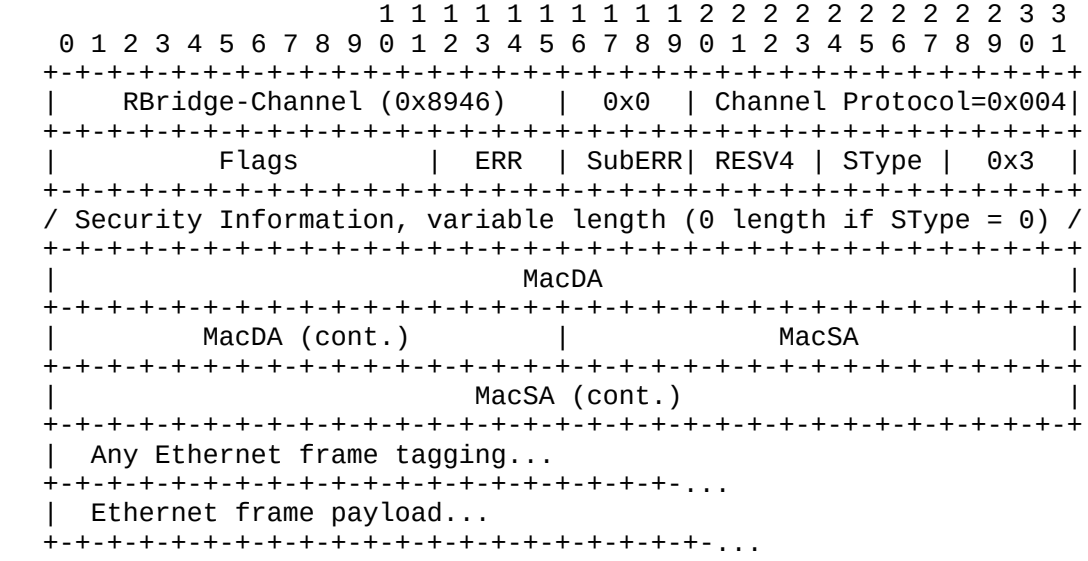


Figure 8: Message Structure with Ethernet Frame Payload

In the case of a non-Ethernet link, such as a PPP (Point-to-Point Protocol) link [RFC6361], the ports on the link are considered to have link-local synthetic 48-bit MAC addresses constructed as described below. Such a constructed address MAY be used as a MacSA. If the RBridge Channel message is individually addressed to a link-local port, 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. By the use of such a MacSA and either such a unicast MacDA or a group-addressed MacDA, an Ethernet frame can be sent between two TRILL switch ports connected by a non-Ethernet link.

These synthetic TRILL switch port MAC addresses for non-Ethernet ports are constructed as follows (and as shown in Figure 9): 0xFEFF, the nickname of the TRILL switch used in TRILL Hellos sent on that port, and the Port ID that the TRILL switch has assigned to that port. (Both the Port ID of the port on which a TRILL Hello is sent and the nickname of the sending TRILL switch appear in the Special VLANs and Flags sub-TLV [RFC7176] in TRILL IS-IS Hellos.) The resulting MAC address has the Local bit on and the Group bit off [RFC7042]. However, since there will be no Ethernet end stations on a non-Ethernet link in a TRILL campus, such synthetic MAC addresses cannot conflict on the link with a real Ethernet port address regardless of their values.

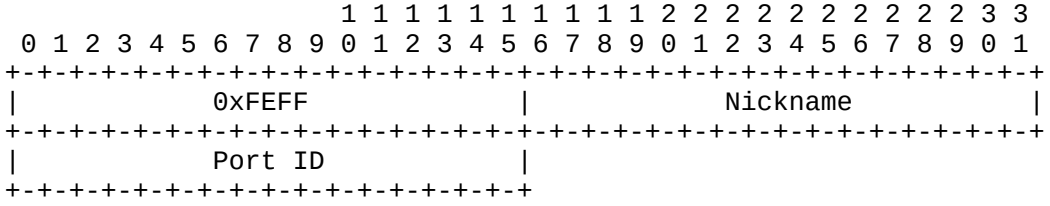


Figure 9: Synthetic MAC Address

## 4. Extended RBridge Channel Security

Table 2 below gives the assigned values of the SType (Security Type) field and their meaning. Use of DTLS Pairwise Security (SType = 2) or Composite Security (SType = 3) is RECOMMENDED.

