

TRILL Working Group  
INTERNET-DRAFT  
Intended status: Proposed Standard

Donald Eastlake  
Mingui Zhang  
Huawei  
Puneet Agarwal  
Broadcom  
Dinesh Dutt  
Cisco  
Radia Perlman  
Intel Labs  
July 11, 2011

Expires: January 10, 2012

**RBridges: Fine-Grained Labeling**  
<[draft-eastlake-trill-rbridge-fine-labeling-01.txt](#)>

## Abstract

The IETF has standardized RBridges (Routing Bridges), devices that implement the TRILL (TRansparent Interconnection of Lots of Links) protocol, a solution for least cost transparent frame routing in multi-hop networks with arbitrary topologies, using link-state routing and encapsulation with a hop count.

The TRILL base protocol standard supports up to 4K VLAN IDs (Virtual Local Area Network IDentifiers). However, there are applications that require more fine-grained labeling of data and end stations. This document specifies extensions to the TRILL protocol to accomplish this.

## Status of This Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Distribution of this document is unlimited. Comments should be sent to the TRILL working group mailing list.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at  
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at

<http://www.ietf.org/shadow.html>

D. Eastlake, et al

[Page 1]

## Table of Contents

<a href="#">1. Introduction.....</a>	<a href="#">3</a>
<a href="#">1.1 Terminology.....</a>	<a href="#">3</a>
<a href="#">2. Fine-Grained Labeling.....</a>	<a href="#">4</a>
<a href="#">2.1 Requirements.....</a>	<a href="#">4</a>
<a href="#">2.2 Existing TRILL VLAN Labeling.....</a>	<a href="#">5</a>
<a href="#">2.3 Fine-Grained Labeling.....</a>	<a href="#">6</a>
<a href="#">3. Coexistence with ST RBridges.....</a>	<a href="#">8</a>
<a href="#">4. Processing Finely Labeled Frames.....</a>	<a href="#">9</a>
<a href="#">4.1 Ingress Processing.....</a>	<a href="#">9</a>
<a href="#">4.2 Transit Processing.....</a>	<a href="#">10</a>
<a href="#">4.2.1 Unicast Transit Processing.....</a>	<a href="#">10</a>
<a href="#">4.2.2 Multi-Destination Transit Processing.....</a>	<a href="#">10</a>
<a href="#">4.3 Egress Processing.....</a>	<a href="#">11</a>
<a href="#">4.4 Address Learning.....</a>	<a href="#">12</a>
<a href="#">5. IS-IS Extensions.....</a>	<a href="#">13</a>
<a href="#">5.1 Announcing RBridge DT Support.....</a>	<a href="#">13</a>
<a href="#">5.2 Interested Labels and Bridge Roots sub-TLV.....</a>	<a href="#">13</a>
<a href="#">5.3 The Group Labeled MAC Address sub-TLV.....</a>	<a href="#">14</a>
<a href="#">6. IANA Considerations.....</a>	<a href="#">16</a>
<a href="#">7. Security Considerations.....</a>	<a href="#">16</a>
<a href="#">7.1 Ingress Forgery and Egress Compromise.....</a>	<a href="#">16</a>
<a href="#">8. References.....</a>	<a href="#">17</a>
<a href="#">8.1 Normative References.....</a>	<a href="#">17</a>
<a href="#">8.2 Informative References.....</a>	<a href="#">17</a>



## **1. Introduction**

The IETF has standardized RBridges (Routing Bridges), devices that implement the TRILL (TRansparent Interconnection of Lots of Links) protocol [[RFCtrill](#)], a solution for least cost transparent frame routing in multi-hop networks with arbitrary topologies, using link-state routing and encapsulation with a hop count.

The TRILL base protocol standard supports up to 4,094 VLAN IDs (Virtual Local Area Network IDentifiers). However, there are applications that require more fine-grained labeling of data and end stations. This document specifies extensions to the TRILL protocol to accomplish this.

Familiarity with [[RFCtrill](#)] and [[ISIStrill](#)] is assumed in this document.

### **1.1 Terminology**

The terminology and acronyms of [[RFCtrill](#)] are used in this document with the additions listed below.

DT - Double Tagging or Double Tagged or Double Tag

Edge RBridge - An RBridge announcing VLAN or fine-grained label connectivity in its LSP

ST - Single Tagging or Single Tagged or Single Tag

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



## 2. Fine-Grained Labeling

The essence of fine-grained labeling is that (a) when TRILL Data frames are ingressed or created they may incorporate a label from a set of significantly more than 4K labels, (b) RBridge ports can be labeled with a set of such labels, and (c) a TRILL Data frame cannot be egressed through such an RBridge port unless its label matches one of the labels of the port.

