

Internet Engineering Task Force (IETF)
Request for Comments: 8302
Category: Standards Track
ISSN: 2070-1721

Y. Li
D. Eastlake 3rd
L. Dunbar
Huawei Technologies
R. Perlman
Dell EMC
M. Umair
Cisco
January 2018

Transparent Interconnection of Lots of Links (TRILL):
ARP and Neighbor Discovery (ND) Optimization

Abstract

This document describes mechanisms to optimize the Address Resolution Protocol (ARP) and Neighbor Discovery (ND) traffic in a Transparent Interconnection of Lots of Links (TRILL) campus. TRILL switches maintain a cache of IP / Media Access Control (MAC) address / Data Label bindings that are learned from ARP/ND requests and responses that pass through them. In many cases, this cache allows an edge Routing Bridge (RBridge) to avoid flooding an ARP/ND request by either responding to it directly or encapsulating it and unicasting it. Such optimization reduces packet flooding over a TRILL campus.

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 <https://www.rfc-editor.org/info/rfc8302>.

Copyright Notice

Copyright (c) 2018 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 (<https://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.

Table of Contents

1. Introduction	3
1.1. Terminology	3
2. ARP/ND Optimization Requirement and Solution	4
3. IP/MAC Address Mappings	5
4. Handling ARP/ND/SEND Messages	6
4.1. SEND Considerations	7
4.2. Address Verification	7
4.3. Extracting Local Mapping Information for End-Station IP/MAC Addresses	8
4.4. Determining How to Reply to ARP/ND	9
4.5. Determining How to Handle the ARP/ND Response	10
5. Handling of Reverse Address Resolution Protocol (RARP) Messages	11
6. Handling of DHCP Messages	11
7. Handling of Duplicate IP Addresses	11
8. RBridge ARP/ND Cache Liveness and MAC Mobility	12
9. Security Considerations	13
9.1. Data-Plane-Based Considerations	13
9.2. Directory-Based Considerations	14
9.3. Use of the Confidence Level Feature	14
10. IANA Considerations	14
11. References	15
11.1. Normative References	15
11.2. Informative References	16
Acknowledgments	17
Authors' Addresses	18

1. Introduction

ARP [RFC826] and ND [RFC4861] messages are normally sent by broadcast and multicast, respectively. To reduce the burden on a TRILL campus caused by these multi-destination messages, RBridges MAY implement an "optimized ARP/ND response", as specified herein, when the target's location is known by the ingress RBridge or can be obtained from a directory. This avoids ARP/ND query and answer flooding.

1.1. Terminology

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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

The abbreviations and terminology in [RFC6325] are used herein. Some of these are listed below for convenience along with some additions:

APPsub-TLV	Application sub-Type-Length-Value [RFC6823]
ARP	Address Resolution Protocol [RFC826]
Campus	A TRILL network consisting of RBridges, links, and possibly bridges bounded by end stations and IP routers [RFC6325]
DAD	Duplicate Address Detection [RFC3756] [RFC5227]
Data Label	VLAN or Fine-Grained Label (FGL)
DHCP	Dynamic Host Configuration Protocol. In this document, DHCP refers to both DHCP for IPv4 [RFC2131] and DHCPv6 [RFC3315]
ESADI	End Station Address Distribution Information [RFC7357]
FGL	Fine-Grained Label [RFC7172]
IA	Interface Address; a TRILL APPsub-TLV [RFC7961]
IP	Internet Protocol, both IPv4 and IPv6
MAC	Media Access Control [RFC7042]
ND	Neighbor Discovery [RFC4861]

RBridge	A contraction of "Routing Bridge". A device implementing the TRILL protocol.
SEND	Secure Neighbor Discovery [RFC3971]
TRILL	Transparent Interconnection of Lots of Links or Tunneled Routing in the Link Layer [RFC6325] [RFC7780]

2. ARP/ND Optimization Requirement and Solution

IP address resolution can create significant issues in data centers due to flooded packets, as discussed in [RFC6820]. Such flooding can be avoided by a proxy ARP/ND function on edge RBridges as described in this document, particularly in Section 4. This section is a general discussion of this problem and is not intended to be normative.

