

INTERNET-DRAFT
Intended Status: Proposed Standard
Updates: 7177, 7178

Margaret Cullen
Painless Security
Donald Eastlake
Mingui Zhang
Huawei
Dacheng Zhang
Alibaba
October 19, 2015

Expires: April 18, 2016

Transparent Interconnection of Lots of Links (TRILL) over IP
<draft-ietf-trill-over-ip-05.txt>

Abstract

The Transparent Interconnection of Lots of Links (TRILL) protocol supports both point-to-point and multi-access links and is designed so that a variety of link protocols can be used between TRILL switch ports. This document standardizes methods for encapsulating TRILL in IP (v4 or v6) so as to use IP as a TRILL link protocol in a unified TRILL campus. It updates RFC 7177 and updates RFC 7178.

Status of This Document

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 author or the DNSEXT mailing list <dnsext@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.....	4
2. Terminology.....	5
3. Use Cases for TRILL over IP.....	6
3.1 Remote Office Scenario.....	6
3.2 IP Backbone Scenario.....	6
3.3 Important Properties of the Scenarios.....	6
3.3.1 Security Requirements.....	7
3.3.2 Multicast Handling.....	7
3.3.3 Neighbor Discovery.....	8
4. TRILL Packet Formats.....	9
4.1 General Packet Formats.....	9
4.2 General TRILL Over IP Packet Formats.....	10
4.2.1 Without Security.....	10
4.2.2 With Security.....	10
4.3 QoS Considerations.....	11
4.4 Broadcast Links and Multicast Packets.....	12
4.5 TRILL Over IP IS-IS SubNetwork Point of Attachment.....	13
5. TRILL over IP Encapsulation Formats.....	14
5.1 Encapsulation Considerations.....	14
5.2 Encapsulation Agreement.....	15
5.3 Broadcast Link Encapsulation Considerations.....	16
5.4 Native Encapsulation.....	16
5.5 VXLAN Encapsulation.....	17
5.6 Other Encapulsations.....	18
6. Handling Multicast.....	19
7. Use of IPsec and IKEv2.....	20
7.1 Keying.....	20
7.1.1 Pairwise Keying.....	20
7.1.2 Group Keying.....	21
7.2 Mandatory-to-Implement Algorithms.....	21
8. Transport Considerations.....	22
8.1 Congestion Considerations.....	22
8.2 Recursive Ingress.....	23
8.3 Fat Flows.....	24
8.4 MTU Considerations.....	25
8.5 Middlebox Considerations.....	25
9. TRILL over IP Port Configuration.....	27
9.1 Per IP Port Configuration.....	27
9.2 Additional per IP Address Configuration.....	27
9.2.1 Native Multicast Configuration.....	27

Table of Contents (continued)

9.2.2 Serial Unicast Configuration.....	28
9.2.3 Encapsulation Specific Configuration.....	28
9.2.3.1 VXLAN Configuration.....	28
9.2.3.2 Other Encapsulation Configuration.....	29
9.2.4 Security Configuration.....	29
10. Security Considerations.....	30
10.1 IPsec.....	30
10.2 IS-IS Security.....	31
11. IANA Considerations.....	32
11.1 Port Assignments.....	32
11.2 Multicast Address Assignments.....	32
11.3 Encapsulation Method Support Indication.....	32
Normative References.....	34
Informative References.....	36
Acknowledgements.....	38

1. Introduction

TRILL switches (RBridges) are devices that implement the IETF TRILL protocol [RFC6325] [RFC7177] [rfc7180bis]. TRILL provides transparent forwarding of frames within an arbitrary network topology, using least cost paths for unicast traffic. It supports VLANs and Fine Grained Labels [RFC7172] as well as multipathing of unicast and multi-destination traffic. It uses IS-IS [RFC7176] link state routing and encapsulation with a hop count.

RBridges ports can communicate with each other over various protocols, such as Ethernet [RFC6325], pseudowires [RFC7173], or PPP [RFC6361].

This document defines a method for RBridge ports to communicate over IP (v4 or v6). TRILL over IP allows Internet-connected RBridges to form a single TRILL campus, or multiple TRILL over IP networks within a campus to be connected as a single TRILL campus via a TRILL over IP backbone.

TRILL over IP connects RBridge ports using IPv4 or IPv6 as a transport in such a way that the ports appear to TRILL to be connected by a single multi-access link. If more than two RBridge ports are connected via a single TRILL over IP link, any pair of them can communicate.

To support the scenarios where RBridges are connected via IP paths (such as over the public Internet) that are not under the same administrative control as the TRILL campus and/or not physically secure, this document specifies the use of IPsec [RFC4301] Encapsulating Security Protocol (ESP) [RFC4303] to secure such paths.

To dynamically select a mutually supported TRILL over IP encapsulation, normally one with good fast path hardware support, a method is provided for agreement between adjacent TRILL switch ports as to what encapsulation to use. This document updates [RFC7177] and [RFC7178] as described in Section 5 by making adjacency between TRILL over IP ports dependent on having a method of encapsulation in common and by redefining an interval of RBridge Channel protocol numbers to indicate encapsulation method support for TRILL over IP.

2. Terminology

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 RFC 2119 [RFC2119].

The following terms and acronyms have the meaning indicated:

DRB - Designated RBridge. The RBridge (TRILL switch) elected to be in charge of certain aspects of a TRILL link that is not configured as a point-to-point link [RFC6325] [RFC7177].

ENCAP Hdr - Encapsulation headers in use between the IP Header and the TRILL Header. See Section 5.

ESP - IPsec Encapsulating Security Protocol [RFC4303].

FGL - Fine Grained Label [RFC7172].

Hdr - Used herein as an abbreviation for "Header".

HKDF - Hash based Key Derivation Function [RFC5869].

MTU - Maximum Transmission Unit.

RBridge - Routing Bridge. An alternative term for a TRILL switch.

TRILL - Transparent Internconnection of Lots of Links or Tunneled Routing in the Link Layer. The protocol specified in [RFC6325], [RFC7177], [rfc7180bis], and related RFCs.

TRILL switch - A device implementing the TRILL protocol.

VNI - Virtual Network Identifier. In VXLAN [RFC7348], the VXLAN Network Identifier.

3. Use Cases for TRILL over IP

This section introduces two application scenarios (a remote office scenario and an IP backbone scenario) which cover typical situations where network administrators may choose to use TRILL over an IP network to connect TRILL switches.

3.1 Remote Office Scenario

In the Remote Office Scenario, a remote TRILL network is connected to a TRILL campus across a multihop IP network, such as the public Internet. The TRILL network in the remote office becomes a part of TRILL campus, and nodes in the remote office can be attached to the same VLANs or Fine Grained Labels [RFC7172] as local campus nodes. In many cases, a remote office may be attached to the TRILL campus by a single pair of RBridges, one on the campus end, and the other in the remote office. In this use case, the TRILL over IP link will often cross logical and physical IP networks that do not support TRILL, and are not under the same administrative control as the TRILL campus.

