TRILL Working Group                                         W. Hao
INTERNET-DRAFT                                               Y. Li
Intended Status: Standards Track            Huawei Technologies
                                                    M. Durrani
                                                        Equinix
                                                      S. Gupta
                                                     IP Infusion
                                                          A. Qu
                                                       MediaTec
Expires: July 2018                          January 05, 2018

**Centralized Replication for Active-Active BUM Traffic**
**draft-ietf-trill-centralized-replication-12.txt**

Abstract

   In TRILL active-active access, an RPF (Reverse Path Forwarding) check
   failure issue may occur when using the pseudo-nickname mechanism
   specified in RFC 7781. This draft describes a solution to resolve
   this RPF check failure issue through centralized replication. All
   ingress RBridges send BUM (Broadcast, Unknown unicast and Multicast)
   traffic to a centralized node with unicast TRILL encapsulation.
   When the centralized node receives the BUM traffic, it decapsulates
   the packets and forwards them to their destination RBridges using
   a distribution tree established as per TRILL base protocol RFC 6325.
   To avoid RPF check failure on a RBridge sitting between the ingress
   RBridge and the centralized replication node, some change in the
   RPF calculation algorithm is required. RPF checks on each RBridge
   MUST be calculated as if the centralized node was the ingress
   RBridge, instead of being calculated using the actual ingress
   RBridge.

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/1id-abstracts.html

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

Copyright Notice

Table of Contents

**1. Introduction**

   The IETF TRILL (Transparent Interconnection of Lots of Links)
   [RFC6325] protocol provides loop free and per hop based multipath
   data forwarding with minimum configuration. TRILL uses IS-IS
   [RFC6165] [RFC7176] as its control plane routing protocol and
   defines a TRILL specific header for user data.

   Customer Equipment (CE) devices can be multi-homed to a set of edge
   RBridges forming an edge group where active-active service can be
   provided. In that case, all of the uplinks from a CE are handled via
   a Local Active-Active Link Protocol (LAALP [RFC7379]) such as Multi-
   Chassis Link Aggregation (MC-LAG) or Distributed Resilient Network
   Interconnect (DRNI) [802.1AX]. An active-active flow-based load
   sharing mechanism  can achieve better load balancing and high
   reliability. A CE device can be a layer 3 end system by itself or a
   bridge switch through which layer 3 end systems access to TRILL
   campus.

   In active-active access, the pseudo-nickname solution in [RFC7781]
   can be used to avoid MAC flip-flop on remote RBridges. The basic
   idea is to use a virtual RBridge (RBv) with a single pseudo-nickname
   to represent an edge group. Any member RBridge of that edge group
   uses this pseudo-nickname rather than its own nickname as the
   ingress nickname when it injects TRILL data frames to TRILL campus.
   The use of the nickname solves the address flip-flop issue by
   setting the MAC address learnt by a remote RBridge to the pseudo-
   nickname. However, it introduces another issue of incorrect packet
   dropping as follows: When a pseudo-nickname is used by an edge
   RBridge as the ingress nickname to forward BUM (Broadcast, Unknown
   unicast and Multicast) traffic, any RBridges (RBn) sitting between
   the ingress RBridge and the distribution tree root will treat the
   traffic as if it was ingressed from the virtual RBridge RBv. If the
   same distribution tree is used by different edge RBridges of the
   same RBv, the traffic may arrive at some RBn from different ports.
   Then the RPF (Reverse Path Forwarding) check required by TRILL

[RFC6325] fails, and the BUM traffic received on unexpected ports will be dropped by RBn.

This document specifies a centralized replication solution for broadcast, unknown unicast and multicast (BUM) traffic forwarding to resolve the issue of incorrect packet drop caused by the RPF check failure in the virtual RBridge case. The basic idea is that all ingress RBridges send BUM traffic to a centralized node, which MUST be a distribution tree root, using unicast TRILL encapsulation. When the centralized node receives the packets, it decapsulates and forwards them to their destination RBridges using a distribution tree established as per the TRILL base protocol.

