

Internet Engineering Task Force (IETF)
Request for Comments: 8397
Category: Standards Track
ISSN: 2070-1721

M. Zhang
D. Eastlake 3rd
Huawei
R. Perlman
Dell EMC
H. Zhai
JIT
D. Liu
China Telecom Co., Ltd
May 2018

Transparent Interconnection of Lots of Links (TRILL) Multilevel Using Unique Nicknames

Abstract

TRILL (Transparent Interconnection of Lots of Links) routing can be extended to support multiple levels by building on the multilevel feature of IS-IS routing. Depending on how nicknames are managed, there are two primary alternatives to realize TRILL multilevel: the unique nickname approach and the aggregated nickname approach as discussed in RFC 8243. This document specifies a unique nickname approach. This approach gives unique nicknames to all TRILL switches across the multilevel 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/rfc8397>.

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
2. Acronyms and Terminology	4
3. Data Routing	4
3.1. Unicast Routing	4
3.2. Multi-destination Routing	5
3.2.1. Local Distribution Trees	6
3.2.2. Global Distribution Trees	6
4. Protocol Basics and Extensions	8
4.1. Multilevel TRILL Basics	8
4.2. Nickname Allocation	9
4.3. Nickname Announcements	9
4.4. Capability Indication	11
5. Mix with Aggregated Nickname Areas	11
6. Security Considerations	12
7. IANA Considerations	13
8. References	13
8.1. Normative References	13
8.2. Informative References	14
Contributors	15
Authors' Addresses	15

1. Introduction

The multiple-level feature of [IS-IS] can increase the scalability of TRILL as discussed in [RFC8243]. However, multilevel IS-IS needs some extensions to support the TRILL multilevel feature. The two most significant extensions are how TRILL switch nicknames are managed and how distribution trees are handled [RFC8243].

There are two primary alternatives to realize TRILL multilevel [RFC8243]. One approach, which is referred to as the "aggregated nickname" approach, involves assigning nicknames to the areas, and allowing nicknames to be reused in different areas by having the border TRILL switches rewrite nickname fields when entering or leaving an area. For more description of the aggregated nickname approach, one can refer to [RFC8243] and [SingleN]. The other approach, which is referred to as the "unique nickname" approach, is specified in this document. The unique nickname approach gives unique nicknames to all the TRILL switches in the multilevel campus by having the TRILL switches at the Level 1 / Level 2 border advertise into the Level 1 area those nicknames are not available for assignment in that area and advertising into the Level 2 area those nicknames that are used by the Level 1 area so that other areas cannot use them anymore. The advertising of Level 1 nicknames informs the rest of the campus how to reach the nicknames residing in that area. In this document, protocol extensions that support such advertisement are specified.

Each RBridge in a unique nickname area calculates two types of trees: local distribution trees and global distributions trees. For multi-destination traffic that is limited to an area, the packets will be flooded on a local distribution tree. Otherwise, the multi-destination packets will be flooded along a global distribution tree.

In the unique nickname approach, nicknames are globally valid so that border RBridges do not rewrite the nickname field of TRILL data packets that transition between Level 1 and Level 2, as border RBridges do in the aggregated nickname approach. If a border RBridge is a transit node on a forwarding path, it does not learn MAC addresses of the TRILL data packets forwarded along this path. Testing and maintenance operations that originate in one area and terminate in a different area are also simplified [RFC8243]. For these reasons, the unique nickname approach might realize simpler border RBridges than the aggregated nickname approach. However, the unique nickname approach is less scalable and may be less well suited for very large campuses.

Let us say that S transmits a frame to destination D and let us say that D's location is learned by the relevant TRILL switches already. These relevant switches have learned the following:

- 1) RB27 has learned that D is connected to nickname 44.
- 2) RB2 has learned that nickname 44 is accessible through RB3.