3.2 IP Backbone Scenario

In the IP Backbone Scenario, TRILL over IP is used to connect a number of TRILL networks to form a single TRILL campus. For example, a TRILL over IP backbone could be used to connect multiple TRILL networks on different floors of a large building, or to connect TRILL networks in separate buildings of a multi-building site. In this use case, there may often be several TRILL switches on a single TRILL over IP link, and the IP link(s) used by TRILL over IP are typically under the same administrative control as the rest of the TRILL campus.

3.3 Important Properties of the Scenarios

There are a number of differences between the above two application scenarios, some of which drive features of this specification. These differences are especially pertinent to the security requirements of the solution, how multicast data frames are handled, and how the TRILL switch ports discover each other.

3.3.1 Security Requirements

In the IP Backbone Scenario, TRILL over IP is used between a number of RBridge ports, on a network link that is in the same administrative control as the remainder of the TRILL campus. While it is desirable in this scenario to prevent the association of unauthorized RBridges, this can be accomplished using existing IS-IS security mechanisms. There may be no need to protect the data traffic, beyond any protections that are already in place on the local network.

In the Remote Office Scenario, TRILL over IP may run over a network that is not under the same administrative control as the TRILL network. Nodes on the network may think that they are sending traffic locally, while that traffic is actually being sent, in an IP tunnel, over the public Internet. It is necessary in this scenario to protect the integrity and confidentiality of user traffic, as well as ensuring that no unauthorized RBridges can gain access to the RBridge campus. The issues of protecting integrity and confidentiality of user traffic are addressed by using IPsec for both TRILL IS-IS and TRILL Data packets between RBridges in this scenario.

3.3.2 Multicast Handling

In the IP Backbone scenario, native IP multicast may be supported on the TRILL over IP link. If so, it can be used to send TRILL IS-IS and multicast data packets, as discussed later in this document. Alternatively, multi-destination packets can be transmitted serially by IP unicast to the intended recipients.

In the Remote Office Scenario there will often be only one pair of RBridges connecting a given site and, even when multiple RBridges are used to connect a Remote Office to the TRILL campus, the intervening network may not provide reliable (or any) multicast connectivity. Issues such as complex key management also make it difficult to provide strong data integrity and confidentiality protections for multicast traffic. For all of these reasons, the connections between local and remote RBridges will commonly be treated like point-to-point links, and all TRILL IS-IS control messages and multicast data packets that are transmitted between the Remote Office and the TRILL campus will be serially transmitted by IP unicast, as discussed later in this document.

3.3.3 Neighbor Discovery

In the IP Backbone Scenario, TRILL switches that use TRILL over IP can use the normal TRILL IS-IS Hello mechanisms to discover the existence of other TRILL switches on the link [RFC7177], and to establish authenticated communication with them.

In the Remote Office Scenario, an IPsec session will need to be established before TRILL IS-IS traffic can be exchanged, as discussed below. In this case, one end will need to be configured to establish a IPSEC session with the other. This will typically be accomplished by configuring the TRILL switch or a border device at a Remote Office to initiate an IPsec session and subsequent TRILL exchanges with a TRILL over IP-enabled RBridge attached to the TRILL campus.


4. TRILL Packet Formats

To support the TRILL protocol [RFC6325], two types of TRILL packets are transmitted between TRILL switches: TRILL Data packets and TRILL IS-IS packets.

Section 4.1 describes general TRILL packet formats for data and IS-IS independent of link technology. Section 4.2 specifies general TRILL over IP packet formats including IPsec ESP encapsulation. Section 4.3 provides QoS Considerations. Section 4.4 discusses broadcast links and multicast packets. And Section 4.5 provides TRILL IS-IS Hello SubNetwork Point of Attachment (SNPA) considerations for TRILL over IP.


4.1 General Packet Formats

The on-the-wire form of a TRILL Data packet in transit between two neighboring TRILL switch ports is as shown below:



The encapsulated Native Frame Payload is similar to an Ethernet frame with a VLAN tag or Fine Grained Label [RFC7172] but with no trailing Frame Check Sequence (FCS).

TRILL IS-IS packets are formatted on-the-wire as follows:



The Link Header and Link Trailer in these formats depend on the specific link technology. The Link Header contains one or more fields that distinguish TRILL Data from TRILL IS-IS. For example, over Ethernet, the Link Header for TRILL Data ends with the TRILL Ethertype while the Link Header for TRILL IS-IS ends with the L2-IS-IS Ethertype; on the other hand, over PPP, there are no Ethernets in the Link Header but PPP protocol code points are included that distinguish TRILL Data from TRILL IS-IS.


4.2 General TRILL Over IP Packet Formats

In TRILL over IP, we will use an IP (v4 or v6) header as the link header. (On the wire, the IP header will normally be preceded by the lower layer header of a protocol that is carrying IP; however, this does not concern us at the level of this document.)

There are multiple IP based encapsulations usable for TRILL over IP that differ in exactly what appears after the IP header and before the TRILL Header or the TRILL IS-IS Payload. These encapsulations are further detailed in Section 5. In the general specification below, those encapsulation fields will be represented as "ENCAP Hdr". See Section 5 for details.

4.2.1 Without Security

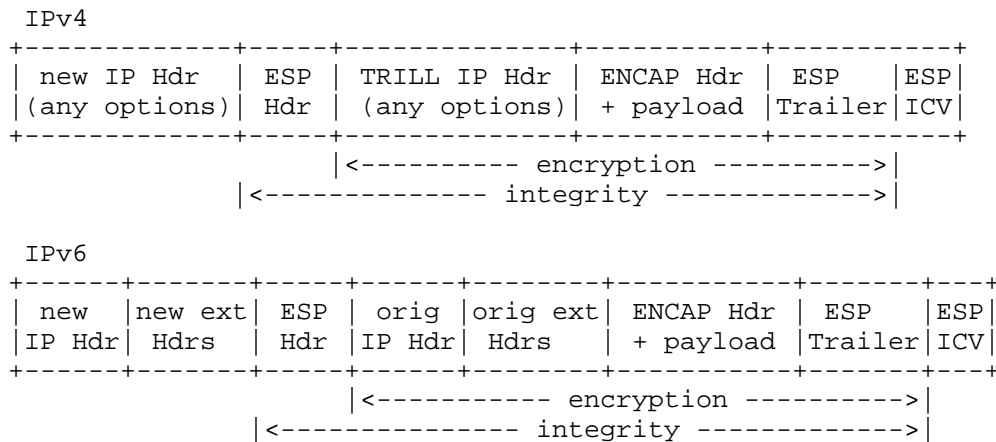
When TRILL over IP link security is not being used, a TRILL over IP packet on the wire looks like the following:



As discussed above and further specified in Section 5, the ENCAP Hdr indicates whether the packet is TRILL Data or IS-IS.