## [2](#). Conventions used in this document

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 acronyms and terminology in [RFC6325] are used herein with the following additions:

   BUM - Broadcast, Unknown unicast, and Multicast

   CE - As in [RFC7783], Customer Equipment device (end station or bridge). The device can be either physical or virtual equipment.

   Data Label - VLAN or FGL [RFC7172]

   DF - Designated Forwarder [RFC7781]

   FGL - Fine Grained Label [RFC7172].

   LAALP - Local Active-Active Link Protocol [RFC7379].

   MAC flip flop - A problem where the attachment point of a MAC address appears to a remote switch to keep changing. See Section 3.3 of [RFC7379].

   MC-LAG - Multi-Chassis Link Aggregation.

   RPF - Reverse Path Forwarding.

3. **Centralized replication solution overview**

   When an edge RBridge receives BUM traffic from a CE device, it uses
   unicast TRILL encapsulation instead of multicast encapsulation to
   send the packets to a centralized node. The centralized node MUST be
   a distribution tree root. Distribution tree roots are normally
   chosen to be high capacity core RBridges with many high bandwidth
   adjacencies and this constraint makes its practical, as described
   below, to support centralized replication with only software changes
   to transit RBridges.

   The TRILL header of the unicast TRILL encapsulation contains an
   "ingress RBridge nickname" field and an "egress RBridge nickname"
   field [RFC6325]. If the ingress RBridge receives the BUM packet from
   a port that is in an active-active edge group using [RFC7781], it
   sets the ingress RBridge nickname to be the pseudo-nickname rather
   than its own nickname to avoid MAC flip-flop (see Section 3.3 of
   [RFC7379]) on remote RBridges. The egress RBridge nickname is set
   to a special nickname of the centralized node which is used to
   differentiate the centralized replication purpose unicast TRILL
   encapsulation from a normal unicast TRILL encapsulation. This
   special nickname is called an R-nickname.

   When the centralized RBridge receives a unicast TRILL encapsulated
   packet with its R-nickname as egress nickname, it decapsulates the
   packet. Then the centralized RBridge replicates and forwards the
   BUM packet to the packet's destination RBridges using one of the
   distribution trees established as per TRILL base protocol. It MUST
   use a distribution tree whose tree root is the centralized RBridge
   itself. (An RBridge may by the root of more than one tree.) When the
   centralized RBridge forwards the BUM traffic, it simply sends it on
   the distribution tree as if it was a locally ingressed frame except
   that the ingress nickname remains the same as that in the packet it
   received to ensure that the MAC address learning by all egress
   RBridges is bound to the pseudo-nickname.

   When the replicated packet is forwarded by each RBridge along the
   distribution tree starting from the centralized node, an RPF check
   is performed as per [RFC6325]. For any RBridge sitting between the
   ingress RBridge and the centralized replication node, the incoming
   port of such BUM packet should be the centralized node facing port
   as the multicast traffic always comes from the centralized node in
   this solution. However the RPF port as the result of distribution
   tree calculation as specified in [RFC6325] will be the real ingress
   RBridge facing port as it uses the edge group's virtual RBridge as
   the ingress RBridge, so the RPF check will fail.

To solve this problem, some change in the RPF test is required. In this case, the RPF calculation on each RBridge should use the centralized node as the ingress RBridge for each tree for which that node is the root instead of the real ingress virtual RBridge to perform the calculation. As a result, RPF check will accept traffic on the centralized node facing port of the RBridge for multi-destination traffic. This prevents incorrect frame drops by the RPF check.

The change in the actual RPF check on a received multi-destination TRILL data packet is easy. The [RFC6325] RPF check is a check to see if a triple of {ingress nickname, tree, receiving RBridge port} is allowed. (The tree is indicated by the nickname of its root which is stored in the TRILL Header "egress nickname" field.)  When determining the RFC check, if "ingress nickname" is using centralized replication (indicated by a C-nickname, see Section 9), then the check is based on distribution from the tree root. If "ingress nickname" is not using centralized replication, then the check is based on distribution from the RBridge having the ingress nickname.