The following sequence of events will occur:

- S transmits an Ethernet frame with source MAC = S and destination MAC = D.
- RB27 encapsulates with a TRILL header with ingress RBridge nickname 27, and egress RBridge nickname 44 producing a TRILL Data packet.
- RB2 has announced in the Level 1 IS-IS instance in Area X that it owns all nicknames of other areas, including 44. Therefore, IS-IS routes the packet to RB2.
- The packet is forwarded through Level 2, from RB2 to RB3, which has advertised, in Level 2, it owns the nickname 44.
- RB3, when forwarding into Area Y, does not change the ingress nickname 27 or the egress nickname 44.
- RB44, when decapsulating, learns that S is attached to nickname 27.

3.2. Multi-destination Routing

The scope of Multi-destination routing is defined by the tree root nickname. A tree with a Level 2 tree root nickname is global, and a tree with a Level 1 tree root nickname is local. See Section 4.2 for the Level 1 and Level 2 nickname allocation.

Border RBridges announce the global trees to be calculated only for those Data Labels that span across areas. APPsub-TLVs as specified in Section 3.2 of [RFC7968] will be advertised for this purpose. Based on the Data Label, an ingress RBridge can determine whether a global tree or a local tree is to be used for a TRILL multi-destination Data packet.

If there are legacy TRILL switches that do not understand the APPsub-TLVs for tree selection, configuration MUST guarantee that Data Labels [RFC7172] being used globally in Level 2 are disabled on these legacy TRILL switches. (Otherwise, the legacy TRILL switches might use local trees for multi-destination traffic with a global scope.)

These legacy TRILL switches may use global trees to flood multi-destination packets with a scope of the local area. Those global trees MUST be pruned at the border TRILL switches based on Data Labels.

3.2.1. Local Distribution Trees

The root RBridge RB1 of a local distribution tree resides in the area. RBridges in this area calculate this local tree based on the link state information of this area, using RB1's nickname as the root. Protocol behaviors for local distribution trees have been specified in Section 4.5 of [RFC6325]. The sole difference is that the local distribution tree spans this area only. A multi-destination packet with an egress nickname of the root RBridge of a local tree MUST NOT be leaked into Level 2 at the border RBridge.

3.2.2. Global Distribution Trees

Within Level 2, the RBridge with the highest tree root priority advertises the set of global trees by providing a list of Level 2 RBridge nicknames as defined in Section 4.5 of [RFC6325].

According to [RFC6325], the RBridge with the highest root priority advertises the tree roots for a Level 1 area. There has to be a border RBridge with the highest root tree priority in each area so that it can advertise the global tree root nicknames into the area. Also, this border RBridge MUST advertise the set of local distribution trees by providing another set of nicknames. Since nicknames of global tree roots and local tree roots indicate different flooding scopes, these two sets MUST NOT overlap. If a border RBridge has been assigned both as a global tree root and a local tree root, it MUST acquire both global tree root nickname(s) and local tree root nickname(s). However, non-border RBridges in an area do not differentiate between a global tree root nickname and a local tree root nickname.

Suppose RB3 is the RBridge with the highest tree root priority within Level 2, and RB2 is the highest tree root priority in Area X. RB2 advertises in Area X that nickname RB3 is the root of a distribution tree. Figures 2 through 5 illustrate how different RBridges view the global distribution tree.




Figure 2: RB27's View of the Global Distribution Tree




Figure 3: RB2's View of the Global Distribution Tree




Figure 4: RB3's View of the Global Distribution Tree




Figure 5: RB44's View of the Global Distribution Tree

The following sequence of events will occur when a multi-destination TRILL Data packet is forwarded using the global distribution tree:

- RB27 produces a multi-destination (M bit is one) TRILL Data packet with ingress RBridge nickname 27 and egress RBridge nickname 3. RB27 floods this packet using the segment of the global distribution tree that resides in Area X.
- RB2, when flooding the packet in Level 2, uses the segment of the global distribution tree that resides in Level 2.
- RB3, when flooding the packet into Area Y, uses the segment of the global distribution tree that resides in Area Y.
- The multicast listener RB44, when decapsulating the received packet, learns that S is attached to nickname 27.