To support such ARP/ND optimization, edge RBridges need to know an end station's IP/MAC address mapping through manual configuration (management), control-plane mechanisms such as directories [RFC8171], or data-plane learning by snooping of messages such as ARP/ND (including DHCP or gratuitous ARP messages).

When all the end station's IP/MAC address mappings are known to edge RBridges, provisioned through management, or learned via the control plane on the edge RBridges, it should be possible to completely suppress flooding of ARP/ND messages in a TRILL campus. When all end station MAC addresses are similarly known, it should be possible to suppress unknown unicast flooding by dropping any unknown unicast received at an edge RBridge.

An ARP/ND optimization mechanism should include provisions for an edge RBridge to issue an ARP/ND request to an attached end station to confirm or update information and should allow an end station to detect duplication of its IP address.

Most of the end station hosts send either DHCP messages requesting an IP address or gratuitous ARP or Reverse Address Resolution Protocol (RARP) requests to announce themselves to the network right after they come online. Thus, the local edge RBridge will immediately have the opportunity to snoop and learn their MAC and IP addresses and distribute this information to other edge RBridges through the TRILL control-plane End Station Address Distribution Information (ESADI) [RFC7357] protocol. Once remote RBridges receive this information via the control plane, they should add IP-to-MAC mapping information to their ARP/ND cache along with the nickname and Data Label of the address information. Therefore, most active IP hosts in the TRILL

network can be learned by the edge RBridges through either local learning or control-plane-based remote learning. As a result, ARP suppression can vastly reduce the network flooding caused by host ARP learning behavior.

When complete directory information is available, the default data-plane learning of end-station MAC addresses is not only unnecessary but could be harmful if there is learning based on frames with forged source addresses. Such data-plane learning can be suppressed because TRILL already provides an option to disable data-plane learning from the source MAC address of end-station frames (see Section 5.3 of [RFC6325]).

3. IP/MAC Address Mappings

By default, an RBridge [RFC6325] [RFC7172] learns egress nickname mapping information for the MAC address and Data Label (VLAN or FGL) of TRILL data frames it receives and decapsulates. No IP address information is learned directly from the TRILL data frame. The IA APPsub-TLV [RFC7961] enhances the TRILL base protocol by allowing IP/MAC address mappings to be distributed in the control plane by any RBridge. This APPsub-TLV appears inside the TRILL GENINFO TLV in ESADI [RFC7357], but the value data structure it specifies may also occur in other application contexts. Edge RBridge Directory Assist Mechanisms [RFC8171] make use of this APPsub-TLV for its push model and use the value data structure it specifies in its pull model.

An RBridge can easily know the IP/MAC address mappings of the local end stations that it is attached to via its access ports by receiving ARP [RFC826] or ND [RFC4861] messages. If the edge RBridge has extracted the sender's IP/MAC address pair from the received data frame (either ARP or ND), it may save the information and then use the IA APPsub-TLV to link the IP and MAC addresses and distribute it to other RBridges through ESADI. Then, the relevant remote RBridges (normally those interested in the same Data Label as the original ARP/ND messages) also receive and save such mapping information. There are other ways that RBridges save IP/MAC address mappings in advance, e.g., importing them from the management system and distributing them by directory servers [RFC8171].

The examples given above show that RBridges might have saved an end station's triplet of {IP address, MAC address, ingress nickname} for a given Data Label (VLAN or FGL) before that end station sends or receives any real data packet. Note that such information might or might not be a complete list and might or might not exist on all RBridges; the information could possibly be from different sources. RBridges can then use the Flags field in an IA APPsub-TLV to identify

if the source is a directory server or local observation by the sender. A different confidence level may also be used to indicate the reliability of the mapping information.

4. Handling ARP/ND/SEND Messages

A native frame that is an ARP [RFC826] message is detected by its Ethertype of 0x0806. A native frame that is an ND [RFC4861] is detected by being one of five different ICMPv6 packet types. ARP/ND is commonly used on a link to (1) query for the MAC address corresponding to an IPv4 or IPv6 address, (2) test if an IPv4/IPv6 address is already in use, or (3) announce the new or updated info on any of the following: IPv4/IPv6 address, MAC address, and/or point of attachment.