To differentiate the centralized replication unicast TRILL encapsulation from normal unicast TRILL encapsulation, the R-nickname is introduced for centralized replication. When the centralized node receives unicast TRILL encapsulation traffic with the egress nickname R-nickname, it decapsulates the packet and then forwards the packet to the destination RBridges through a distribution tree for which it is the root by re-encapsulation as aforementioned. In TRILL, RBridges can hold multiple nicknames so the centralized RBridge simply obtains another nickname to use as the R-nickname. The centralized RBridge or RBridges should announce their R-nickname to all TRILL campus through the TRILL LSP extension specified in Section 11.


4. **Frame duplication from remote RBridge**

Frame duplication may occur when a remote host sends a multi-destination frame to a local CE which has an active-active connection to the TRILL campus. To avoid local CE receiving multiple copies from a remote RBridge, the designated forwarder (DF) mechanism is supported for egress direction multicast traffic.

The DF election mechanism [RFC7781] allows only one port of one RBridge in an active-active group to forward multicast traffic from the TRILL campus to the local access side for each VLAN. The basic idea of DF is to elect one RBridge per VLAN from an edge group to be

responsible for egressing the BUM traffic. [RFC7781] describes the
DF election mechanism among member RBridges involving in an edge
group.

If the DF election mechanism is used for frame duplication
prevention, access ports on an RBridge are categorized as one of
three types: non-group, group DF port and group non-DF port. The
last two types can be called group ports. Each of the group ports
is associated with a pseudo-nickname. If consistent nickname
allocation to edge group RBridges is used, it is possible that same
pseudo-nickname is associated with more than one port on a single
RBridge. A typical scenario is that CE1 is connected to RB1 & RB2
by LAALP1 while CE2 is connected to RB1 & RB2 by LAALP2.
In order to conserve the number of pseudo-nicknames used, member
ports for both LAALP1 and LAALP2 on RB1 & RB2 are all associated
with the same pseudo-nickname.

## 5. Local forwarding behavior on ingress RBridge

When an ingress RBridge (RB1) receives BUM traffic from a local
active-active connected CE (CE1) device, the traffic will be
injected into the TRILL campus with TRILL encapsulation, and it will
be replicated and forwarded to all destination RBridges through
central replication, including the ingress RBridge itself, along a
TRILL distribution tree. To avoid the traffic looping back to the
original sender CE, an ingress nickname of the CE group's pseudo-
nickname is used for traffic filtering.

However, if there are two CEs, say CE1 and CE2, connecting to the
ingress RB1 and each associated with the same pseudo-nickname, RB1
needs to locally replicate and forward to CE2, because another copy
of the BUM traffic between CE1 and CE2 through TRILL campus will be
blocked by the traffic filtering.

If CE1 and CE2 are not associated with the same pseudo-nickname, the
copy of the BUM traffic between CE1 and CE2 through TRILL campus
won't be blocked by the traffic filtering. To avoid duplicated
traffic on receiver CE, there cannot be local replicated BUM traffic
between these two CEs on ingress RB1.

In summary, to ensure correct BUM traffic forwarding behavior for
each CE, the local replication behavior on the ingress RBridge is as
follows:

   1. Replicate to the active-active group ports associated with the
   same pseudo-nickname as that associated with the incoming port.

   2. Do not replicate to active-active group ports associated with
   other pseudo-nicknames.

   3. Do not replicate to non-edge-group ports.

   The above local forwarding behavior on the ingress RBridge of RB1
   can be called centralized replication local forwarding behavior A.

   If ingress RBridge RB1 itself is the centralized replication node,
   BUM traffic injected by RB1 into the TRILL campus won't loop back to
   RB1. In this case, the local forwarding behavior is called
   centralized replication local forwarding behavior B. Behavior B on
   RB1 is as follows:

   1. Local replication to the ports associated with the same
   pseudo-nickname as that associated with the incoming port.

   2. Local replication to the group DF port associated with
   different pseudo-nicknames. Do not replicate to group non-DF port
   associated with different pseudo-nicknames.