While IS-IS CRYPTO_AUTH-based authentication is also specified and can be used for both pairwise and multi-destination traffic, it provides only authentication and is not considered to meet current security standards. For example, it does not provide for key negotiation; thus, its use is NOT RECOMMENDED.

The Extended RBridge Channel DTLS-based security specified in Section 4.4 and the Composite Security specified in Section 4.5 are intended for pairwise (known unicast) use. That is, the case where the M bit in the TRILL Header is zero and any Outer.MacDA is individually addressed.

Multi-destination Extended RBridge Channel packets would be those with the M bit in the TRILL Header set to one or, in the native RBridge Channel case, the Outer.MacDA would be group addressed. The DTLS Pairwise Security and Composite Security STypes can also be used in the multi-destination case by serially unicasting the messages to all data-accessible RBridges (or stations in the native RBridge Channel case) in the recipient group. For TRILL Data packets, that group is specified by the Data Label; for native frames, the group is specified by the groupcast destination MAC address. It is intended to specify a true group keyed SType to secure multi-destination packets in a separate document [GroupKey].

| SType | Description | Reference |
|---|---|---|
| 0 | None | Section 4.2 of RFC 7978 |
| 1 | IS-IS CRYPTO_AUTH-Based Authentication | Section 4.3 of RFC 7978 |
| 2 | DTLS Pairwise Security | Section 4.4 of RFC 7978 |
| 3 | Composite Security | Section 4.5 of RFC 7978 |
| 4-14 | Unassigned | |
| 15 | Reserved | |

Table 2: SType Values

## 4.1. Derived Keying Material

In some cases, it is possible to use material derived from IS-IS CRYPTO_AUTH keying material [RFC5310] as an element of Extended RBridge Channel security. It is assumed that the IS-IS keying material is of high quality. The material actually used is derived from the IS-IS keying material as follows:

```
Derived Material =
   HKDF-Expand-SHA256 ( IS-IS-key, "Extended Channel" | 0x0S, L )
```

where "|" indicates concatenation, HKDF is as in [RFC5869], SHA256 is as in [RFC6234], IS-IS-key is the input IS-IS keying material, "Extended Channel" is the 16-character ASCII [RFC20] string indicated without any leading length byte or trailing zero byte, 0x0S is a single byte where S is the SType for which this key derivation is being used and the upper nibble is zero, and L is the length of the output-derived material needed.

Whenever IS-IS keying material is being used as above, the underlying IS-IS CRYPTO_AUTH keying material [RFC5310] might expire or be invalidated. At the time of or before such expiration or invalidation, the use of the Derived Material from the IS-IS keying material MUST cease. Continued security MAY use new derived material from currently valid IS-IS CRYPTO_AUTH keying material.

## 4.2. SType None

No security services are being invoked. The length of the Security Information field (see Figure 4) is zero.

### 4.3. IS-IS CRYPTO_AUTH-Based Authentication

This SType provides security for Extended RBridge Channel messages similar to that provided for [IS-IS] PDUs by the [IS-IS] Authentication TLV. The Security Information (see Figure 4) is as shown in Figure 10.

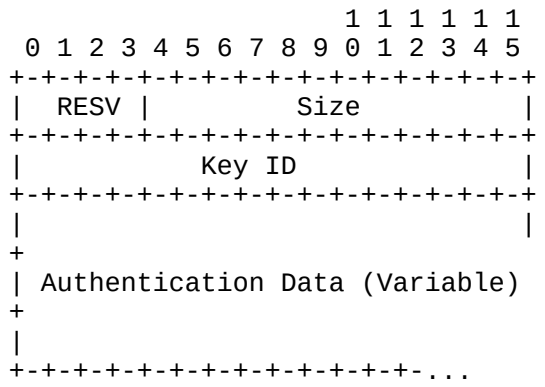


Figure 10: SType 1 Security Information

- RESV: Four bits that MUST be sent as zero and ignored on receipt.
- Size: Set to 2 + the size of Authentication Data in bytes.
- Key ID: specifies the keying value and authentication algorithm that the Key ID specifies for TRILL IS-IS LSP [RFC5310] Authentication TLVs. The keying material actually used is always derived as shown in Section 4.1.
- Authentication Data: The authentication data produced by the derived key and algorithm associated with the Key ID acting on the part of the TRILL Data packet shown. Length of the authentication data depends on the algorithm. The authentication value is included in the security information field and is treated as zero when authentication is calculated.

