

TRILL Working Group
INTERNET-DRAFT
Intended status: Proposed Standard

Radia Perlman
Intel Labs
Anil Rijhsinghani
HP Networking
Donald Eastlake
Huawei
Ayan Banerjee
Dinesh Dutt
Cumulus Networks
January 1, 2013

Expires June 30, 2013

TRILL: Campus VLAN and Priority Regions
<[draft-ietf-trill-rbridge-vlan-mapping-08.txt](#)>

Abstract

Within a TRILL campus, the VLAN and priority of TRILL encapsulated frames is preserved. However, in some cases it may be desired that data VLAN and/or priority be mapped at the boundary between regions of such a campus. This document describes an optional RBridge feature to provide this function.

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/1id-abstracts.html>. The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Table of Contents

1. Introduction.....	3
1.1 TRILL Campus Regions.....	4
1.2 Terminology.....	5
2. Internal and Cut Set Configuration and Mappings.....	6
2.1 Multiple Crossings.....	7
2.2 Native Frame Considerations.....	8
2.3 More than Two Regions.....	8
2.4 Mapping Implementation.....	9
3. End Node Address Learning Between Regions.....	11
4. Cut Set Attraction of VLANs and Multicast.....	12
5. Advertisement of VLAN and Priority Mappings.....	13
6. IANA Considerations.....	13
7. Security Considerations.....	13
8. Normative References.....	14
9. Informative References.....	14
Appendix Z: Change Summary.....	15

1. Introduction

The IETF TRILL protocol provides transparent forwarding, with a number of additional features, by use of link state routing and encapsulation with a hop count as specified in [\[RFC6325\]](#).

Devices implementing the TRILL protocol are called TRILL switches or R Bridges (Routing Bridges). A TRILL campus is an area of TRILL switches and possibly bridges bounded by and interconnecting end stations and Layer 3 routers, analogous to a customer bridge LAN (which is an area of bridges interconnecting end stations, routers, and TRILL switches). In a TRILL campus, native frames (as defined in [\[RFC6325\]](#)), when they arrive at their first or ingress R Bridge, are encapsulated, routed in encapsulated form via zero or more transit R Bridges, and finally decapsulated and delivered by their egress R Bridge or R Bridges.

TRILL switch ports have some features specified in IEEE 802.1Q as described in [\[RFC6325\]](#), with TRILL being implemented above those ports. Such ports provide for the association of incoming frames with a particular frame priority and customer VLAN. (See [Appendix D of \[RFC6325\]](#).)

Bridge ports can map frame priorities, a process called "priority regeneration" in IEEE 802.1. In addition, some bridge products provide a feature to map the customer VLAN of incoming VLAN tagged frames, a process of the type called "VLAN ID translation" in IEEE 802.1.

Using such port features, it is possible to configure R Bridge ports to map the priority and/or VLAN of native frames being received for ingress or to map the priority and/or VLAN of the frame inside a TRILL data frame (as defined in [\[RFC6325\]](#)) after it has been decapsulated for egress through an output port. But priority and/or VLAN mapping of the outer priority and VLAN (Outer.VLAN) of a TRILL encapsulated data frame has no effect on the Inner.VLAN tag in the encapsulated frame. In TRILL, the Inner.VLAN tag gives the real VLAN and priority of the data and these are unaffected by any port features that change only the Outer.VLAN priority or VLAN.

(Note: VLAN mapping is also referred to in [\[RFC6325\]](#). However, that reference concerns Outer VLAN mapping within a link between neighbor R Bridges, a condition that may require the R Bridges connected to such a link to take precautions as described in [Section 4.4.5 of \[RFC6325\]](#).)

The default for TRILL is to provide connectivity between all end station and router ports in the same VLAN. However, there are cases

where it may be desirable to have the same VLAN in different regions of a TRILL campus mean different things. In that case, it would be

necessary for end stations or Layer 3 routers in that VLAN not to be connected if they are in different regions. It might also be desirable to have connectivity between end stations in different regions that are in different VLANs if those different VLANs in their different regions actually indicate membership in the same Layer 2 community. Similar circumstances can arise for priority. This document describes how to achieve this through an optional TRILL feature.

An example of where this feature might be useful would be the merger of two organizations which previously had separate networks. They might desire to combine these networks into a new unified network under unified control; however, for some period of time, there might be disagreements between the previously separate networks as to VLAN and/or priority assignments requiring mapping at any points of interconnection. If these were Layer 2 networks, and particularly if they were TRILL campuses, combination into a single unified TRILL campus would be natural; but, this would probably require mapping facilities, such as those specified herein, between the regions of the new unified campus that had previously been separate networks.

Considerations related to service or S-VLANs are beyond the scope of this document.

1.1 TRILL Campus Regions

The set of RBridges interconnecting different regions of an TRILL campus are known as the "cut set", meaning that if that set of RBridges is removed, the regions are disconnected from each other.

RBridges in the cut set can be configured to translate some set of VLAN IDs in one region to different VLAN IDs when forwarding from that region to another region and/or to block encapsulated frames with certain VLAN IDs. They can be similarly configured for priority.