To simplify the text, we use the following terms in this section.

1. IP address -- indicated protocol address that is normally an IPv4 address in ARP or an IPv6 address in ND.
2. sender's IP/MAC address -- sender IP/MAC address in ARP, source IP address, and source link-layer address in ND.
3. target's IP/MAC address -- target IP/MAC address in ARP, target address, and target link-layer address in ND.

When an ingress RBridge receives an ARP/ND/SEND message, it can perform the steps described within the subsections below. In particular, Section 4.4 describes the options for such an ingress handling an ARP/ND message and, in the cases where it forwards the message, Section 4.5 describes how to handle any response that may be returned due to the forwarded message.

Section 4.3 describes the extraction of address information by an RBridge from ARP/ND messages it handles. Under some circumstances, this extraction may prompt verification with an end station. Section 4.2 describes an optional use of ARP/ND messages originated by RBridges to verify addresses or liveness.

As described in Section 4.1, SEND messages are not optimized by the mechanisms specified in this document but are snooped on.

4.1. SEND Considerations

Secure Neighbor Discovery (SEND) [RFC3971] is a method of securing ND that addresses the threats discussed in [RFC3756]. Typical TRILL campuses are inside data centers, Internet exchange points, or carrier facilities. These are generally controlled and protected environments where these threats are of less concern. Nevertheless, SEND provides an additional layer of protection.

Secure SEND messages require knowledge of cryptographic keys. Methods of communicating such keys to RBridges for use in SEND are beyond the scope of this document. Thus, using the optimizations in this document, RBridges do not attempt to construct SEND messages and are generally transparent to them. RBridges only construct ARP, RARP, or insecure ND messages, as appropriate. Nevertheless, RBridges implementing ARP/ND optimization SHOULD snoop on SEND messages to extract the addressing information that would be present if the SEND had been sent as an insecure ND message and is still present in the SEND message.

4.2. Address Verification

RBridges may use ARP/ND to probe directly attached or remote end stations for address or liveness verification. This is typically most appropriate in less-managed and/or higher-mobility environments. In strongly managed environments, such as a typical data center, where a central orchestration/directory system has complete addressing knowledge [RFC7067], optimized ARP/ND responses can use that knowledge. In such cases, there is little reason for verification except for debugging operational problems or the like.

4.3. Extracting Local Mapping Information for End-Station IP/MAC Addresses

Edge RBridges extract and use information about the correspondence between local end-station IP and MAC addresses from the ARP/ND messages those end stations send, as described below. An apparent zero source IP address means that the end station is probing for duplicate IP addresses, and messages with such a zero source IP address are not used for the extraction of IP/MAC address mapping information.

- o If the sender's IP is not present in the ingress RBridge's ARP/ND cache, populate the information of the sender's IP/MAC address mapping in its ARP/ND cache table. The ingress RBridge correlates its nickname and that IP/MAC address mapping information. Such a triplet of {IP address, MAC address, ingress nickname} information is saved locally and can be distributed to other RBridges, as explained later.
- o Else, if the sender's IP has been saved before but with a different MAC address mapped or a different ingress nickname associated with the same pair of IP/MAC, the RBridge SHOULD verify if a duplicate IP address has already been in use or an end station has changed its attaching RBridge. The RBridge may use different strategies to do so. For example, the RBridge might ask an authoritative entity like directory servers or it might encapsulate and unicast the ARP/ND message to the location where it believes the address is in use (Section 4.2). RBridge SHOULD update the saved triplet of {IP address, MAC address, ingress nickname} based on the verification results. An RBridge might not verify an IP address if the network manager's policy is to have the network behave, for each Data Label, as if it were a single link and just believe an ARP/ND it receives.

The ingress RBridge MAY use the IA APPsub-TLV [RFC7961] with the Local flag set in ESADI [RFC7357] to distribute any new or updated triplet of {IP address, MAC address, ingress nickname} information obtained. If a Push Directory server is used, such information can be distributed as specified in [RFC8171].