[Section 2.1](#) lists fine-grained labeling requirements. [Section 2.2](#) briefly outlines VLAN labeling in the TRILL base protocol standard [[RFCtrill](#)]. And [Section 2.3](#) then outlines a method of fine-grained labeling of TRILL Data frames.

In the remainder of this document, we commonly refer to the simple VLAN labeling provided by the TRILL base protocol standard as single tagging (ST) or coarse labeling and refer to fine-grained labeling as double tagging (DT).

### 2.1 Requirements

There are several requirements that should be met by fine-grained labeling in TRILL. They are briefly described in the list below in approximate order by priority with the most important first.

#### 1. Fine-Grained

Some networks have a large number of entities that need configurable isolation, whether those entities are independent customers, applications, or branches of a single endeavor or some combination of these or other entities. The VLAN tags supported by [[RFCtrill](#)] provides for only  $(2^{12} - 2)$  valid VLAN identifiers. A substantially larger number is required.

#### 2. Silicon Considerations

Fine-grained labeling should, to the extent practical, use existing features, processing, and fields that are already supported in at least some of the many existing TRILL fast path silicon implementations.

#### 3. Base RBridge Compatibility

To support some incremental conversion scenarios, it is desirable that not all RBridges in a campus using fine-grained labeling be required to be fine-grained label aware. That is, it is desirable that RBridges not implementing the fine-grained labeling feature

and performing at least the transit forwarding function can

usefully process TRILL Data frames that incorporate fine-grained labeling.

#### 4. Alternate Priority

It would be desirable for an ingress RBridge to be able to assign a different priority to a fine grain labeled TRILL Data frame for its ingress-to-egress propagation from the priority of the original native frame. The original priority should be restored on egress.

## 2.2 Existing TRILL VLAN Labeling

This section provides a brief review of existing TRILL Data frame coarse VLAN labeling.

Currently TRILL Data frames have the single tagged (ST) structure shown below:

```

+-----+
| Link Header |
+-----+
| TRILL Header |
+-----+
| Inner.MacDA |
+-----+
| Inner.MacSA |
+-----+
| Inner.VLAN   | <-- Coarse VLAN Label
+-----+
| Payload      |
+-----+
| Link Trailer |
+-----+

```

The Inner.VLAN tag is always present and is specified as a C-tag [802.1Q] providing (  $2^{12} - 2$  ) labels (the values 0 and 0xFFF are reserved) that is structured as follows:

```

  0 1 2 3 4 5 6 7 8 9 A B C D E F
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Ethertype 0x8100           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| PRI |C|           VLAN ID           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The PRI field above is the 3-bit unsigned priority field where larger

numbers represent higher priority except that the default zero

priority is above priority 1 and below priority 2 [[802.1Q](#)]. Under the TRILL base protocol [[RFCtrill](#)], in the Inner.VLAN the C bit is required to be set to zero, transparently forwarded, and ignored on receipt by RBridges.

For an RBridge conformant to the TRILL base protocol, incoming frames are classified as to their VLAN ID and priority by the port on which they are received as described in [Appendix D](#) of [[RFCtrill](#)].

### 2.3 Fine-Grained Labeling

In the proposed form, fine-grained labeling expands the 12-bit coarse VLAN label available under the TRILL base protocol standard to a 24-bit label. In this document, fine-grained labels are sometimes denoted as "(X.Y)" where X is the high order 12 bits and Y is the low order 12 bits. The fine grained label information appears in the same location in a TRILL Data frame as the coarse VLAN label did, as shown below, although it is encoded as two consecutive VLAN tags (DT).

```
+-----+
| Link Header |
+-----+
| TRILL Header |
+-----+
| Inner.MacDA |
+-----+
| Inner.MacSA |
+-----+
| Inner.Label | <-- Fine-Grained Label
+-----+
| Payload      |
+-----+
| Link Trailer |
+-----+
```

The fine-grained label is encoded as two sequential C-tags as shown below. The high order 12 bits of the fine-grained label appear in the VLAN ID field of the first C-tag and the low order 12 bits appear in the VLAN ID field of the second. Because some silicon might subject the high order part of the fine-grained label to the same constraints as VLAN IDs and for other reasons such as the reporting described in [Section 4.2.2](#), the values zero and 0xFFF are reserved for the high order part of a TRILL fine-grained label. [[[ Should 0 and 0xFFF be prohibited in the low order 12 bits also? ]]]

[[[ Alternative Ethertype sequences could be specified. Perhaps the most obvious alternative would be for the first VLAN tag to be as S-

tag (Ethertype 0x88A8) and the second a C-tag. However, this might

cause problems for some ST RBridges; if they check the Ethertype of the first VLAN tag, they might reject such frames. ]]]