4. Protocol Basics and Extensions

4.1. Multilevel TRILL Basics

Multilevel TRILL builds on the multilevel feature of [IS-IS]. Border RBridges are in both a Level 1 area and in Level 2. They establish adjacency with Level 1 RBridges as specified in [RFC7177] and [RFC6325]. They establish adjacency with Level 2 RBridges in exactly the same way except that (1) for a LAN link, the IS-IS Hellos used are Level 2 Hello PDUs [IS-IS] and (2) for a point-to-point link, the Level is configured and indicated in flags in the point-to-point Hello. The state machines for Level 1 and Level 2 adjacency are independent, and two RBridges on the same LAN link can have any adjacency state for Level 1 and, separately, any adjacency state for Level 2. Level 1 and Level 2 link state flooding are independent using Level 1 and Level 2 versions of the relevant IS-IS PDUs (LSP, CSNP, PSNP, FS-LSP, FS-CSNP, and FS-PSNP [RFC7356] [RFC7780]). Thus, Level 1 link state information stays within a Level 1 area and Level 2 link state information stays in Level 2 unless there are specific provisions for leaking (copying) information between levels. This is why multilevel can address the TRILL scalability issues as specified in Section 2 of [RFC8243].

The former "campus wide" minimum acceptable link size S_z is calculated as before: by Level 1 RBridges (including border RBridges) using the `originatingLSPBufferSize` advertised in the Level 1 LSP so it is area local in multilevel TRILL. A minimum acceptable link size in Level 2, called S_{z2} , is calculated by the RBridges participating in Level 2 in the same way as S_z is calculated but using the `originatingLSPBufferSize` distributed in Level 2 LSPs.

4.2. Nickname Allocation

Level 2 RBridges contend for nicknames in the range from 0xF000 through 0xFFBF the same way as specified in [RFC6325]: using Level 2 LSPs. The highest-priority border router for a Level 1 area should contend with others in Level 2 for blocks of nicknames for the range from 0x0001 to 0xEFFF. Blocks of 64 aligned on boundaries of multiples of 64 are RECOMMENDED in this document.

The nickname contention in Level 2 will determine which blocks of nicknames are available for an area and which blocks of nicknames are used elsewhere. The `NickBlockFlags` APPsub-TLV as specified in Section 4.3 will be used by the border RBridge(s) to announce the nickname availability.

4.3. Nickname Announcements

Border RBridges need to exchange nickname information between Level 1 and Level 2; otherwise, forwarding paths inward or outward will not be calculated. For this purpose, border RBridges need to fabricate nickname announcements. Sub-TLVs used for such announcements are specified as follows.

Besides its own nickname(s), a border RBridge MUST announce, in its area, the ownership of all external nicknames that are reachable from this border RBridge. These external nicknames include nicknames used in other unique nickname areas and nicknames in Level 2. Non-border RBridge nicknames within aggregated nickname areas are excluded. Also, a border RBridge MUST announce, in Level 2, the ownership of all nicknames within its area. From listening to these Level 2 announcements, border RBridges can figure out the nicknames used by other areas.

RBridges in the TRILL base protocol use the Nickname Sub-TLV as specified in Section 2.3.2 of [RFC7176] to announce the ownership of nicknames. However, it becomes uneconomic to use this Sub-TLV to announce a mass of internal/external nicknames. To address this issue, border RBridges SHOULD make use of the `NickBlockFlags`