4.4. Determining How to Reply to ARP/ND

The options for an edge RBridge to handle a native ARP/ND are given below. For generic ARP/ND requests seeking the MAC address corresponding to an IP address, if the edge RBridge knows the IP address and corresponding MAC, behavior is as in item (a), otherwise behavior is as in item (b). Behavior for gratuitous ARP and ND unsolicited Neighbor Advertisements (NAs) [RFC4861] is given in item (c). And item (d) covers the handling of an Address Probe ARP query. Within each lettered item, it is an implementation decision as to which numbered strategy to use for any particular ARP/ND query falling under that item.

- a. If the message is a generic ARP/ND request, and the ingress RBridge knows the target's IP address and associated MAC address, the ingress RBridge MUST take one or a combination of the actions below. In the case of SEND [RFC3971], cryptography would prevent a local reply by the ingress RBridge, since the RBridge would not be able to sign the response with the target's private key, and only action a.2 or a.5 is valid.
 - a.1. Send an ARP/ND response directly to the querier, using the target's MAC address present in the ingress RBridge's ARP/ND cache table. Because the edge RBridge might not have an IPv6 address, the source IP address for such an ND response MUST be that of the target end station.
 - a.2. Encapsulate the ARP/ND/SEND request to the target's Designated RBridge and have the egress RBridge for the target forward the query to the target. This behavior has the advantage that a response to the request is authoritative. If the request does not reach the target, then the querier does not get a response.
 - a.3. Block ARP/ND requests that occur for some time after a request to the same target has been launched, and then respond to the querier when the response to the recently launched query to that target is received.
 - a.4. Reply to the querier based on directory information [RFC8171] such as information obtained from a Pull Directory server or directory information that the ingress RBridge has requested to be pushed to it.
 - a.5. Flood the ARP/ND/SEND request as per [RFC6325].

- b. If the message is a generic ARP/ND/SEND request and the ingress RBridge does not know the target's IP address, the ingress RBridge MUST take one of the following actions. In the case of SEND [RFC3971], cryptography would prevent local reply by the ingress RBridge, since the RBridge would not be able to sign the response with the target's private key; therefore, only action b.1 is valid.
 - b.1. Flood the ARP/ND/SEND message as per [RFC6325].
 - b.2. Use a directory server to pull the information [RFC8171] and reply to the querier.
 - b.3. Drop the message if there should be no response because the directory server gives authoritative information that the address being queried is nonexistent.
- c. If the message is a gratuitous ARP, which can be identified by the same sender's and target's "protocol" address fields, or an unsolicited Neighbor Advertisement [RFC4861] in ND/SEND then:

The RBridge MAY use an IA APPsub-TLV [RFC7961] with the Local flag set to distribute the sender's IP/MAC address mapping information. When one or more directory servers are deployed and complete Push Directory information is used by all the RBridges in the Data Label, a gratuitous ARP or unsolicited NA SHOULD be discarded rather than ingressed. Otherwise, they are either ingressed and flooded as per [RFC6325] or discarded depending on local policy.
- d. If the message is an Address Probe ARP query [RFC5227], which can be identified by the sender's protocol (IPv4) address field being zero and the target's protocol address field being the IPv4 address to be tested or a Neighbor Solicitation for Duplicate Address Detection (DAD) that has the unspecified source address [RFC4862], it SHOULD be handled as the generic ARP message as in (a) or (b) above.

4.5. Determining How to Handle the ARP/ND Response

If the ingress RBridge R1 decides to unicast the ARP/ND request to the target's egress RBridge R2 as discussed in Section 4.4, item a.2 or to flood the request as per item a.5 and [RFC6325], then R2 decapsulates the query and initiates an ARP/ND query on the target's link. If and when the target responds, R2 encapsulates and unicasts the response to R1, which decapsulates the response and sends it to the querier. R2 SHOULD initiate a link state update to inform all the other RBridges of the target's location, Layer 3 address, and

Layer 2 address, in addition to forwarding the reply to the querier. The update uses an IA APPsub-TLV [RFC7961] (so the Layer 3 and Layer 2 addresses can be linked) with the Local flag set in ESADI [RFC7357] or as per [RFC8171] if the Push Directory server is in use.