```

      0 1 2 3 4 5 6 7 8 9 A B C D E F
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           Ethertype 0x8100           |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| PRI |C| High Order Label Bits |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           Ethertype 0x8100           |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| PRI |C| Low Order Label Bits  |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

The appropriate DT for an ingress native frame is determined by the input RBridge port as specified in [Section 4.1](#). The priority in the second tag is that associated by the ingress port with the native frame as with ST ingress. The priority in the first tag is either a copy of the second tag priority or that priority mapped at ingress, depending on the capabilities of the ingress RBridge. Ports of RBridges supporting DT also have capabilities to transmit frames being forwarded or egressed as untagged or C-tagged as specified in [Section 4.3](#).

Use of S-tags or tags stacked beyond that indicated are beyond the scope of this document but are an obvious extension.



### 3. Coexistence with ST RBridges

ST (single tag) RBridges will operate properly as transit RBridges. Transit RBridges look at the Inner.VLAN ID only for the filtering of multi-destination frames. If an RBridge does not perform filtering, or filters on only some of the fields in the packet, the only consequence is that multi-destination frames will use more bandwidth than necessary. Because ST RBridges could only look at the initial VLAN tag in the fine-grained label of a DT (double tag) multi-destination frame, they will not be able to prune as effectively as transit DT RBridges could.

It would be more serious if an ST edge RBridge, RB1, unaware of the double tag, forwarded a DT frame with DT label (X.Y) onto a link configured as ST VLAN-X, with RB1 stripping the "X" and forwarding the packet. This violates the separation of VLANs, and might cause other problems on a link in which the VLAN tag should have been stripped. It would also be problematic if a malicious end station could forge an apparent DT label (X.Y) frame by including extra tags in native frames ingressed by an ST edge RBridge. Therefore, it is highly desirable for all the edge RBridges to be DT RBridges.

DT RBridges will report the DT capability in LSPs, so DT RBridges (and any management system with access to the link state database) will be able to detect the existence of ST edge RBridges.

It might be useful, in a particular campus with mixed DT and ST RBridges, to have some end station VLANs accessible via ST edge RBridges. This is supported by reserving some number of VLANs (say the first k), to be ST-addressable. These VLANs will be specified with a single Inner.VLAN tag, whether or not the edge RBridges attached to these VLANs are DT-capable. When ST-specifiable VLANs are used in a DT campus, and where there are ST edge RBridges advertising connectivity to those VLANs, the first VLAN tag in a double tag MUST NOT be equal to the value of any ST-specifiable VLAN.

If this rule is violated, the network misconfiguration is detected by the DT RBridges that will then refuse ingress to or egress from label (X.Y) while VLAN X connectivity is being advertised by an ST edge RBridge.



## **4. Processing Finely Labeled Frames**

This section specifies ingress, transit, and egress processing of TRILL Data frames with regard to fine-grained labels, also known as double tagging (DT). A transit or egress DT RBridge detects DT TRILL Data frames by noticing that the Ethertype immediately after the first Inner.Label VLAN tag is the C-tag Ethertype.

### **4.1 Ingress Processing**

There is no change in Appointed Forwarder logic [[RFCaf](#)] for the ports of a DT RBridge.

A DT RBridge may be configured, on one or more ports, to double tag ingressed native frames. There is no change in ST ingress processing, which is the default unless a port has been configured for DT.

DT RBridges MUST remove any extra C-tags from incoming native frames being ingressed, regardless of whether the ingress port is configured as ST or DT (see [Section 7.1](#)).

DT RBridges MUST support configurable per port mapping from the C-VLAN ID associated with a native frame to a 24-bit fine-grained label. DT RBridges MAY support other methods to determine the DT ID of an incoming native frame. If the resulting label (X.Y) is such that VLAN X connectivity is being advertised by an ST edge RBridge in the campus, the ingressed frame MUST be dropped.

The DT ingress process MUST place the priority associated with an ingressed native frame in the second Inner.Label C-tag. It SHOULD also associate a possibly different mapped priority with an ingressed frame. The mapped priority is placed in the initial Inner.Label C-tag. If such mapping is not supported then the original priority is also placed in the initial inner C-tag.