APPsub-TLV to advertise into the Level 1 area the inclusive range of nicknames that are or are not available for self allocation by the Level 1 RBridges in that area. Its structure is as follows:

```

      0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Type = 24      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Length         |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|OK|      RESV        |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Nickname Block 1      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      ...                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Nickname Block K      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

- o Type: 24 (TRILL NickBlockFlags)
- o Length: 2 + 4*K, where K is the number of nickname blocks.
- o OK:
 - When this bit is set to 1, the blocks of nicknames in this APPsub-TLV are associated to the border RBridge's attached Level 1 area. The APPsub-TLV will be advertised in both Level 1 and Level 2. For nicknames that fall in the ranges of the nickname blocks, RBridges of Level 2 always route to the originating border RBridge, just as if this border RBridge owns these nicknames.
 - When this bit is set to 0, it indicates that the nicknames covered by the nickname blocks are being used in Level 2 or other areas so that they are not available for use in the border RBridge's attached Level 1 area. The APPsub-TLV will be advertised into Level 1 only. For nicknames that fall in the ranges of the nickname blocks, RBridges of the area always route to the originating border RBridge, just as if this border RBridge owns these nicknames. For nicknames in these ranges, other RBridges will deem that they are owned by the originating border RBridge. The paths to nicknames that fall in these ranges will be calculated to reach the originating border RBridge. TRILL Data packets with egress nicknames that are neither in these ranges nor announced by any RBridge in the area MUST be discarded.

- o RESV: reserved for future flag allocation. MUST be sent as zero and ignored on receipt.
- o Nickname Block: a starting and ending nickname as follows:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   starting nickname   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   ending nickname    |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Nickname Sub-TLV as specified in Section 2.3.2 of [RFC7176] is still allowed to be used, given the above NickBlockFlags APPsub-TLV is being used.

There might be multiple border RBridges connected to the same area. Each border RBridge may advertise a subset of the entire internal/external nickname space in order to realize load balance. However, optimization of such load balance is an implementation issue and is outside the scope of this document.

As specified in Section 4.2.6 of [RFC6325], multiple border RBridges may claim the same nicknames outwardly and/or inwardly. Other RBridges add those nicknames as if they are attached to all of those border RBridges.

4.4. Capability Indication

All border RBridges MUST understand the NickBlockFlags APPsub-TLV. Non-border RBridges in an area should understand the NickBlockFlags APPsub-TLV. If an RBridge within an area understands the NickBlockFlags APPsub-TLV, it MUST indicate this capability by announcing it in its TRILL-VER Sub-TLV. (See Section 7.)

If there are RBridges that do not understand the NickBlockFlags APPsub-TLV, border RBridges of the area MUST also use the traditional Nickname Sub-TLV [RFC7176] to announce into the area those nicknames covered by the nickname blocks of the NickBlockFlags APPsub-TLV whose OK is 0. The available range of nicknames for this area should be configured on these traditional RBridges.

5. Mix with Aggregated Nickname Areas

The design of TRILL multilevel allows a mixture of unique nickname areas and aggregated nickname areas (see Section 1.2 of [RFC8243]). Usage of nickname space MUST be planned so that nicknames used in any one unique nickname area and Level 2 are never used in any other areas, including unique nickname areas as well as aggregated nickname

areas. In other words, nickname reusage is merely allowed among aggregated nickname areas.

Border RBridges of an aggregated area MUST announce nicknames heard from Level 2 into their area like just like a unique nickname border RBridge. However, these RBridges do not announce nicknames of their area into Level 2.

Each border RBridge of the aggregated areas will appear on the global tree, as specified in Section 4.1, as a single node. The global trees for unique nickname areas span unique nickname areas and Level 2 but never reach the inside of aggregated areas.

6. Security Considerations