As show in Figure 11, the area covered by this authentication starts with the byte immediately after the TRILL Header optional Flag Word if it is present. If the Flag Word is not present, it starts after the TRILL Header Ingress Nickname. In either case, it extends to just before the TRILL Data packet link trailer. For example, for an Ethernet packet it would extend to just before the FCS.

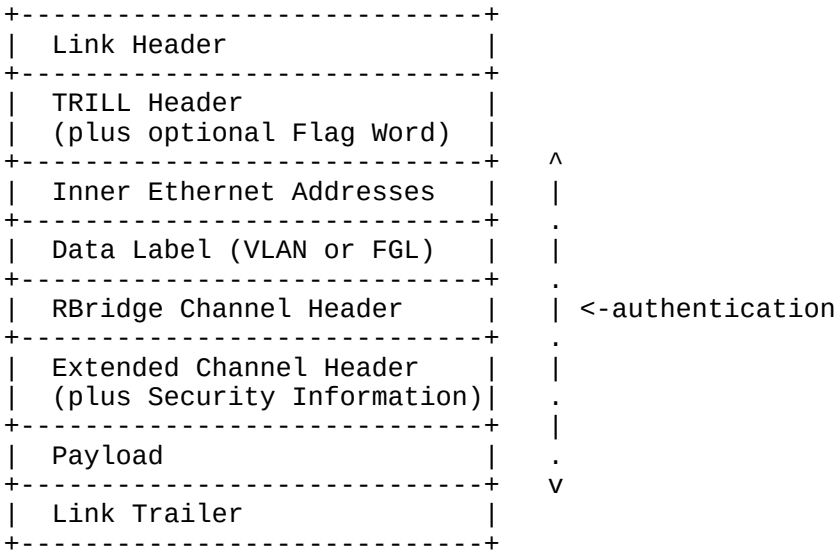


Figure 11: SType 1 Authentication Coverage

In the native RBridge Channel case, this authentication coverage is as specified in the above paragraph except that it starts with the RBridge-Channel Ethertype, since there is no TRILL Header, inner Ethernet addresses, or inner Data Label (see Figure 12).

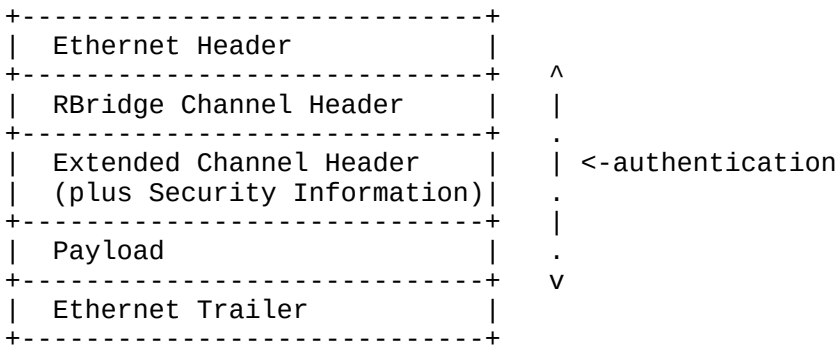


Figure 12: Native SType 1 Authentication Coverage

RBridges, which are IS-IS routers, can reasonably be expected to hold IS-IS CRYPTO_AUTH keying material [RFC5310] so that this SType can be used for RBridge Channel messages, which go between RBridges. How end stations might come to hold IS-IS CRYPTO_AUTH keying material is

beyond the scope of this document. Thus, this SType might not be applicable to native RBridge Channel messages, which are between an RBridge and an end station.