A DT ingress RBridge MAY serially unicast a multi-destination DT frame to the relevant egress RBridge or RBridges after encapsulating it as a TRILL known unicast data frame. The relevant egress RBridges are determined by starting with those announcing connectivity to the frame's (X.Y) label. That set SHOULD be further filtered based on multicast listener and router connectivity if the native frame was a multicast frame.



## **4.2 Transit Processing**

TRILL Data frame transit processing is fairly straightforward as described in [Section 4.2.1](#) for known unicast TRILL Data frames and in [Section 4.2.2](#) for multi-destination TRILL Data frames.

### **4.2.1 Unicast Transit Processing**

There is almost no change in TRILL unicast transit processing. A transit RBridge forwards any TRILL unicast data frame to the next hop towards the egress RBridge as specified in the TRILL Header. Just as RBridges conformant to the TRILL base protocol standard [[RFCtrill](#)] do not examine the Inner.VLAN ID of ST transit known unicast TRILL Data frames, DT RBridges do not examine either the high or low order part of the 24-bit ID in the Inner.Label for transit DT known unicast TRILL Data frames.

However, as provided in the TRILL base protocol standard [[RFCtrill](#)], all transit RBridges, whether ST or DT, MUST take the priority used for a forwarded frame from the Inner.VLAN tag, which will be the first of the two DT VLAN tags for a DT TRILL Data frame.

### **4.2.2 Multi-Destination Transit Processing**

All multi-destination TRILL Data frames are forwarded on a distribution tree selected by the ingress RBridge. The distribution trees for DT multi-destination frames are the same trees as for ST multi-destination frames, calculated as provided for in the TRILL base protocol standard [[RFCtrill](#)]. There is no change in the Reverse Path Forwarding Check.

A DT RBridge, say RB1, having a DT multi-destination frame for label (X.Y) to forward, SHOULD prune as in the base specification, based on whether there are any edge RBridges on the tree branch that are connected to label (X.Y). In addition, RB1 SHOULD prune multicast frames based on reported multicast listener and multicast router attachment in (X.Y). Finally, a transit DT RBridge MAY drop any multi-destination frame for label (X.Y) if some DT RBridge is advertising connectivity to VLAN X. "MAY" is chosen in this case to minimize the mandatory burden on transit RBridges.

To ensure that a transit ST RBridge does not falsely filter traffic for DT label (X.Y), a DT edge RBridge attached to DT label (X.Y) MUST report connection to VLAN X, as if X were a ST VLAN, in addition to reporting connectivity to label (X.Y). Because of this, DT transit

RBridges can safely apply pruning to all TRILL Data frames, both ST

and DT, based on the first Inner.VLAN ID and the reported VLAN-X connectivity of all downstream RBridges.

To ensure that a transit ST RBridge does not falsely prune traffic for DT label (X.Y) based on multicast filtering, a DT edge RBridge attached to label (X.Y) MUST report for VLAN X either (1) that it is attached to both IPv4 and IPv6 multicast routers or (2) its actual DT label (X.Y) multicast listener and router connectivity situation.

### **4.3 Egress Processing**

Egress processing is generally the reverse of ingress processing described in [Section 4.1](#).

If any ST RBridge in the campus is announcing connectivity to VLAN-X, a DT RBridge MUST NOT egress a frame with DT label (X.Y) but must drop such a frame.

A DT RBridge MUST be able to configurably convert the 24-bit fine grained label in a DT TRILL Data frame it is egressing to a 12-bit C-VLAN ID for the resulting native frame on a per port basis. A port MAY be configured to strip such tagging. It is the responsibility of the network manager to properly configure the DT RBridges and ports in the campus to obtain the desired mappings.

A DT RBridge egresses DT frames with the above tag conversion similarly to the egressing of ST frames, as follows:

1. A known unicast DT frame is egressed to the DT port matching its fine-grained label and Inner.MacDA. Or, if there is no such port, it is flooded out all DT ports with its fine-grained label unless the RBridge has knowledge that the frames Inner.MacDA cannot be out that port.
2. A multi-destination DT frame is decapsulated and flooded out all ports with its fine-grained label subject to multicast pruning.

DT RBridges MUST accept multi-destination encapsulated frames that are sent to them as TRILL unicast frames (TRILL Header M bit = 0). They locally egress such frames, if appropriate, and MUST NOT forward them (other than egressing them as native frames on their local links).