Since TRILL multilevel uses the existing IS-IS multilevel facilities [IS-IS], flooding of control traffic for link-state information is automatically confined to a Level 1 area or to Level 2 (except for limited types of information that can be specifically flagged for wider flooding). This addresses the TRILL scalability issues as specified in Section 2 of [RFC8243], and also, except for the wider flooding case, this confines the scope of the effects of malicious events that could be communicated through the link state. However, due to the fact that unique nickname areas share a common nickname space, border RBridges still have to leak nickname information between levels. Such leaking means that nickname-related events in one area can affect other areas.

For this purpose, border RBridges need to fabricate the nickname announcements as specified in Section 4.3. Malicious devices may also fake the NickBlockFlags APPsub-TLV to announce a range of nicknames. By doing this, the attacker can attract TRILL data packets that were originally sent to a bunch of other RBridges. For this reason, RBridges SHOULD be configured to use the IS-IS Authentication TLV (10) in the IS-IS PDUs, particularly those containing the NickBlockFlags APPsub-TLV, so that IS-IS security [RFC5310] can be used to authenticate those PDUs and discard them if they are forged.

If border RBridges do not prune multi-destination distribution tree traffic in Data Labels that are configured to be area local, then traffic that should have been contained within an area might be wrongly delivered to end stations in that Data Label in other areas. That is, the Data Label would no longer be area local. This would generally violate security constraints that require traffic to be delivered only to end stations in that Data Label in a single area.

For general TRILL Security Considerations, see [RFC6325].

7. IANA Considerations

IANA has registered a new flag bit under the "TRILL-VER Sub-TLV Capability Flags" registry.

Bit	Description	Reference
---	-----	-----
5	Able to handle the NickBlockFlags APPsub-TLV	RFC 8397

IANA has assigned a new type for the NickBlockFlags APPsub-TLV from the range available below 256 and add the following entry to the "TRILL APPsub-TLV Types under IS-IS TLV 251 Application Identifier 1" registry as follows:

Type	Name	Reference
----	-----	-----
24	NickBlockFlags	RFC 8397

8. References

8.1. Normative References

- [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>>.
- [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>>.
- [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, <<https://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, <<https://www.rfc-editor.org/info/rfc7177>>.
- [RFC7968] Li, Y., Eastlake 3rd, D., Hao, W., Chen, H., and S. Chatterjee, "Transparent Interconnection of Lots of Links (TRILL): Using Data Labels for Tree Selection for Multi-Destination Data", RFC 7968, DOI 10.17487/RFC7968, August 2016, <<https://www.rfc-editor.org/info/rfc7968>>.
- [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>>.
- [IS-IS] International Organization for Standardization, "Information technology -- Telecommunications and information exchange between systems -- Intermediate System to Intermediate System intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)", ISO/IEC 10589:2002, Second Edition, November 2002.

8.2. Informative References

- [SingleN] Zhang, M., Eastlake, D., et al, "Transparent Interconnection of Lots of Links (TRILL) Single Area Border RBridge Nickname for Multilevel", draft-ietf-trill-multilevel-single-nickname-05, Work in Progress, January 2018.
- [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, <<https://www.rfc-editor.org/info/rfc5310>>.
- [RFC7356] Ginsberg, L., Previdi, S., and Y. Yang, "IS-IS Flooding Scope Link State PDUs (LSPs)", RFC 7356, DOI 10.17487/RFC7356, September 2014, <<https://www.rfc-editor.org/info/rfc7356>>.
- [RFC7780] Eastlake 3rd, D., Zhang, M., Perlman, R., Banerjee, A., Ghanwani, A., and S. Gupta, "Transparent Interconnection of Lots of Links (TRILL): Clarifications, Corrections, and Updates", RFC 7780, DOI 10.17487/RFC7780, February 2016, <<https://www.rfc-editor.org/info/rfc7780>>.

[RFC8243] Perlman, R., Eastlake 3rd, D., Zhang, M., Ghanwani, A.,
and H. Zhai, "Alternatives for Multilevel Transparent
Interconnection of Lots of Links (TRILL)", RFC 8243,
DOI 10.17487/RFC8243, September 2017,
<<https://www.rfc-editor.org/info/rfc8243>>.

Contributors

Margaret Cullen
Painless Security
14 Summer St. Suite 202
Malden, MA 02148
United States of America

Email: margaret@painless-security.com

Authors' Addresses

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

Phone: +86-13810702575
Email: zhangmingui@huawei.com

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

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

Radia Perlman
Dell EMC
176 South Street
Hopkinton, MA 01748
United States of America

Email: radia@alum.mit.edu

Hongjun Zhai
Jinling Institute of Technology
99 Hongjing Avenue, Jiangning District
Nanjing, Jiangsu 211169
China

Email: honjun.zhai@tom.com

Dongxin Liu
China Telecom Co., Ltd
109 West Zhongshan Ave, Tianhe District
Guangzhou 510630
China

Email: liudx@gsta.com