### 4.4. DTLS Pairwise Security

DTLS [RFC6347] supports key negotiation and provides both encryption and authentication. The RBridge Channel Extended Header DTLS Pairwise SType uses a negotiated DTLS version that MUST NOT be less than 1.2.

When DTLS pairwise security is used, the entire payload of the Extended RBridge Channel packet, starting just after the null Security Information and ending just before the link trailer, is one or more DTLS records [RFC6347]. As specified in [RFC6347], DTLS records MUST be limited by the path MTU, in this case so that each record fits entirely within a single Extended RBridge Channel message. A minimum path MTU can be determined from the TRILL campus minimum MTU Sz, which will not be less than 1470 bytes, by allowing for the TRILL Data packet, extended RBridge Channel, and DTLS framing overhead. With this SType, the security information between the extended RBridge Channel header and the payload is null because all the security information is in the payload area.

The DTLS Pairwise keying is set up between a pair of RBridges, independent of Data Label, using messages of a priority configurable at the RBridge level, which defaults to priority 6. DTLS message types other than application_data can be the payload of an extended RBridge Channel message with a TRILL Header using any Data Label, and, for such DTLS message types, the PType in the RBridge Channel Header Extension is ignored.

Actual application_data sent within such a message using this SType SHOULD use the Data Label and priority as specified for that application_data. In this case, the PType value in the RBridge Channel Header Extension applies to the decrypted application_data.

TRILL switches that implement the extended RBridge Channel DTLS Pairwise SType SHOULD support the use of certificates for DTLS, but certificate size may be limited by the DTLS requirement that each record fit within a single message. Appropriate certificate contents are out of scope for this document.

TRILL switches that support the extended RBridge Channel DTLS Pairwise SType MUST support the use of pre-shared keys. If the psk_identity (see [RFC4279]) is two bytes, it is interpreted as a Key ID as defined in [RFC5310], and the value derived as shown in Section 4.1 from that key is used as a pre-shared key for DTLS

negotiation. A psk_identity with a length other than two bytes MAY be used to indicate other implementation-dependent pre-shared keys. Pre-shared keys used for DTLS negotiation SHOULD be shared only by the pair of endpoints; otherwise, security could be attacked by diverting messages to another endpoint holding that pre-shared key.

### 4.5. Composite Security

Composite Security (SType = 3) is the combination of DTLS Pairwise Security and IS-IS CRYPTO_AUTH-Based Authentication. On transmission, the DTLS record or records to be sent are secured as specified in Section 4.4 then used as the payload for the application of Authentication as specified in Section 4.3. On reception, the IS-IS CRYPTO_AUTH-based authentication is verified first and an error is returned if it fails. If the IS-IS CRYPTO_AUTH-based authentication succeeds, then the DTLS record or records are processed.

An advantage of Composite Security is that the payload is authenticated and encrypted with a modern security protocol; in addition, the RBridge Channel Header and (except in the native case) preceding the MAC addresses and Data Label are provided with some authentication.

## 5. Extended RBridge Channel Errors

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

| Value | RBridge Channel Error Code Meaning |
|---|---|
| 6 | Unknown or unsupported field value |
| 7 | Authentication failure |
| 8 | Error in nested RBridge Channel message |

Table 3: Additional ERR Values

### 5.1. SubERRs

If the ERR field is 6, the SubERR field indicates the problematic field or value as shown in the table below. At this time no suberrror codes are assigned under any other ERR field value.

| Err | SubERR | Meaning (for ERR = 6) |
|---|---|---|
| 0 | | No Error; suberrors not allowed |
| 1-5 | | (no suberrors assigned) |
| 6 | 0 | Reserved |
| 6 | 1 | Non-zero RESV4 nibble |
| 6 | 2 | Unsupported SType |
| 6 | 3 | Unsupported PType |
| 6 | 4 | Unknown Key ID |
| 6 | 5 | Unsupported Ethertype with PType = 2 |
| 6 | 6 | Unsupported authentication algorithm for SType = 1 |
| 6 | 7 | Non-zero SubERR with zero ERR field |
| 7-14 | | (no suberrors assigned) |
| 15 | | Reserved |

Table 4: SubERR Values

### 5.2. Secure Nested RBridge Channel Errors

If
- an extended RBridge Channel message is sent with security and with a payload type (PType) indicating an Ethertyped payload and the Ethertype indicates 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 extended RBridge Channel message using the same type of security. In this case, the ERR field in the Extended RBridge Channel envelope is set to 8 indicating that there is a nested error in the message being tunneled back.