#### **4.4 Address Learning**

A DT RBridge learns addresses on DT ports based on the fine-grained label rather than VLAN ID. Addresses learned from ingressed native frames are logically represented by { MAC address, fine-grained label, port, confidence, timer } while remote addresses learned from egressing DT frames are logically represented by { MAC address, fine grained label, remote RBridge nickname, confidence, timer }.



## 5. IS-IS Extensions

[[[ Most of the following may be moved to an ISIS draft. ]]]

### 5.1 Announcing RBridge DT Support

An RBridge announces that it is DT in its LSP by ... TBD.

### 5.2 Interested Labels and Bridge Roots sub-TLV

A DT RBridge announces its DT connectivity and related information in the "Interested Labels and Bridge Spanning Tree Roots sub-TLV" (INT-LABEL) which is a variation of the "Interested VLANs and Spanning Tree Roots sub-TLV" (INT-VLAN) structured as below. All fields not defined here are as specified in [[ISIStrill](#)].

```

+---+---+---+---+---+
|Type= INT-LABEL|          (1 byte)
+---+---+---+---+---+
|  Length      |          (1 byte)
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  Interested Labels                                     | (7 bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  Appointed Forwarder Status Lost Counter               | (4 bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Root Bridges                                 | (6*n bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

- o Type: Router Capability sub-TLV Type, set to TBD (INT-LABEL).
- o Length: 13 + 6\*n where n is the number of root bridge IDs.
- o Interested Labels: The Interested Labels field is seven bytes long and formatted as shown below.

```

      0  1  2  3  4  5  6  7
+---+---+---+---+---+---+---+
|M4|M6| R| R| R| R| R| R|
+---+---+---+---+---+---+---+
|          Label.start - 24 bits          |
+---+---+---+---+---+---+---+
|          Label.end - 24 bits            |
+---+---+---+---+---+---+---+

```

- M4, M6: These bits indicate, respectively, that there is an

IPv4 or IPv6 multicast router on a link for which the originating IS is appointed forwarder for every label in the indicated range.

- R: These reserved bits MUST be sent as zero and are ignored on receipt.
- Label.start and Label.end: This fine-grained label ID range is inclusive. A range of one label ID is indicated by setting them both to that label ID value.

### 5.3 The Group Labeled MAC Address sub-TLV

The existing GMAC-ADDR sub-TLV of the Group Address (GADDR) TLV is specified in [[ISIStrill](#)]. It provides for only a 12-bit VLAN-ID. The Group Labeled MAC Address sub-TLV, below, extends this to a 24-bit label.