   3. Local replication to non-edge-group ports.

## [6]. Loop prevention among RBridges in an edge group

   If a CE sends a broadcast, unknown unicast, or multicast (BUM)
   packet through a DF port to an ingress RBridge, that RBridge will
   forward that packet to all or a subset of the other RBridges that
   only have non-DF ports for that active-active group. Because BUM
   traffic forwarding to non-DF ports isn't allowed, in this case the
   frame won't loop back to the CE.

   If a CE sends a BUM packet through a non-DF port to an ingress
   RBridge, say RB1, then RB1 will forward that packet to other
   RBridges that have a DF port for that active-active group. In this
   case the frame will loop back to the CE and the traffic split-
   horizon filtering mechanism is used to avoid looping back among
   RBridges in the edge group.

   This split-horizon mechanism relies on the ingress nickname field in
   the TRILL header to check if a packet's egress port belongs to a
   same active-active group as the packet's incoming port to the TRILL
   campus.

When the ingress RBridge receives BUM traffic from an active-active
connected CE device, the traffic will be sent through the TRILL
campus with TRILL encapsulation to a centralized RBridge. There it
will be replicated and forwarded to its destination RBridges, which
include ingress RBridge itself, through a TRILL distribution tree.
If the same pseudo-nickname is used for two active-active access CEs
as ingress nickname, an egress RBridge can use that nickname to
filter traffic forwarding to all local CEs. In this case, the
traffic between these two CEs goes through the local RBridge and
another copy of the traffic from the TRILL campus is filtered. If
different ingress nicknames are used for two connecting CE devices,
the access ports connecting to these two CEs should be isolated from
each other. The BUM traffic between these two CEs should go through
the TRILL campus, otherwise the destination CE connected to same
RBridge with the sender CE will receive two copies of the traffic.

## 7. Centralized replication forwarding process

```
                           +-----------+
                           |   (RB5)   |
                           +-----------+
                                |
                           +-----------+
                           |   (RB4)   |
                           +-----------+

                      |      |      |
                   --------      |     --------
                   |             |            |
                +------+     +------+     +------+
                |(RB1) |     |(RB2) |     | (RB3)|
                +------+     +------+     +------+
                  *    |       *   |        * |  ^
                  *    |       *   |        * |   ^
                  *   ----------*-------------*--    ^
                 *************************** |     ^
             LAALP1 *                    LAALP2 |      ^
               +------+                   +------+    +------+
               | CE1 |                    | CE2  |    | CE3  |
               +------+                   +------+    +------+
           Figure 1  TRILL Active-active access
```
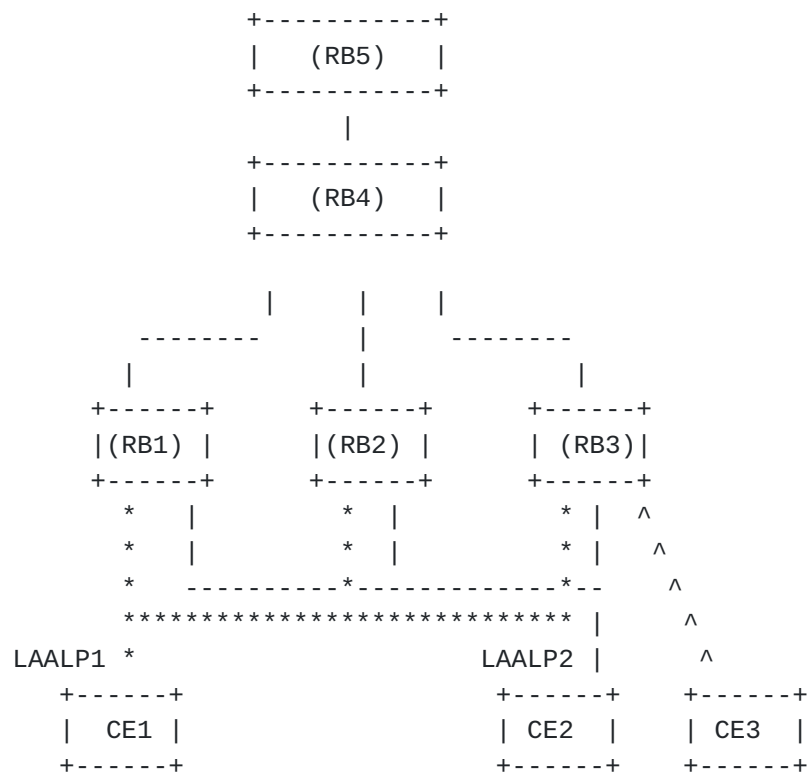
Assuming the centralized replication solution is used in the example
network of above figure 1, RB5 is the distribution tree root and