4.2.2 With Security

TRILL over IP link security uses IPsec Encapsulating Security Protocol (ESP) in tunnel mode [RFC4303]. Since TRILL over IP always starts with an IP Header (on the wire this appears right after any lower layer header that might be required), the modifications for IPsec are independent of the TRILL over IP ENCAP Hdr that occurs after that IP Header. The resulting packet formats are as follows for IPv4 and IPv6:




As shown above, IP Header options are considered part of the IPv4 Header but are extensions ("ext") of the IPv6 Header. For further information on the IPsec ESP Hdr, Trailer, and ICV, see [RFC4303] and Section 7. "ENCAP Hdr + payload" is the encapsulation header (Section 5) and TRILL data or IS-is payload, that is, the material after the IP Header in the diagram in Section 4.2.1.

This architecture permits the ESP tunnel end point to be separated from the TRILL over IP RBridge port (see, for example, Section 1.1.3 of [RFC7296]).

4.3 QoS Considerations

In IP, QoS handling is indicated by the Differential Services Code Point (DSCP [RFC2474] [RFC3168]) in the TRILL Header. The former Type of Service (TOS) octet in the IPv4 Header and the Traffic Class octet in the IPv6 Header has been divided as shown in the following diagram adapted from [RFC3168]. (TRILL support of ECN is beyond the scope of this document.)



DSCP: Differentiated Services Codepoint
ECN: Explicit Congestion Notification

Within a TRILL switch, priority is indicated by configuration for TRILL IS-IS packets and for TRILL Data packets by a three bit (0 through 7) priority field and a Drop Eligibility Indicator bit (see Sections 8.2 and 7 of [rfc7180bis]). (Typically TRILL IS-IS is

configured to use the highest priority or, alternatively, the highest two priorities depending on the IS-IS PDU.) The priority affects queuing behavior at TRILL switch ports and may be encoded into the link header, particularly if there could be priority sensitive devices within the link. For example, if the link is a bridged LAN, it is commonly encoded into an Outer.VLAN tag's priority and DEI fields.

TRILL over IP implementations MUST support setting the DSCP value in the outer IP Header of TRILL packets they send by mapping the TRILL priority and DEI to the DSCP. They MAY support, for a TRILL Data packet where the native frame payload is an IP packet, copying the DSCP in this inner IP packet to the outer IP Header.

The default TRILL priority and DEI to DSCP mapping, which may be configured per TRILL over IP port, is as follows. Note that the DEI value does not affect the default mapping and, to provide a potentially lower priority service than the default 0, priority 1 is considered lower priority than 0. So the priority sequence from lower to higher priority is 1, 0, 2, 3, 4, 5, 6, 7.

TRILL Priority	DEI	DSCP Field (Binary/decimal)
0	0/1	001000 / 8
1	0/1	000000 / 0
2	0/1	010000 / 16
3	0/1	011000 / 24
4	0/1	100000 / 32
5	0/1	101000 / 40
6	0/1	110000 / 48
7	0/1	111000 / 56

4.4 Broadcast Links and Multicast Packets

TRILL supports broadcast links. These are links to which more than two TRILL switch ports can be attached and where a packet can be broadcast or multicast from a port to all or a subset of the other ports on the link as well as unicast to a specific single other port on the link.

As specified in [RFC6325], TRILL Data packets being forwarded between TRILL switches can be unicast on a link to a specific TRILL switch port or multicast on a link to all TRILL switch ports. TRILL IS-IS packets are always multicast to all other TRILL switches on the link except for IS-IS MTU PDUs, which may be unicast [RFC7177]. This distinction is not significant if the link is inherently point-to-point, such as a PPP link; however, on a broadcast link there will be a packet outer link address that is unicast or multicast as

appropriate. For example, over Ethernet links, the Ethernet multicast addresses All-RBridges and All-IS-IS-RBridges are used for multicasting TRILL Data and TRILL IS-IS respectively. For details on TRILL over IP handling of multicast, see Section 6.

4.5 TRILL Over IP IS-IS SubNetwork Point of Attachment

IS-IS routers, such as TRILL switches, establish adjacency through the exchange of Hello PDUs on a link [IS-IS] [RFC7177]. The Hellos transmitted out a port indicate what neighbor ports that port can see on the link by listing what IS-IS refers to as the neighbor port's SubNetwork Point of Attachment (SNPA). (For an Ethernet link, which may be a bridged LAN, the SNPA is the port MAC address.)

In TRILL Hello PDUs on a TRILL over IP link, the IP addresses of the IP ports connected to that link are their actual SNPA (SubNetwork Point of Attachment [IS-IS]) addresses and, for IPv6, the 16-byte IPv6 address is used as the SNPA; however, for easy in re-using code designed for the common case of 48-bit SNPAs, in TRILL over IPv4 a 48-bit synthetic SNPA that looks like a unicast MAC address is constructed for use in the SNPA field of TRILL Neighbor TLVs [RFC7176] [RFC7177] in such Hellos. This synthetic SNPA is derived from the port IPv4 address is as follows:

```

          1 1 1 1 1 1
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-----+
| 0xFE          | 0x00          |
+-----+
| IPv4 upper half          |
+-----+
| IPv4 lower half         |
+-----+
```

This synthetic SNPA (MAC) address has the local (0x02) bit on in the first byte and so cannot conflict with any globally unique 48-bit Ethernet MAC. However, when TRILL operates on an IP link, TRILL sees only IP stations, not MAC stations, even if the TRILL over IP Link is being carried over Ethernet. Therefore conflict on the link in TRILL IS-IS between a real MAC address and the synthetic SNPA (MAC) address as above would be impossible in any case.

5. TRILL over IP Encapsulation Formats

There are a variety of TRILL over IP encapsulation formats possible. By default TRILL over IP adopts a hybrid encapsulation approach.

There is one format, called "native encapsulation" that MUST be implemented. Although native encapsulation does not typically have good fast path support, as a lowest common denominator it can be used by low bandwidth control traffic to determine a preferred encapsulation with better performance. In particular, by default, all TRILL IS-IS Hellos are sent using native encapsulation and those Hellos are used to determine the encapsulation used for all TRILL Data packets and all other TRILL IS-IS PDUs (with the possible exception of IS-IS MTU-probe and MTU-ack PDUs).

Alternatively, the network operator can pre-configure a TRILL over IP port to use a particular encapsulation chosen for their particular network needs and port capabilities. That encapsulation is then used for all TRILL Data and IS-IS packets on ports so configured.

Section 5.1 discusses general consideration for the TRILL over IP encapsulation format. Section 5.2 discusses encapsulation agreement. Section 5.3 discusses broadcast link encapsulation considerations. The subsequent subsections discuss particular encapsulations.