```

+---+---+---+---+---+
|Type=GLMAC-ADDR|                (1 byte)
+---+---+---+---+---+
|  Length      |                (1 byte)
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  RESV  |      Topology-ID      |  (2 bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          24-Bit Label          | (3 bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Num Group Recs |                (1 byte)
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          GROUP RECORDS (1)          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          .....          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          GROUP RECORDS (N)          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where each group record is of the form:



```

+---+---+---+---+---+
| Num of Sources|          (1 byte)
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Group Address          (6 bytes)          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Source 1 Address      (6 bytes)          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Source 2 Address      (6 bytes)          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          .....                  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Source M Address      (6 bytes)          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

- o Type: GADDR sub-TLV Type, set to TBD (GLMAC-ADDR).
- o Length: Variable, minimum 6.
- o RESV: Reserved. 4-bit field that MUST be sent as zero and ignored on receipt.
- o Topology-ID: This field is not currently used in TRILL, where it is sent as zero and ignored on receipt, but is included for use by other technologies.
- o Label: This carries the 24-bit fine-grained label identifier for all subsequent MAC addresses in this sub-TLV, or the value zero if no label is specified.
- o Number of Group Records: A 1-byte integer that is the number of group records in this sub-TLV.
- o Group Record: Each group record carries the number of sources. It then has a 48-bit multicast address followed by 48-bit source MAC addresses. If the sources do not fit in a single sub-TLV, the same group address may be repeated with different source addresses in another sub-TLV of another instance of the Group Address TLV.

[[[ Most of the above may be moved to an ISIS draft. ]]]



## **6. IANA Considerations**

TBD

## **7. Security Considerations**

See [[RFCtrill](#)] for general RBridge Security Considerations.

As with any communications system, end-to-end encryption and authentication should be considered for particularly sensitive data.

More TBD??

### **7.1 Ingress Forgery and Egress Compromise**

Confusion between a frame with VLAN-X coarse labeling and DT label (X.Y) is a potential problem.

An end station might try to cause a forged DT TRILL Data frame by sending a double C-tagged frame to a port configured for ST ingress.

The requirement in [Section 4.1](#) that all extra C-tags be removed from native frames on input solves this for DT RBridges. After such removal, the DT RBridge will properly add ST or DT to the encapsulated frame. Thus there is no ingress forgery problem for DT RBridges. However, this does not help for ST RBridges.

ST RBridges need only conform to the [[RFCtrill](#)] standard and are not subject to the requirement herein to remove extra C-tags. Thus they might ingress in VLAN-X a native frame double tagged by the end station as (X.Y), removing only the first tag, and then re-insert a VLAN-X tag in the encapsulated frame. The result would be an encapsulated frame that looks like a frame with DT label (X.Y). DT RBridges will think this is a DT frame in (X.Y) and might egress it because they could not distinguish it from a coarsely labeled VLAN-X frame.

Additionally, a TRILL Data frame with DT label (X.Y) could be egressed to VLAN-X by an ST RBridge that is Appointed Forwarder for VLAN-X on one of its ports. Such a frame should not arrive at such an ST RBridge as egress unless the frame is multi-destination.

The above problems are both solved by the prohibition against DT RBridges ingressing to or egressing from DT labeling (X.Y) if the RBridge campus is misconfigured so that an ST edge RBridge is

reporting connectivity to VLAN-X while label (X.Y) is in use.

## **8. References**

The following sections list normative and informative references for this document.

### **8.1 Normative References**

[802.1Q] - IEEE 802.1, "IEEE Standard for Local and metropolitan area networks - Virtual Bridged Local Area Networks", IEEE Std 802.1Q-2011, May 2011.

[RFC2119] - Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997

[RFCtrill] - R. Perlman, D. Eastlake, D. Dutt, S. Gai, and A. Ghanwani, "RBridges: Base Protocol Specification", [draft-ietf-trill-rbridge-protocol-16.txt](#), in RFC Editor's queue.

[ISIStrill] - Eastlake, D., A. Banerjee, D. Dutt, R. Perlman, A. Ghanwani, "TRILL Use of IS-IS", [draft-ietf-isis-trill-05.txt](#), in RFC Editor's queue.

### **8.2 Informative References**

[RFCaf] - Perlman, R., D. Eastlake, A. Banerjee, H. Fangwei, "RBridges: Appointed Forwarders", [draft-ietf-trill-rbridge-af-03.txt](#), work in progress.



## Acknowledgements

The comments and contributions of the following are gratefully acknowledged:

Anoop Ghanwani, Sujay Gupta, Jon Hudson, Vishwas Manral, and Erik Nordmark.

## Authors' Addresses

Donald Eastlake 3rd  
Huawei Technologies  
155 Beaver Street  
Milford, MA 01757 USA

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

Mingui Zhang  
Huawei Technologies Co., Ltd  
HuaWei Building, No.3 Xinxu Rd., Shang-Di  
Information Industry Base, Hai-Dian District,  
Beijing, 100085 P.R. China

Email: zhangmingui@huawei.com

Puneet Agarwal  
Broadcom Corporation  
3151 Zanker Road  
San Jose, CA 95134 USA

Phone: +1-949-926-5000  
Email: pagarwal@broadcom.com

Dinesh G. Dutt  
Cisco Systems  
170 Tasman Drive  
San Jose, CA 95134-1706 USA

Phone: +1-408-527-0955  
Email: ddutt@cisco.com

Radia Perlman



INTERNET-DRAFT

RBridges: Fine-Grained Labeling

2200 Mission College Blvd.  
Santa Clara, CA 95054 USA

Phone: +1-408-765-8080

Email: [Radia@alum.mit.edu](mailto:Radia@alum.mit.edu)



## Copyright, Disclaimer, and Additional IPR Provisions

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License. The definitive version of an IETF Document is that published by, or under the auspices of, the IETF. Versions of IETF Documents that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of IETF Documents. The definitive version of these Legal Provisions is that published by, or under the auspices of, the IETF. Versions of these Legal Provisions that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of these Legal Provisions. For the avoidance of doubt, each Contributor to the IETF Standards Process licenses each Contribution that he or she makes as part of the IETF Standards Process to the IETF Trust pursuant to the provisions of [RFC 5378](#). No language to the contrary, or terms, conditions or rights that differ from or are inconsistent with the rights and licenses granted under [RFC 5378](#), shall have any effect and shall be null and void, whether published or posted by such Contributor, or included with or in such Contribution.