5. Handling of Reverse Address Resolution Protocol (RARP) Messages

RARP [RFC903] uses the same packet format as ARP but different Ethertype (0x8035) and opcode values. Its processing is similar to the generic ARP request/response as described in Section 4.4, items a. and b. The difference is that it is intended to query for the target "protocol" (IP) address corresponding to the target "hardware" (MAC) address provided. It SHOULD be handled by doing a local cache or directory server lookup on the target "hardware" address provided to find a mapping to the desired "protocol" address.

6. Handling of DHCP Messages

When a newly connected end station exchanges messages with a DHCP [RFC3315][RFC2131] server, an edge RBridge should snoop them (mainly the DHCPACK message) and store IP/MAC address mapping information in its ARP/ND cache and should also send the information out through the TRILL control plane using ESADI.

7. Handling of Duplicate IP Addresses

Duplicate IP addresses within a Data Label can occur due to an attacker sending fake ARP/ND messages or due to human/configuration errors. If complete, trustworthy directory information is available, then, by definition, the IP location information in the directory is correct. Any appearance of an IP address in a different place (different edge RBridge or port) from other sources is not correct.

Without complete directory information, the ARP/ND optimization function should support duplicate IP detection. This is critical in a data center to stop an attacker from using ARP/ND spoofing to divert traffic from its intended destination.

Duplicate IP addresses can be detected when an existing active IP/MAC address mapping gets modified. Also, an edge RBridge may send a query called a Duplicate Address Detection query (DAD-query) asking about the IP address in question to the former owner of that IP address by using the MAC address formerly associated with that IP address. A DAD-query is a unicast ARP/ND message with sender IP 0.0.0.0 in case of ARP (or a configurable IP address per RBridge called the DAD-Query source IP) and an IPv6 Link Local Address in case of ND with the source MAC set to the DAD-querier RBridge's MAC. If the querying RBridge does not receive an answer within a given

time, it may be a case of mobility; in any case, the new IP entry will be confirmed and activated in its ARP/ND cache.

In the case where the former owner replies, a duplicate address has been detected. In this case, the querying RBridge SHOULD log the duplicate so that the network administrator can take appropriate action.

It is an implementation choice how to respond to a query for an address that is duplicated in the network when authoritative information is not available from a directory or configuration. Typically, the information most recently snooped is returned.

8. RBridge ARP/ND Cache Liveness and MAC Mobility

A maintenance procedure is needed for RBridge ARP/ND caching to ensure IP end stations connected to ingress RBridges are still active.

Some links provide a physical-layer indication of link liveness. A dynamic proxy ARP/ND entry (one learned from data-plane observation) MUST be removed from the table if the link over which it was learned fails.

Similarly, a dynamic proxy ARP/ND entry SHOULD be flushed out of the table if the IP/MAC address mapping has not been refreshed within a given age-time. The entry is refreshed if an ARP or ND message is received for the same IP/MAC address mapping entry from any location. The IP/MAC address mapping information Ageing Timer is configurable per RBridge and defaults to 3/4 of the MAC address learning Ageing Timer [RFC6325].

For example, end station "A" is connected to edge-RBridge1 (RB1) and has been learned as a local entry on RB1. If end station "A" moves to some other location (MAC / Virtual Machine (VM) Mobility) and gets connected to edge-RBridge (RB2), after learning on RB2's access port, RB2 advertises this entry through the TRILL control plane, and it is learned on RB1 as a remote entry. The old entry on RB1 SHOULD get replaced, and all other edge-RBridges with end-station service enabled for that Data Label should update the entry to show reachability from RB2 instead of RB1.

If an ARP/ND entry in the cache is not refreshed, then the RBridge connected to that end station MAY send periodic refresh messages (ARP/ND "probes") to that end station, so that the entries can be refreshed before they age out. The end station would reply to the ARP/ND probe, and the reply resets the corresponding entry age-timer.