centralized replication node, CE1 and CE2 are active-active accessed
to RB1, RB2, and RB3 through LAALP1 and LAALP2 respectively, CE3 is
single homed to RB3. The RBridge's own nicknames of RB1 to RB5 are
nick1 to nick5 respectively. RB1, RB2, and RB3 use the same pseudo-
nickname for LAALP1 and LAALP2; that pseudo-nickname is P-nick. The
R-nickname on the centralized replication node of RB5 is S-nick.

The BUM traffic forwarding process from CE1 to CE2 and CE3 is as
follows:


    1. CE1 sends BUM traffic to RB3.

    2. RB3 replicates and sends the BUM traffic to CE2 locally. RB2
also sends the traffic to RB5 using unicast TRILL encapsulation. In
the TRILL Header, the ingress nickname is set as P-nick and the
egress nickname is set as S-nick.

    3. RB5 decapsulates the unicast TRILL Data packet. Then it uses a
distribution tree for which it is the root to forward the packet as
a multi-destination TRILL Data packet. The egress nickname in the
multi-destination TRILL Header is the nick5 and the ingress nickname
is still P-nick. If RB3 had sent the unicast to some nickname that
was not an R-nickname, the packet will not be re-encapsulate. If it
is sent to an R-nickname that is not a tree root, it will either be
not forwarded at all or, if it is re-encapsulated and forwarded,
will be subject to incorrect pruning and will not be delivered to
all of its intended recipients.

    4. RB4 receives multicast TRILL traffic from RB5. Traffic
incoming port is the up port facing the distribution tree root,
RB4's RPF check will be correct based on the changed RPF port
calculation algorithm in this document. After the RPF check is
performed, it forwards the traffic to all other egress RBridges (RB1,
RB2, and RB3).

    5. RB3 receives multicast TRILL traffic from RB4. It decapsulates
the multi-destination TRILL Data packet. Because the ingress
nickname of P-nick is equivalent to the nickname of local LAALPs
connecting to CE1 and CE2, RB3 doesn't forward the traffic to CE1
and CE2 to avoid duplicated frame. RB3 only forwards the packet to
CE3.

    6. RB1 and RB2 receive multicast TRILL traffic from RB4. The
forwarding process is similar to the process on RB3, i.e., because
the ingress nickname of P-nick is equivalent to the nickname of the

local LAALPs connecting CE1 and CE2, they also don't forward the traffic to local CE1 and CE2.

**8. BUM traffic load balancing among multiple centralized nodes**

To support unicast TRILL encapsulation BUM traffic load balancing, multiple centralized replication nodes can be deployed and the traffic can be spread over these nodes based on data label (VLAN or FGL). Furthermore, if it was desirable for a centralized node to be sent more of this BUM traffic, it could hold two or more R-nicknames. The share of BUM traffic it would receive would be proportional to the number of R-nicknames it held.