5.1 Encapsulation Considerations

In all cases, there must be a method specified to distinguish TRILL Data packets and TRILL IS-IS packets, or that encapsulation is not useful for TRILL. In addition, the following criteria can be helpful in choosing between different encapsulations:

- a) Fast path support - For many applications, it is highly desirable to be able to encapsulate/decapsulate TRILL over IP at line speed so a format where existing or anticipated fast path hardware can do that is best. This is commonly a dominant consideration.
- b) Ease of multi-pathing - The IP path between TRILL over IP ports may include equal cost multipath routes internal to the IP link so a method of encapsulation that provides variable fields available for existing or anticipated fast path hardware multi-pathing is better.
- c) Robust fragmentation and re-assembly - MTU of the IP link may require fragmentation in which case an encapsulation with robust fragmentation and re-assembly is important. There are known problems with IPv4 fragmentation and re-assembly [RFC6864] which generally do not apply to IPv6. Some encapsulations can fix these

problems but the two encapsulations specified in this document do not. Therefore, if fragmentation is anticipated with the encapsulations specified in this document, the use of IPv6 is RECOMMENDED.

- d) Checksum strength - Depending on the particular circumstances of the TRILL over IP link, a checksum provided by the encapsulation may be an important factor. Use of IPsec can also provide a strong integrity check.

5.2 Encapsulation Agreement

TRILL Hellos sent out a TRILL over IP port indicate the encapsulations that port is willing to support through a mechanism initially specified in [RFC7178] and [RFC7176] that is hereby extended. Specifically, RBridge Channel Protocol numbers 0xFD0 through 0xFF7 are redefined to be link technology dependent flags that, for TRILL over IP, indicate support for different encapsulations, allowing for up to 40 encapsulations to be specified. Support for an encapsulation is indicated in the Hello PDU in the same way that support for an RBridge Channel was indicated. (See also section 11.3.) "Support" indicates willingness to use that encapsulation for TRILL Data and TRILL IS-IS packets (although TRILL IS-IS Hellos are still sent in native encapsulation by default).

If, in a TRILL Hello on a TRILL over IP link, support is not indicated for any encapsulation, then the port from which it was sent is assumed to support only native encapsulation (see Section 5.4).

An adjacency is formed between two TRILL over IP ports if the intersection of the sets of encapsulation methods they support is not null. If that intersection is null, then no adjacency is formed. In particular, for a TRILL over IP link, the adjacency state machine MUST NOT advance to the Report state unless the ports share an encapsulation [RFC7177]. If no encapsulation is shared, the adjacency state machine remains in the state from which it would otherwise have transitioned to the Report state.

If any TRILL over IP packet, other than an IS-IS Hello or MTU PDU in native encapsulation, is received in an encapsulation for which support is not being indicated, it MUST be discarded (see Section 5.3).

If there are two or more encapsulations in common between two adjacent ports for unicast or the set of adjacent ports for multicast, a transmitter is free to choose whichever of the encapsulations it wishes to use. Thus transmissions between adjacent ports P1 and P2 could use different encapsulations depending on which

port is transmitting and which is receiving.

It is expected to be the normal case in a well configured network that all the TRILL over IP ports connected to an IP link (i.e., an IP network) that are intended to communicate with each other will support the same encapsulation(s).


5.3 Broadcast Link Encapsulation Considerations

To properly handle TRILL protocol packets on a TRILL over IP link in the general case, either native IP multicast mode is used on that link or multicast must be simulated using serial IP unicast, as discussed in Section 6. (Of course, if the IP link happens to actually be point-to-point no special provision is needed for handling multicast addressed packets.)


It is possible for the Hellos from a TRILL over IP port P1 to establish adjacency with multiple other TRILL over IP ports (P2, P3, ...) on broadcast link. In a well configured network one would expect all of the IP ports involved to support the same encapsulation(s); but, if P1 supports multiple encapsulations, it is possible that P2 and P3, for example, do not have an encapsulation in common that is supported by P1. IS-IS can handle such non-transitive adjacencies which are reported as specified in [RFC7177]. If serial IP unicast is being used by P1, it can use different encapsulations for different transmissions. If native IP multicast is being used by P1, it will have to send one transmission per encapsulation method by which it has an adjacency on the link. (It is for this reason that a TRILL over IP port MUST discard any packet received with the wrong encapsulation. Otherwise, packets would be duplicated.)

5.4 Native Encapsulation

The mandatory to implement "native encapsulation" format of a TRILL over IP packet, when used without security, is TRILL over UDP as shown below.



Where the UDP Header is as follows:



Source Port - see Section 8.3

Destination Port - indicates TRILL Data or IS-IS, see Section 11


UDP Length - as specified in [RFC0768]

UDP Checksum - as specified in [RFC0768]

The TRILL Payload starts with the TRILL Header (not including the TRILL Ethertype) for TRILL Data packets and starts with the 0x83 Intradomain Routing Protocol Discriminator byte (thus not including the L2-IS-IS Ethertype) for TRILL IS-IS packets.

5.5 VXLAN Encapsulation

VXLAN [RFC7348] IP encapsulation of TRILL looks, on the wire, like TRILL over Ethernet over VXLAN over UDP over IP.



The outer UDP uses a destination port number indicating VXLAN and the outer UDP source port MAY be used for entropy as with native encapsulation (see Section 5.4). The VXLAN header after the outer UDP header adds a 24 bit Virtual Network Identifier (VNI). The Ethernet header after the VXLAN header and before the TRILL header consists of source MAC address, destination MAC address, and Ethertype. The Ethertype distinguishes TRILL Data from TRILL IS-IS; however, the destination and source MAC addresses in this inner Ethernet header are not used and are 12 wasted bytes.

A TRILL over IP port using VXLAN encapsulation by default uses a VNI of 1 but can be configured as described in Section 9.2.3.1 to use some other fixed VNI or to map from VLAN/FGL to VNI.

5.6 Other Encapsulations

It is anticipated that additional TRILL over IP encapsulations will be specified in future documents and allocated a bit in the TRILL Hello as per Section 11.3. A primary consideration for whether it is worth the effort to specify an encapsulation is good existing or anticipated fast path support.

6. Handling Multicast

By default, both TRILL IS-IS packets and multi-destination TRILL Data packets are sent to an All-RBridges IPv4 or IPv6 IP multicast Address as appropriate (see Section 11.2); however, a TRILL over IP port may be configured (see Section 9) to use a different multicast address or to use serial IP unicast with a list of one or more unicast IP addresses of other TRILL over IP ports to which multi-destination packets are sent. In the serial unicast case the outer IP header of each copy of the packet sent shows an IP unicast destination address even though the TRILL header has the M bit set to one to indicate multi-destination. Serial unicast configuration is necessary if the TRILL over IP port is connected to an IP network that does not support IP multicast. In any case, unicast TRILL packets are sent by unicast IP.