9. Security Considerations

There are generally two modes of learning the address information that is the basis of ARP/ND optimization: data-plane mode and directory mode. The data-plane mode is the traditional bridge address learning [IEEE802.1Q] that is also implemented in TRILL switches [RFC6325] and is discussed in Section 9.1. The directory mode uses data obtained from a directory [RFC8171] and is discussed in Section 9.2. The TRILL confidence-level feature, which can help arbitrate between conflicting address information, is discussed in Section 9.3.

RBridges should rate limit ARP/ND queries injected into the TRILL campus to limit some potential denial-of-service attacks.

9.1. Data-Plane-Based Considerations

Generally speaking, when ARP/ND optimization is operating in the data-plane mode, the information learned by RBridges is the same as that which is learned by end stations. Thus, the answers generated by RBridges to the query messages being optimized are generally those that would be generated by end stations in the absence of optimization, and the security considerations are those of the underlying ARP/ND protocols.

RBridges that snoop on DHCPack messages respond to ARP/ND messages in essentially the same way that the end stations sending those DHCPack messages would. Thus, for security considerations of ARP/ND optimization for DHCP messages that may be snooped, see the Security Considerations sections of [RFC3315] and [RFC2131].

Unless SEND [RFC3971] is used, ARP and ND messages can be easily forged. Therefore, the learning of IP/MAC addresses by RBridges from ARP/ND is hackable, but this is what is available for data-plane learning without SEND. See "SEND Considerations", Section 4.1.

Since end stations communicate with edge RBridges using Ethernet, some security improvements could be obtained by the use of [IEEE802.1AE] between end stations and edge RBridges. Such link security is beyond the scope of this document and would impose requirements on edge stations, while TRILL is generally designed to operate with unmodified, TRILL-ignorant end stations.

ARP/ND address mapping information learned locally at an RBridge can be distributed to other RBridges using the TRILL ESADI protocol that can be secured as specified in [RFC7357]. (ESADI is also used for Push Directories with flags in the data indicating whether data comes from a directory or from data-plane learning, as well as from a confidence level (see Section 9.3).)

9.2. Directory-Based Considerations

ARP/ND optimization can be based on directory information [RFC8171]. If the directory information is known to be trustworthy and complete, then trustworthy responses to ARP/ND queries can be entirely based on this information. This bounds the damage that forged ARP/ND messages can do to the local link between end stations and edge RBridges. (In TRILL, such a "link" can be a bridged LAN.)

Of course, there can also be incomplete and/or unreliable directory address mapping data. The network administrator can configure their TRILL campus to use such directory data in place of data-plane-learned data. Alternatively, such directory data can be used along with data-plane-learned data arbitrated by the confidence level as discussed in Section 9.3.

9.3. Use of the Confidence Level Feature

An RBridge can use the confidence level in IA APPsub-TLV information received via ESADI or Pull Directory retrievals to determine the configured relative reliability of IP/MAC address mapping information from those sources and from locally learned address information. Push Directory information is sent via ESADI, which can be secured as provided in [RFC7357]; Pull Directory information can be secured as provided in [RFC8171]. The implementation decides if an RBridge will distribute the IP and MAC address mappings received from local native ARP/ND messages to other RBridges in the same Data Label, and with what confidence level it does so. Thus, the implementer can, to some extent, cause sources that they know are more reliable to dominate those they know to be less reliable. How the implementer determines this is beyond the scope of this document.

10. IANA Considerations

This document does not require any IANA actions.

11. References

11.1. Normative References

- [RFC826] Plummer, D., "Ethernet Address Resolution Protocol: Or Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware", STD 37, RFC 826, DOI 10.17487/RFC0826, November 1982, <<https://www.rfc-editor.org/info/rfc826>>.
- [RFC903] Finlayson, R., Mann, T., Mogul, J., and M. Theimer, "A Reverse Address Resolution Protocol", STD 38, RFC 903, DOI 10.17487/RFC0903, June 1984, <<https://www.rfc-editor.org/info/rfc903>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <<https://www.rfc-editor.org/info/rfc2131>>.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <<https://www.rfc-editor.org/info/rfc3315>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<https://www.rfc-editor.org/info/rfc4862>>.
- [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, <<https://www.rfc-editor.org/info/rfc6325>>.