Assuming there are k different R-nicknames held by centralized nodes in a TRILL campus. The VLAN-based (or FGL-based [RFC7172]) load balancing algorithm used by ingress active-active access RBridge is as follows:

    1. All R-nicknames are ordered and numbered from 0 to k-1 in ascending order treating the nicknames as unsigned 16-bit integers.

    2. For data label ID m, choose the R-nickname whose number equals (m mod k) as egress nickname for BUM traffic unicast TRILL encapsulation.

For examples, there are 3 R-Nicknames (RN). The RNs will be ordered RN0 to RN2. Assuming there are 5 VLANs from VLAN ID 1 to ID 5 spreading among edge RBridges, the traffic in VLAN 1 will go to RN1, VLAN 2 will go to RN2, and so on.

When an ingress RBridge participating in active-active connection receives BUM traffic from local CE, the RBridge decides which R-nickname to send the traffic to based on the VLAN-based load spreading algorithm, thus data-label-based load balancing for the BUM traffic can be achieved using multiple centralized nodes/multiple R-nicknames.
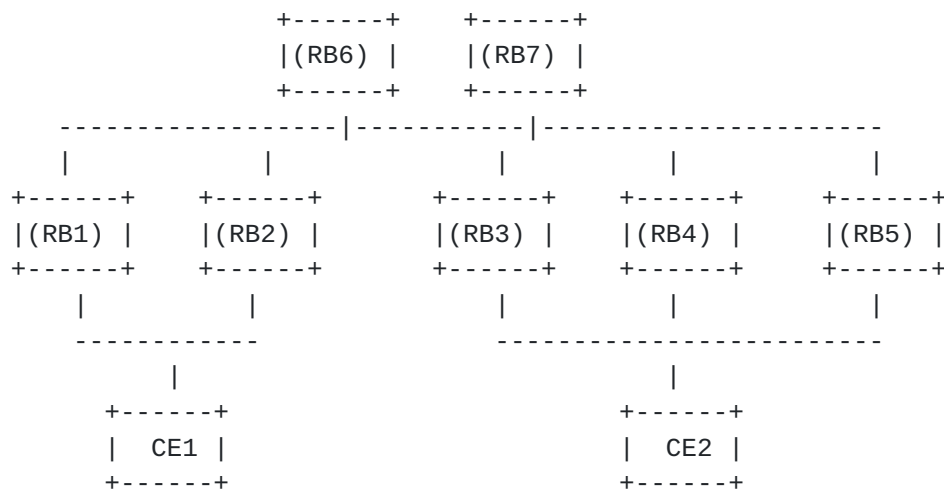
## 9. Co-existing with the CMT (RFC 7783) solution

```
                +------+     +------+
                |(RB6) |     |(RB7) |
                +------+     +------+
     -----------------|-----------|----------------------
       |         |           |          |            |
   +------+   +------+     +------+    +------+     +------+
   |(RB1) |   |(RB2) |     |(RB3) |    |(RB4) |     |(RB5) |
   +------+   +------+     +------+    +------+     +------+
      |          |           |           |            |
      ------------           ------------------------
           |                             |
        +------+                      +------+
        | CE1 |                       | CE2 |
        +------+                      +------+
```

         Figure 2 CMT and centralized replication co-existing scenario


   Both the centralized replication solution and the Coordinated
   Multicast Trees (CMT) [RFC7783] solution rely on using
   pseudo-nicknames to avoid MAC flip-flop on remote RBridges.
   These two solutions can co-exist in a single TRILL campus.
   Each solution can be selected by each active-active edge
   group of RBridges independently.

   As illustrated in Figure 2, RB1 and RB2 use CMT for CE1's active-
   active access, RB3, RB4, and RB5 use the centralized replication for
   CE2's active-active access.

   For the centralized replication solution, edge group RBridges MUST
   announce the local pseudo-nickname using Nickname Flags APPsub-TLV
   with C-flag set. A nickname with the C-flag set is called a "C-
   nickname". A transit RBridge will perform the centralized
   replication specific RPF check algorithm if it receives TRILL Data
   packets with a C-nickname as ingress nickname.