Even if a TRILL over IP port is configured to send multi-destination packets with serial unicast, it MUST be prepared to receive IP multicast TRILL packets. All TRILL over IP ports default to periodically transmitting appropriate IGMP (IPv4 [RFC3376] or MLD (IPv6 [RFC2710]) packets, so that the TRILL multicast IP traffic will be sent to them, unless they are configured not to do so.

Although TRILL fully supports broadcast links with more than 2 RBridges connected to the link there may be good reasons for configuring TRILL over IP ports to use serial unicast even where native IP multicast is available. Use of serial unicast provides the network manager with more precise control over adjacencies and how TRILL over IP links will be formed in an IP network. In some networks, unicast is more reliable than multicast. If multiple point-to-point TRILL over IP connections between parts of a TRILL campus are configured, TRILL will in any case spread traffic across them, treating them as parallel links, and appropriately fail over traffic if a link fails or incorporate a new link that comes up.

7. Use of IPsec and IKEv2

All TRILL switches (RBridges) that support TRILL over IP MUST implement IPsec [RFC4301] and support the use of IPsec Encapsulating Security Protocol (ESP [RFC4303]) in tunnel mode to secure both TRILL IS-IS and TRILL data packets. When IPsec is used to secure a TRILL over IP link and no IS-IS security is enabled, the IPsec session MUST be fully established before any TRILL IS-IS or data packets are exchanged. When there is IS-IS security [RFC5310] provided, implementers SHOULD use IS-IS security to protect TRILL IS-IS packets. However, in this case, the IPsec session still MUST be fully established before any data packets transmission since IS-IS security does not provide any protection to data packets.

All RBridges that support TRILL over IP MUST implement the Internet Key Exchange Protocol version 2 (IKEv2) for automated key management.

7.1 Keying

The following subsections discuss pairwise and group keying for TRILL over IP IPsec.

7.1.1 Pairwise Keying

When IS-IS security is in use, IKEv2 will use a pre-shared key that incorporates the IS-IS shared key in order to bind the TRILL data session to the IS-IS session. The pre-shared key that will be used for IKEv2 exchanges for TRILL over IP is determined as follows:

```
HKDF-Expand-SHA256 ( IS-IS-key,  
  "TRILL IP" | P1-System-ID | P1-Port | P2-System-ID | P2-Port )
```

In the above "|" indicates concatenation, HKDF is as in [RFC5869], SHA256 is as in [RFC6234], and "TRILL IP" is the eight byte US ASCII [RFC0020] string indicated. "IS-IS-key" is an IS-IS key usable for IS-IS security of link local IS-IS PDUs such as Hello, CSNP, and PSNP. This SHOULD be a link scope IS-IS key. With [RFC5310] there could be multiple keys identified with 16-bit key IDs. In this case, the Key ID of IS-IS-key is also used to identify the derived key. P1-System-ID and P2-System ID are the System IDs of the two TRILL RBridges, and P1-Port and P2-Port are the ports in use on each end. System IDs are guaranteed to be unique within the TRILL campus. Both of the RBridges involved treat the larger magnitude System ID, comparing System IDs as unsigned integers, as P1 and the smaller as P2 so both will derive the same key.

When IS-IS security is in use, the IS-IS-shared key from which the IKEv2 shared secret is derived might expire and be updated as described in [RFC5310]. The IKEv2 pre-shared keys derived from the IS-IS shared key MUST expire within the same lifetime as the IS-IS-shared key from which they were derived. When the IKEv2 pre-shared key expires, the IKEv2 Security Association must be rekeyed using a new shared secret derived from the new IS-IS shared key.

When IS-IS security is not in use, IKEv2 will not use a pre-shared key.

7.1.2 Group Keying

In the case of a TRILL over IP port configured as point-to-point (see Section 4.2.4.1 of [RFC6325]), there is no group keying and the pairwise key determined as in Section 7.1.1 is used for IP multicast traffic.

In the case of a TRILL over IP port configured as broadcast but where the port is configured to use serial unicast (see Section 8), there is no group keying and the pairwise keying determined as in Section 7.1.1 is used for IP multicast traffic.

In the case of a TRILL over IP port configured as broadcast and using native multicast, ... tbd ...

7.2 Mandatory-to-Implement Algorithms

All RBridges that support TRILL over IP MUST implement IPsec ESP [RFC4303] in tunnel mode. The implementation requirements for ESP cryptographic algorithms are as specified for IPsec. That specification is currently [RFC7321].

8. Transport Considerations

This section discusses a variety of important transport considerations.

8.1 Congestion Considerations

Section 3.1.3 of [RFC5405] discussed the congestion implications of UDP tunnels. As discussed in [RFC5405], because other flows can share the path with one or more UDP tunnels, congestion control [RFC2914] needs to be considered.

The default initial determination of the TRILL over IP encapsulation to be used through the exchange of TRILL IS-IS Hellos is a low bandwidth process. Hellos are not permitted to be sent any more often than once per second, and so are unlikely to cause congestion.

One motivation for including UDP in a TRILL encapsulation is to improve the use of multipath (such as ECMP) in cases where traffic is to traverse routers which are able to hash on UDP Port and IP address. In many cases this may reduce the occurrence of congestion and improve usage of available network capacity. However, it is also necessary to ensure that the network, including applications that use the network, responds appropriately in more difficult cases, such as when link or equipment failures have reduced the available capacity.

The impact of congestion must be considered both in terms of the effect on the rest of the network of a UDP tunnel that is consuming excessive capacity, and in terms of the effect on the flows using the UDP tunnels. The potential impact of congestion from a UDP tunnel depends upon what sort of traffic is carried over the tunnel, as well as the path of the tunnel.

TRILL is used to carry a wide range of traffic. In many cases TRILL is used to carry IP traffic. IP traffic is generally assumed to be congestion controlled, and thus a tunnel carrying general IP traffic (as might be expected to be carried across the Internet) generally does not need additional congestion control mechanisms. As specified in [RFC5405]:

"IP-based traffic is generally assumed to be congestion-controlled, i.e., it is assumed that the transport protocols generating IP-based traffic at the sender already employ mechanisms that are sufficient to address congestion on the path. Consequently, a tunnel carrying IP-based traffic should already interact appropriately with other traffic sharing the path, and specific congestion control mechanisms for the tunnel are not necessary".

For this reason, where TRILL is sent using UDP and used to carry IP traffic that is known to be congestion controlled, the UDP paths MAY be used across any combination of a single or cooperating service providers or across the general Internet.

However, TRILL is also used to carry traffic that is not necessarily congestion controlled. For example, TRILL may be used to carry traffic where specific bandwidth guarantees are provided.

In such cases congestion may be avoided by careful provisioning of the network and/or by rate limiting of user data traffic. Where TRILL is carried, directly or indirectly, over UDP over IP, the identity of each individual TRILL flow is in general lost.