- [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, <<https://www.rfc-editor.org/info/rfc7172>>.
- [RFC7357] Zhai, H., Hu, F., Perlman, R., Eastlake 3rd, D., and O. Stokes, "Transparent Interconnection of Lots of Links (TRILL): End Station Address Distribution Information (ESADI) Protocol", RFC 7357, DOI 10.17487/RFC7357, September 2014, <<https://www.rfc-editor.org/info/rfc7357>>.
- [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, <<https://www.rfc-editor.org/info/rfc7780>>.
- [RFC7961] Eastlake 3rd, D. and L. Yizhou, "Transparent Interconnection of Lots of Links (TRILL): Interface Addresses APPsub-TLV", RFC 7961, DOI 10.17487/RFC7961, August 2016, <<https://www.rfc-editor.org/info/rfc7961>>.
- [RFC8171] Eastlake 3rd, D., Dunbar, L., Perlman, R., and Y. Li, "Transparent Interconnection of Lots of Links (TRILL): Edge Directory Assistance Mechanisms", RFC 8171, DOI 10.17487/RFC8171, June 2017, <<https://www.rfc-editor.org/info/rfc8171>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

11.2. Informative References

- [IEEE802.1AE]
IEEE, "IEEE Standard for Local and metropolitan area networks: Media Access Control (MAC) Security", IEEE Std 802.1AE.
- [IEEE802.1Q]
IEEE, "IEEE Standard for Local and metropolitan area networks -- Bridges and Bridged Networks", IEEE Std 802.1Q.

- [RFC3756] Nikander, P., Ed., Kempf, J., and E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats", RFC 3756, DOI 10.17487/RFC3756, May 2004, <<https://www.rfc-editor.org/info/rfc3756>>.
- [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, DOI 10.17487/RFC3971, March 2005, <<https://www.rfc-editor.org/info/rfc3971>>.
- [RFC5227] Cheshire, S., "IPv4 Address Conflict Detection", RFC 5227, DOI 10.17487/RFC5227, July 2008, <<https://www.rfc-editor.org/info/rfc5227>>.
- [RFC6820] Narten, T., Karir, M., and I. Foo, "Address Resolution Problems in Large Data Center Networks", RFC 6820, DOI 10.17487/RFC6820, January 2013, <<https://www.rfc-editor.org/info/rfc6820>>.
- [RFC6823] Ginsberg, L., Previdi, S., and M. Shand, "Advertising Generic Information in IS-IS", RFC 6823, DOI 10.17487/RFC6823, December 2012, <<https://www.rfc-editor.org/info/rfc6823>>.
- [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, <<https://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, <<https://www.rfc-editor.org/info/rfc7067>>.

Acknowledgments

The authors would like to thank Igor Gashinsky and Sue Hares for their contributions.

Authors' Addresses

Yizhou Li
Huawei Technologies
101 Software Avenue,
Nanjing 210012
China
Phone: +86-25-56625375
Email: liyizhou@huawei.com

Donald Eastlake 3rd
Huawei R&D USA
155 Beaver Street
Milford, MA 01757
United States of America
Phone: +1-508-333-2270
Email: d3e3e3@gmail.com

Linda Dunbar
Huawei Technologies
5430 Legacy Drive, Suite #175
Plano, TX 75024
United States of America
Phone: +1-469-277-5840
Email: ldunbar@huawei.com

Radia Perlman
Dell EMC
176 South Street
Hopkinton, MA 01748
United States of America
Email: Radia@alum.mit.edu

Mohammed Umair
Cisco
Cessna Business Park, Kadubeesanahalli Village, Hobli,
Sarjapur, Varthur Main Road, Marathahalli,
Bengaluru, Karnataka 560087
India
Email: mohammed.umair2@gmail.com