   An edge group using CMT [RFC7783] MUST NOT set the C-nickname flag
   on the pseudo-nickname it is using. This is already mandatory
   behavior because any RBridge originating a Nickname Flags APPsub-TLV
   is required by [RFC7780] to set any flag bit it does not know about
   to zero. If a edge RBridge using CMT [RFC7783] nevertheless set the
   C-bit for an edge group pseudo-nickname, it is very likely that BUM
   traffic encapsulated with that nickname as ingress would be
   incorrectly pruned early in its distribution and would thus reach
   few (possibly none) of its intended targets. To avoid confusion, a

pseudo-nickname MUST NOT be shared between a centralized replication
edge group and a CMT-based edge group.

## [10](#). Network upgrade analysis

Centralized nodes will typically need software and hardware upgrades
to support centralized replication, which stitches together the
TRILL unicast traffic decapsulation process and the process of
normal TRILL multicast traffic forwarding along distribution tree.

Active-active connection edge RBridges will typically need software
and hardware upgrade to support unicast TRILL encapsulation for BUM
traffic; the process is similar to other head-end replication
processes.

Transit nodes typically need only a software upgrade to support the
changed RPF port calculation algorithm.

## [11](#). TRILL protocol extensions

Two new flags, "R" and "C", are specified in the Nickname Flags
APPsub-TLV [[RFC7780](#)]. A nickname with the "R" flag set is called an
R-nickname and a nickname with the the "C" flag set is called a C-
nickname. The R-nickname is a specialized nickname attached to a
centralized node to differentiate unicast TRILL encapsulated BUM
traffic from normal unicast TRILL traffic. The C-nickname flag is
set on the pseudo-nickname for each edge group that uses the
centralized replication. A C-nickname is a specialized pseudo-
nickname for which transit RBridges perform a different RPF check
algorithm for TRILL data packets with the C-nickname in the ingress
nickname field.

When active-active edge RBridges use centralized replication to
forward BUM traffic, the R-nickname is used as the egress nickname
and the C-nickname is used as ingress nickname in the TRILL header
for the unicast TRILL encapsulation of BUM traffic.

## [11.1](#). "R" and "C" Flag in the Nickname Flags APPsub-TLV

If this APPsub-TLV is being advertised by an RBridge that does not
have the nickname appearing in the Nickname Flags APPsub-TLV, the R
and C flag bits in the APPsub-TLV MUST be treated as if they were
zero. If an RBridge that is not a distribution tree root advertises
an R-nickname, that nickname MUST NOT be treated as an R-nickname
but rather as an ordinary nickname, that is, the R-nickname flag is

ignored for all purposes if the nickname is held by an RBridge that is not a tree root.

```
      0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
     +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
     |    Nickname                                   |
     +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
     |IN|SE|R | C|    RESV                           |
     +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
                    NICKFLAG RECORD
```

o R = If R flag is one, it indicates that the advertising TRILL switch holding Nickname is a centralized replication node, and Nickname is used as egress nickname for edge group RBridges to inject BUM traffic into the TRILL campus when the edge group RBridges use this centralized replication solution for active-active access. If flag is zero, that nickname will not be used for that purpose.

o C = If C flag is one, it indicates that the TRILL traffic with this nickname as an ingress nickname that requires the special RPF check algorithm specified in Section 3. If flag is zero, that nickname will not be used for that purpose.

It is possible, due to errors or due to transient inconsistent LSPs when the link state database is converging after a configuration change or the like for there to be inconsistent Nickname Flags APPsub-TLVs for the same nickname. In this case it is RECOMMENDED that the nickname be treated as if the R / C flag was set if any Nickname Flags APPsub-TLV for that nickname has the R / C flag set.