For this reason, where the TRILL traffic is not congestion controlled, TRILL over UDP/IP MUST only be used within a single service provider that utilizes careful provisioning (e.g., rate limiting at the entries of the network while over-provisioning network capacity) to ensure against congestion, or within a limited number of service providers who closely cooperate in order to jointly provide this same careful provisioning. As such, TRILL over UDP/IP MUST NOT be used over the general Internet, or over non-cooperating service providers, to carry traffic that is not congestion-controlled.

Measures SHOULD be taken to prevent non-congestion-controlled TRILL over UDP/IP traffic from "escaping" to the general Internet, for example the following:

- a. Physical or logical isolation of the TRILL over IP links from the general Internet.
- b. Deployment of packet filters that block the UDP ports assigned for TRILL-over-UDP.
- c. Imposition of restrictions on TRILL over UDP/IP traffic by software tools used to set up TRILL over UDP paths between specific end systems (as might be used within a single data center).
- d. Use of a "Managed Circuit Breaker" for the TRILL traffic as described in [circuit-breaker].

8.2 Recursive Ingress

TRILL is specified to transport data to and from end stations over Ethernet and IP is frequently transported over Ethernet. Thus, an end station native data Ethernet frame EF might get TRILL ingressed to

TRILL(EF) that was then sent out a TRILL over IP over Ethernet port resulting in a packet on the wire of the form Ethernet(IP(TRILL(EF))). There is a risk of such a packet being re-ingressed by the same TRILL campus, due to physical or logical misconfiguration, looping round, being further re-ingressed, and so on. The packet might get discarded if it got too large but if fragmentation is enabled, it would just keep getting split into fragments that would continue to loop and grow and re-fragment until the path was saturated with junk and packets were being discarded due to queue overflow. The TRILL Header TTL would provide no protection because each TRILL ingress adds a new TRILL header with a new TTL.

To protect against this scenario, a TRILL over IP port MUST by default, test whether a TRILL packet it is about to transmit appears to be a TRILL ingress of a TRILL over IP over Ethernet packet. That is, is it of the form TRILL(Ethernet(IP(TRILL(...)))? If so, the default action of the TRILL over IP output port is to discard the packet rather than transmit it. However, there are cases where some level of nested ingress is desired so it MUST be possible to configure the port to allow such packets.

8.3 Fat Flows

For the purpose of load balancing, it is worthwhile to consider how to transport the TRILL packets over the Equal Cost Multiple Paths (ECMPs) existing internal to the IP path between TRILL over IP ports.

The ECMP election for the IP traffic could be based, at least for IPv4, on the quintuple of the outer IP header { Source IP, Destination IP, Source Port, Destination Port, and IP protocol }. Such tuples, however, could be exactly the same for all TRILL Data packets between two RBridge ports, even if there is a huge amount of data being sent between a variety of ingress and egress RBridges. One solution to this is to use the Source Port in as an entropy field. (This idea is also introduced in [gre-in-udp].) For example, for TRILL Data this entropy field could be based on some hash of the Inner.MacDA, Inner.MacSA, and Inner.VLAN or Inner.FGL. Unfortunately, this can conflict with middleboxes inside the TRILL over IP link (see 8.5). Therefore, in order to better support ECMP, a RBridge SHOULD set the Source Port to a range of values as an entropy field for ECMP decisions. However, if there are middleboxes in the path, the range of different Source Port values used MUST be restricted sufficiently to avoid disrupting connectivity.

8.4 MTU Considerations

In TRILL each TRILL switch advertises in its LSP number zero the largest LSP frame it can accept (but not less than 1,470 bytes) on any of its interfaces (at least those interfaces with adjacencies to other TRILL switches in the campus) through the `originatingLSPBufferSize` TLV [RFC6325] [RFC7177]. The campus minimum MTU (Maximum Transmission Unit), denoted *Sz*, is then established by taking the minimum of this advertised MTU for all R Bridges in the campus. Links that do not meet the *Sz* MTU are not included in the routing topology. This protects the operation of IS-IS from links that would be unable to accommodate some LSPs.

A method of determining `originatingLSPBufferSize` for an R Bridge with one or more TRILL over IP ports is described in [rfc7180bis]. However, if an IP link either can accommodate jumbo frames or is a link on which IP fragmentation is enabled and acceptable, then it is unlikely that the IP link will be a constraint on the `originatingLSPBufferSize` of an R Bridge using the link. On the other hand, if the IP link can only handle smaller frames and fragmentation is to be avoided when possible, a TRILL over IP port might constrain the R Bridge's `originatingLSPBufferSize`. Because TRILL sets the minimum values of *Sz* at 1,470 bytes, there may be links that meet the minimum MTU for the IP protocol (1,280 bytes for IPv6, 576 bytes for IPv4) on which it would be necessary to enable fragmentation for TRILL use.

The use of TRILL IS-IS MTU PDUs, as specified in [RFC6325] and [RFC7177] can provide added assurance of the actual MTU of a link.

8.5 Middlebox Considerations

This section gives some middlebox considerations for the IP encapsulations covered by this document, namely native and VXLAN encapsulation.

The requirements on the usage of the zero UDP Checksum in a UDP tunnel protocol are detailed in [RFC6936]. These requirements apply to TRILL over IP the encapsulations specified herein (native and VXLAN), which are applications of UDP tunnel.

Besides the Checksum, the Source Port number of the UDP header is also pertinent to the middlebox behavior. Network Address/Port Translator (NAPT) is the most commonly deployed Network Address Translation (NAT) device [RFC4787]. For a UDP tunnel protocol, the NAPT device establishes a NAT session to translate the {private IP address, private source port number} tuple to a {public IP address, public source port number} tuple, and vice versa, for the duration of

the UDP session. This provides the UDP tunnel protocol application with the "NAT-pass-through" function. NAPT allows multiple internal hosts to share a single public IP address. The port number, i.e., the UDP Source Port number, is used as the demultiplexer of the multiple internal hosts.

However, the above NAPT behavior conflicts with the behavior that the UDP Source Port number is used as an entropy (See Section 8.3). Hence, the tunnel operator **MUST** ensure the TRILL switch ports sending through local or remote NAPT middleboxes disable the entropy usage of the UDP Source Port number.

9. TRILL over IP Port Configuration

This section specifies the configuration information needed at a TRILL over IP port beyond that needed for a general RBridge port.

9.1 Per IP Port Configuration

Each RBridge port used for a TRILL over IP link should have at least one IP (v4 or v6) address. If no IP address is associated with the port, perhaps as a transient condition during re-configuration, the port is disabled. Implementations MAY allow a single port to operate as multiple IPv4 and/or IPv6 logical ports. Each IP address constitutes a different logical port and the RBridge with those ports MUST associate a different Port ID (see Section 4.4.2 of [RFC6325]) with each logical port.