## 6. IANA Considerations

### 6.1. Extended RBridge Channel Protocol Number

IANA has assigned 0x004 from the range assigned by Standards Action [RFC5226] as the RBridge Channel protocol number to indicate RBridge Channel Header Extension.

The added "RBridge Channel Protocols" registry in the TRILL Parameters registry is as follows:

| Protocol | Description | Reference |
|---|---|---|
| 0x004 | RBridge Channel Extension | RFC 7978 |

### 6.2. RBridge Channel Protocol Subregistries

IANA has created three subregistries under the "RBridge Channel Protocols" registry as detailed in the subsections below.

#### 6.2.1. RBridge Channel Error Codes

IANA has assigned three additional code points in the "RBridge Channel Error Codes" subregistry in the "Transparent Interconnection of Lots of Links (TRILL) Parameters" registry. The additional entries are as shown in Table 3 in Section 5 and the "Reference" column value is "RFC 7978" for those rows.

#### 6.2.2. RBridge Channel SubError Codes

IANA has created a subregistry indented under the "RBridge Channel Error Codes" registry, for RBridge Channel SubError Codes. The initial contents of this subregistry are shown in Table 4 in Section 5.1 and the fourth column "Reference" includes value "RFC 7978" for all rows. The header information is as follows:

```
Registry Name: RBridge Channel SubError Codes
Registration Procedures: IETF Review
Reference: RFC 7978
```

#### 6.2.3. Extended RBridge Channel Payload Types Subregistry

IANA has created an "Extended RBridge Channel Payload Types" subregistry after the "RBridge Channel Protocols" registry in the "Transparent Interconnection of Lots of Links (TRILL) Parameters" registry. The header information is as follows:

```
Registration Procedures: IETF Review
Reference: RFC 7978
```

The initial registry content is in Table 1 in Section 3 of this document.

#### 6.2.4. Extended RBridge Channel Security Types Subregistry

IANA has created an "Extended RBridge Channel Security Types" subregistry after the "Extended RBridge Channel Payload Types" registry in the "Transparent Interconnection of Lots of Links (TRILL) Parameters" registry. The header information is as follows:

> Registration Procedures: IETF Review
> Reference: RFC 7978

The initial registry content is in Table 2 in Section 4 of this document.

## 7. Security Considerations

The RBridge Channel Header Extension 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, such as BFD [RFC7175], provide their own security but where this is not true, consideration should be given, when specifying an RBridge Channel protocol, to recommending or requiring use of the security features of the RBridge Channel Header Extension.

On the negative side, the optional ability to tunnel more payload types, and to tunnel them between TRILL switches and to and from end stations, can increase risk unless precautions are taken. The processing of decapsulated extended RBridge Channel payloads is a place where you SHOULD NOT be liberal in what you accept. This is because 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 accept payload types required and then only with adequate security for the particular circumstances.

See the first paragraph of Section 4 for recommendations on SType usage.

See [RFC7457] for security considerations of DTLS.

If IS-IS authentication is not being used, then IS-IS CRYPTO_AUTH keying material [RFC5310] would not normally be available but that presumably represents a judgment by the TRILL campus operator that no security is needed.

See [RFC7178] for general RBridge Channel security considerations and [RFC6325] for general TRILL security considerations.

## 8. Normative References

[IS-IS] International Organization for Standardization, "Information technology -- Telecommunications and information exchange between systems -- Intermediate System to Intermediate System intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)", ISO/IEC 10589:2002, Second Edition, 2002.

[RFC20] Cerf, V., "ASCII format for network interchange", STD 80, RFC 20, DOI 10.17487/RFC0020, October 1969, <http://www.rfc-editor.org/info/rfc20>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC4279] Eronen, P., Ed., and H. Tschofenig, Ed., "Pre-Shared Key Ciphersuites for Transport Layer Security (TLS)", RFC 4279, DOI 10.17487/RFC4279, December 2005, <http://www.rfc-editor.org/info/rfc4279>.

[RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, DOI 10.17487/RFC5310,v February 2009, <http://www.rfc-editor.org/info/rfc5310>.

[RFC5869] Krawczyk, H. and P. Eronen, "HMAC-based Extract-and-Expand Key Derivation Function (HKDF)", RFC 5869, DOI 10.17487/RFC5869, May 2010, <http://www.rfc-editor.org/info/rfc5869>.

[RFC6325] Perlman, R., Eastlake 3rd, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification", RFC 6325, DOI 10.17487/RFC6325, July 2011, <http://www.rfc-editor.org/info/rfc6325>.

[RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <http://www.rfc-editor.org/info/rfc6347>.

[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>.

[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, DOI 10.17487/RFC7176, 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, DOI 10.17487/RFC7178, May 2014, <http://www.rfc-editor.org/info/rfc7178>.

[RFC7356] Ginsberg, L., Previdi, S., and Y. Yang, "IS-IS Flooding Scope Link State PDUs (LSPs)", RFC 7356, DOI 10.17487/RFC7356, September 2014, <http://www.rfc-editor.org/info/rfc7356>.

[RFC7780] Eastlake 3rd, D., Zhang, M., Perlman, R., Banerjee, A., Ghanwani, A., and S. Gupta, "Transparent Interconnection of Lots of Links (TRILL): Clarifications, Corrections, and Updates", RFC 7780, DOI 10.17487/RFC7780, February 2016, <http://www.rfc-editor.org/info/rfc7780>.

## 9. Informative References

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <http://www.rfc-editor.org/info/rfc5226>.

[RFC6234] Eastlake 3rd, D. and T. Hansen, "US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF)", RFC 6234, DOI 10.17487/RFC6234, May 2011, <http://www.rfc-editor.org/info/rfc6234>.

[RFC6361] Carlson, J. and D. Eastlake 3rd, "PPP Transparent Interconnection of Lots of Links (TRILL) Protocol Control Protocol", RFC 6361, DOI 10.17487/RFC6361, August 2011, <http://www.rfc-editor.org/info/rfc6361>.

[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>.

[RFC7067] Dunbar, L., Eastlake 3rd, D., Perlman, R., and I. Gashinsky, "Directory Assistance Problem and High-Level Design Proposal", RFC 7067, DOI 10.17487/RFC7067, November 2013, <http://www.rfc-editor.org/info/rfc7067>.

[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, DOI 10.17487/RFC7175, May 2014, <http://www.rfc-editor.org/info/rfc7175>.

[RFC7457] Sheffer, Y., Holz, R., and P. Saint-Andre, "Summarizing Known Attacks on Transport Layer Security (TLS) and Datagram TLS (DTLS)", RFC 7457, DOI 10.17487/RFC7457, February 2015, <http://www.rfc-editor.org/info/rfc7457>.

[AddrFlush] Hao, W., Eastlake, D., and Y. Li, "TRILL: Address Flush Message", Work in Progress, draft-ietf-trill-address-flush-00, May 2016.

[GroupKey] Eastlake, D., "TRILL: Group Keying", Work in Progress, draft-eastlake-trill-group-keying-00, July 2016.

[TRILL-AF] Eastlake, D., Li, Y., Umair, M., Banerjee, A., and F. Hu, "TRILL: Appointed Forwarders", Work in Progress, draft-ietf-trill-rfc6439bis-03, August 2016.

## Acknowledgements

The contributions of the following are hereby gratefully acknowledged:

Stephen Farrell, Jonathan Hardwick, Susan Hares, Gayle Noble, Alvaro Retana, Yaron Sheffer, and Peter Yee.

## Authors' Addresses

Donald E. Eastlake, 3rd
Huawei Technologies
155 Beaver Street
Milford, MA 01757
United States of America

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

Mohammed Umair
IPinfusion

Email: mohammed.umair2@gmail.com

Yizhou Li
Huawei Technologies
101 Software Avenue
Nanjing 210012
China

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