## 12. Security Considerations

This draft does not introduce any extra security risks. A rogue RBridge that is a tree root can attract traffic by advertising an R-nickname. However, this does not represent a substantial increase in risk as RBridges could cause problems in a number of other ways by advertising low cost adjacencies or making themselves the highest priority tree root or the like. In general, the protection against an untrusted device acting as an RBridge and wrecking havoc is to use IS-IS authentication [RFC5310] and configure and administer the TRILL campus so that only trusted RBridges have the authentication key.

   For general TRILL Security Considerations, see [RFC6325]. For
   Security Considerations related to pseudo-nickname active-active,
   see [RFC7781].

## 13. IANA Considerations

   IANA is requested to assign two bits in the Nickname Flags APPsubTLV
   flags for the R and C bits discussed in Section 11.1 [Bits 3 and 4
   suggested] and update the "NickFlags" Bits registry on the TRILL
   Parameters page as follows:

| Bit | Mnemonic | Description | Reference |
| --- | -------- | ------------------- | ----------- |
| 3 | R | Replication Nickname | [This document] |
| 4 | C | Special RFC Check | [This document] |

## 14. References

## 14.1. 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, <http://www.rfc-editor.org/info/rfc2119>.

   [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>.

   [RFC6165]  Banerjee, A. and D. Ward, "Extensions to IS-IS for
   Layer-2 Systems", RFC 6165, DOI 10.17487/RFC6165, April 2011,
   <http://www.rfc-editor.org/info/rfc6165>.

   [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>.

   [RFC7172] Eastlake 3rd, D., Zhang, M., Agarwal, P., Perlman, R.,
   and D. Dutt, "Transparent Interconnection of Lots of Links (TRILL):
   Fine-Grained Labeling", RFC 7172, DOI 10.17487/RFC7172, May 2014,
   <http://www.rfc-editor.org/info/rfc7172>.

   [RFC7176] Eastlake 3rd, D., Senevirathne, T., Ghanwani, A., Dutt,
   D., and A. Banerjee, "Transparent Interconnection of Lots of Links

(TRILL) Use of IS-IS", RFC 7176, DOI 10.17487/RFC7176, May 2014,
<http://www.rfc-editor.org/info/rfc7176>.

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

[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>.

## 14.2. Informative References

[RFC7781] Zhai, H., Senevirathne, T., Perlman, R., Zhang, M., and
Y. Li, "Transparent Interconnection of Lots of Links (TRILL):
Pseudo-Nickname for Active-Active Access", RFC 7781, DOI
10.17487/RFC7781, February 2016, <http://www.rfc-
editor.org/info/rfc7781>.

[RFC7379] Li, Y., Hao, W., Perlman, R., Hudson, J., and H. Zhai,
"Problem Statement and Goals for Active-Active Connection at the
Transparent Interconnection of Lots of Links (TRILL) Edge", RFC 7379,
DOI 10.17487/RFC7379, October 2014, <http://www.rfc-
editor.org/info/rfc7379>.

[RFC7783] Senevirathne, T., Pathangi, J., and J. Hudson,
"Coordinated Multicast Trees (CMT) for Transparent Interconnection
of Lots of Links (TRILL)", RFC 7783, DOI 10.17487/RFC7783, February
2016, <http://www.rfc-editor.org/info/rfc7783>.

## 15. Acknowledgments

The authors wish to acknowledge the important contributions of
Donald Eastlake, Hongjun Zhai, Xiaomin Wu, Liang Xia, Francis Dupont.

Authors' Addresses

   Weiguo Hao
   Huawei Technologies
   101 Software Avenue,
   Nanjing 210012
   PR China
   Email: haoweiguo@huawei.com

   Yizhou Li
   Huawei Technologies
   101 Software Avenue,
   Nanjing 210012
   PR China
   Email: liyizhou@huawei.com

   Muhammad Durrani
   Equinix
   Email: mdurrani@equinix.com

   Sujay Gupta
   IP Infusion
   RMZ Centennial
   Mahadevapura Post
   Bangalore - 560048
   India
   Email: sujay.gupta@ipinfusion.com

   Andrew Qu
   MediaTec
   Email: laodulaodu@gmail.com