By default a TRILL over IP port discards output packets that fail the possible recursive ingress test (see Section 10.1) unless configured to disable that test.

9.2 Additional per IP Address Configuration

The configuration information specified below is per TRILL over IP port IP address.

The mapping from TRILL packet priority to Differentiated Services Code Point (DSCP [RFC2474]) can be configured (see Section 10.5).

Each TRILL over IP port has a list of acceptable encapsulations it will use. By default this list consists of one entry for native encapsulation (see Section 7). Additional encapsulations MAY be configured. Additional configuration can be required or possible for specific encapsulations as described in Section 9.2.3.

Each IP address at a TRILL over IP port uses native IP multicast by default but may be configured whether to use serial IP unicast (Section 9.2.2) or native IP multicast (Section 9.2.1). Each IP address at a TRILL over IP is configured whether or not to use IPsec (Section 9.2.4).

9.2.1 Native Multicast Configuration

If a TRILL over IP port address is using native IP multicast for multi-destination TRILL packets (IS-IS and data), by default

transmissions from that IP address use the IP multicast address (IPv4 or IPv6) specified in Section 11.2. The TRILL over IP port may be configured to use a different IP address to multicast packets.

9.2.2 Serial Unicast Configuration

If a TRILL over IP port address has been configured to use serial unicast for multi-destination packets (IS-IS and data), it should have associated with it a non-empty list of unicast IP destination addresses with the same IP version as the version of the port's IP address (IPv4 or IPv6). Multi-destination TRILL packets are serially unicast to the addresses in this list. Such a TRILL over IP port will only be able to form adjacencies [RFC7177] with the RBridges at the addresses in this list as those are the only RBridges to which it will send TRILL Hellos.

If this list of destination IP addresses is empty, there is no way to transmit a multi-destination TRILL over IP packet such as a TRILL Hello. Thus it is impossible to achieve adjacency [RFC7177] or if adjacency had been achieved (perhaps the list was non-empty and has just been configured to be empty), no way to maintain such adjacency. Thus, in the empty list case, TRILL Data multi-destination packets cannot be sent and TRILL Data unicast packets will not start flowing or, if they are already flowing, will soon cease, effectively disabling the port.

9.2.3 Encapsulation Specific Configuration

Specific TRILL over IP encapsulation methods may provide for further configuration as specified below.

9.2.3.1 VXLAN Configuration

A TRILL over IP port using VXLAN encapsulation can be configured with a non-default VXLAN Network Identifier (VNI) that is used in that field of the VXLAN header for all TRILL packets sent using the encapsulation and required in all TRILL packets received using the encapsulation. The default VNI is 1. A TRILL packet received with the wrong VNI is discarded.

A TRILL over IP port using VXLAN encapsulation can also be configured to map the Inner.VLAN or Inner.FGL of a TRILL Data packet being transported to the value it places in the VNI field.

9.2.3.2 Other Encapsulation Configuration

Additional encapsulation methods, beyond the native UDP encapsulation and VXLAN encapsulation specified in this document, may be specified in future documents and may require further configuration.

9.2.4 Security Configuration

tbd ...

10. Security Considerations

TRILL over IP is subject to all of the security considerations for the base TRILL protocol [RFC6325]. In addition, there are specific security requirements for different TRILL deployment scenarios, as discussed in the "Use Cases for TRILL over IP" section above.

For communication between end stations in a TRILL campus, security is possible at three levels: end-to-end security between those end stations, edge-to-edge security between ingress and egress R Bridges [LinkSec], and link security to protect a TRILL hop. Any combination of these can be used, including all three.

TRILL over IP link security protects the contents of TRILL Data and IS-IS packets, including the identities of the end stations for data and the identities of the edge R Bridges, from observers of the link and transit devices within the link such as IP routers, but does not encrypt the link local IP addresses used in a packet and does not protect against observation by the sending and receiving R Bridges on the link. Edge-to-edge TRILL security protects the contents of TRILL data packets including the identities of the end stations for data from transit R Bridges but does not encrypt the identities of the edge R Bridges involved and does not protect against observation by those edge R Bridges. End-to-end security does not protect the identities of the end stations or edge R Bridge involved but does protect the content of TRILL data packets from observation by all R Bridges or other intervening devices between the end stations involved. End-to-end security should always be considered as an added layer of security and to protect any particularly sensitive information from unintended disclosure.

If VXLAN encapsulation is used, the unused Ethernet source and destination MAC addresses mentioned in Section 5.5, provide a 96 bit per packet covert path.

10.1 IPsec

This document specifies that all R Bridges that support TRILL over IP links MUST implement IPsec for the security of such links, and makes it clear that it is both wise and good to use IPsec in all cases where a TRILL over IP link will traverse a network that is not under the same administrative control as the rest of the TRILL campus or is not physically secure. IPsec is important, in these cases, to protect the privacy and integrity of data traffic. However, in cases where IPsec is impractical due to lack of fast path support, use of TRILL edge-to-edge security or use by the end stations of end-to-end security can provide significant security.

Further Security Considerations for IPsec ESP and for the cryptographic algorithms used with IPsec can be found in the RFCs referenced by this document.

10.2 IS-IS Security

TRILL over IP is compatible with the use of IS-IS Security [RFC5310], which can be used to authenticate TRILL switches before allowing them to join a TRILL campus. This is sufficient to protect against rogue devices impersonating TRILL switches, but is not sufficient to protect data packets that may be sent in TRILL over IP outside of the local network or across the public Internet. To protect the privacy and integrity of that traffic, use IPsec.

In cases where IPsec is used, the use of IS-IS security may not be necessary, but there is nothing about this specification that would prevent using both IPsec and IS-IS security together.

11. IANA Considerations

IANA considerations are given below.

11.1 Port Assignments

IANA is requested to assign destination UDP Ports for the TRILL IS-IS and Data channels:

UDP Port	Protocol
-----	-----
(TBD1)	TRILL IS-IS Channel
(TBD2)	TRILL Data Channel

11.2 Multicast Address Assignments

IANA is requested to one IPv4 and one IPv6 multicast address, as shown below, which correspond to the All-RBridges and All-IS-IS-RBridges multicast MAC addresses that the IEEE Registration Authority has assigned for TRILL. Because the low level hardware MAC address dispatch considerations for TRILL over Ethernet do not apply to TRILL over IP, one IP multicast address for each version of IP is sufficient.

(Values recommended to IANA in square brackets)

Name	IPv4	IPv6
-----	-----	-----
All-RBridges	TBD3[233.252.14.0]	TBD4[FF0X:0:0:0:0:0:0:205]

The hex digit "X" in the IPv6 address indicates the scope and defaults to 8. The IPv6 All-RBridges IP address may be used with other values of X.

11.3 Encapsulation Method Support Indication

The existing "RBridge Channel Protocols" registry is re-named and a new sub-registry under that registry added as follows:

The TRILL Parameters registry for "RBridge Channel Protocols" is renamed the "RBridge Channel Protocols and Link Technology Specific Flags" registry. [this document] is added as a second reference for this registry. The first part of the table is changed to the following:

Range	Registration	Note
-----	-----	-----
0x002-0x0FF	Standards Action	
0x100-0xFCF	RFC Required	allocation of a single value
0x100-0xFCF	IESG Approval	allocation of multiple values
0xFD0-0xFF7	see Note	link technology dependent, see subregistry

In the existing table of RBridge Channel Protocols, the following line is changed to two lines as shown:

OLD

0x004-0xFF7 Unassigned

NEW

0x004-0xFCF Unassigned

0xFD0-0xFF7 (link technology dependent, see subregistry)

A new subregistry under the re-named "RBridge Channel Protocols and Link Technology Specific Flags" registry is added as follows:

Name: TRILL over IP Link Flags
 Registration Procedure: IETF Review
 Reference: [this document]

Flag	Meaning	Reference
-----	-----	-----
0xFD0	Native encapsulation supported	[this document]
0xFD1	VXLAN encapsulation supported	[this document]
0xFD2-0xFF7	Unassigned	

Normative References

- [IS-IS] - "Intermediate system to Intermediate system routing information exchange protocol for use in conjunction with the Protocol for providing the Connectionless-mode Network Service (ISO 8473)", ISO/IEC 10589:2002, 2002".
- [RFC0020] - Cerf, V., "ASCII format for network interchange", STD 80, RFC 20, DOI 10.17487/RFC0020, October 1969, <<http://www.rfc-editor.org/info/rfc20>>.
- [RFC0768] - Postel, J., "User Datagram Protocol", STD 6, RFC 768, DOI 10.17487/RFC0768, August 1980, <<http://www.rfc-editor.org/info/rfc768>>.
- [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>>.
- [RFC2474] - Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, DOI 10.17487/RFC2474, December 1998, <<http://www.rfc-editor.org/info/rfc2474>>.
- [RFC2710] - Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, DOI 10.17487/RFC2710, October 1999, <<http://www.rfc-editor.org/info/rfc2710>>.
- [RFC2914] - Floyd, S., "Congestion Control Principles", BCP 41, RFC 2914, DOI 10.17487/RFC2914, September 2000, <<http://www.rfc-editor.org/info/rfc2914>>.
- [RFC3168] - Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, DOI 10.17487/RFC3168, September 2001, <<http://www.rfc-editor.org/info/rfc3168>>.
- [RFC3376] - Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, DOI 10.17487/RFC3376, October 2002, <<http://www.rfc-editor.org/info/rfc3376>>.
- [RFC4301] - Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <<http://www.rfc-editor.org/info/rfc4301>>.
- [RFC4303] - Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, DOI 10.17487/RFC4303, December 2005, <<http://www.rfc-editor.org/info/rfc4303>>.

- [RFC5405] - Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, DOI 10.17487/RFC5304, October 2008, <<http://www.rfc-editor.org/info/rfc5304>>.
- [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, 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>>.
- [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>>.
- [RFC7177] - Eastlake 3rd, D., Perlman, R., Ghanwani, A., Yang, H., and V. Manral, "Transparent Interconnection of Lots of Links (TRILL): Adjacency", RFC 7177, DOI 10.17487/RFC7177, May 2014, <<http://www.rfc-editor.org/info/rfc7177>>.
- [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>>.
- [RFC7321] - McGrew, D. and P. Hoffman, "Cryptographic Algorithm Implementation Requirements and Usage Guidance for Encapsulating Security Payload (ESP) and Authentication Header (AH)", RFC 7321, DOI 10.17487/RFC7321, August 2014, <<http://www.rfc-editor.org/info/rfc7321>>.
- [RFC7348] - Mahalingam, M., Dutt, D., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., and C. Wright, "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", RFC 7348, DOI 10.17487/RFC7348, August 2014, <<http://www.rfc-editor.org/info/rfc7348>>.
- [rfc7180bis] - Eastlake, D., et al, "TRILL: Clarifications, Corrections, and Updates", draft-ietf-trill-rfc7180bis, work in progress.

Informative References

- [RFC4787] - Audet, F., Ed., and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, DOI 10.17487/RFC4787, January 2007, <<http://www.rfc-editor.org/info/rfc4787>>.
- [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>>.
- [RFC6864] - Touch, J., "Updated Specification of the IPv4 ID Field", RFC 6864, DOI 10.17487/RFC6864, February 2013, <<http://www.rfc-editor.org/info/rfc6864>>.
- [RFC6936] - Fairhurst, G. and M. Westerlund, "Applicability Statement for the Use of IPv6 UDP Datagrams with Zero Checksums", RFC 6936, DOI 10.17487/RFC6936, April 2013, <<http://www.rfc-editor.org/info/rfc6936>>.
- [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>>.
- [RFC7173] - Yong, L., Eastlake 3rd, D., Aldrin, S., and J. Hudson, "Transparent Interconnection of Lots of Links (TRILL) Transport Using Pseudowires", RFC 7173, DOI 10.17487/RFC7173, May 2014, <<http://www.rfc-editor.org/info/rfc7173>>.
- [RFC7296] - Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, RFC 7296, DOI 10.17487/RFC7296, October 2014, <<http://www.rfc-editor.org/info/rfc7296>>.
- [circuit-breaker] - Fairhurst, G., "Network Transport Circuit Breakers", draft-ietf-tsvwg-circuit-breaker, work in progress.
- [gre-in-udp] - Crabbe, E., Yong, L., and X. Xu, "Generic UDP Encapsulation for IP Tunneling", draft-yong-tsvwg-gre-in-udp-encap, work in progress.
- [LinkSec] - Eastlake, D., D. Zhang, "TRILL: Link Security", draft-

eastlake-trill-link-security, work in progress.

Acknowledgements

The following people have provided useful feedback on the contents of this document: Sam Hartman, Adrian Farrel, and Mohammed Umair.

Some material in Section 10.2 is derived from draft-ietf-mppls-in-udp by Xiaohu Xu, Nischal Sheth, Lucy Yong, Carlos Pignataro, and Yongbing Fan.

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

Authors' Addresses

Margaret Cullen
Painless Security
356 Abbott Street
North Andover, MA 01845
USA

Phone: +1 781 405-7464
Email: margaret@painless-security.com
URI: <http://www.painless-security.com>

Donald Eastlake
Huawei Technologies
155 Beaver Street
Milford, MA 01757
USA

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

Mingui Zhang
Huawei Technologies
No.156 Beiqing Rd. Haidian District,
Beijing 100095 P.R. China

EMail: zhangmingui@huawei.com

Dacheng Zhang
Alibaba
Beijing, Chao yang District
P.R. China

Email: dacheng.zdc@alibaba-inc.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.