This feature is accomplished solely by configuring RBridges in the cut set. No other RBridges need even be aware that the feature is in use. In particular, use of this feature has no effect on the path (sequence of RBridges) followed by TRILL Data frames (except that for multi-destination frames, tree pruning may be affected). The TRILL features of optimum routing and of optional multi-pathing of both unicast and multi-destination frames are unaffected.

This document explains how to implement this feature in RBridges. In this document we will usually assume there are two regions, "East" and "West", and RBridges RB1, RB2, and RB3 that interconnect the two regions and constitute the cut set as shown in Figure 1. Extension to

more than two regions is straightforward and will also be briefly

described.

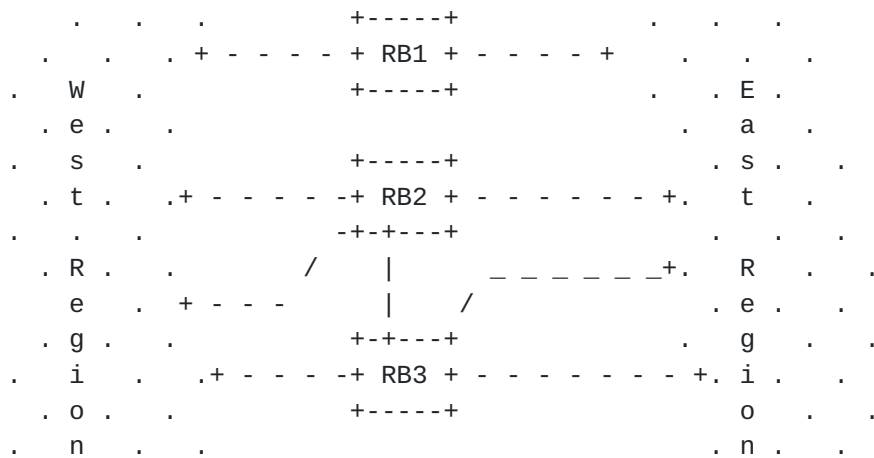


Figure 1.

General familiarity with the TRILL base protocol standard [[RFC6325](#)] is assumed in this document.

1.2 Terminology

The same terminology and acronyms are used in this document as in [[RFC6325](#)]. "Cut set" is defined above. We will refer to RBridges other than the cut set of RBridges as "internal RBridges".

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. Internal and Cut Set Configuration and Mappings

Internal RBridges will not be aware that VLAN and priority mapping is going on and require no configuration. They will behave exactly as they would without mapping. The only evidence they might have of VLAN or priority mapping is the existence of an optional informational sub-TLV that a cut set RBridge, RB1, MAY include in its LSP, listing the mappings that RB1 is configured to be performing. Internal RBridges will ignore this information field. It is there for detection of misconfiguration.

Cut set RBridges are configured as follows:

If VLAN A in region "East" is to be translated into VLAN B in region "West", each cut set RBridge MUST be configured, for every port, as to whether that port is in East or West, and configured with VLAN mappings, such as:

"East/VLAN A -----> West/VLAN B"

That mapping means that a TRILL Data frame with an Inner.VLAN of A received by RB1 on a port configured to be in East and forwarded to a port configured to be in West is forwarded with the Inner.VLAN changed to B. It is possible to configure asymmetric mappings; however, such asymmetric have negative consequences as described below. For the above mapping to be symmetrically configured, it would be necessary to also configure the cut set RBridge in question so that frames arriving from West in VLAN B would also be mapped to VLAN A if they are destined for East, that is

"West/VLAN B <-----> East/VLAN A"

Figure 2.

Mappings of the priority of encapsulated frames are configured in the same way.

The requirement that every port of a cut set RBridge MUST be configured as to which region it is in applies even to ports for a link between cut set RBridges such as the link between RB2 and RB3 in Figure 1. The TRILL encapsulated data frames on that link have a normal Inner.VLAN with a VLAN ID and priority. In a campus with multiple regions, a VLAN ID or priority is, in general, meaningless unless you know the region in which it occurs. So some specific region must be chosen for such a link.

All cut set RBridges between a pair of regions SHOULD be configured similarly if, as is normally the case, it is desired that the mapping of a TRILL Data frame going between those regions will be independent

of which cut set RBridge the frame traverses.

R. Perlman, et al.

[Page 6]

The default VLAN and priority mapping is the mapping that leaves VLAN IDs and priorities unchanged. If a mapping has been specified for both the VLAN and priority of a frame, both mappings are applied.

2.1 Multiple Crossings

Under some circumstances, a frame could pass through cut set RBridges between a pair of regions more than once and thus have its VLAN and priority mapped more than one. This is true of both known unicast and multi-destination frames. For example, in Figure 3, if the link between RBwest1 and RBwest2 fails, then the shortest path from RBwest1 to RBwest2 may be through RBcut1, RBeast1, and RBcut2. In addition, multi-destination frames are sent via a distribution tree which might constrain such frames going between RBwest1 and RBwest2 to be routed through RBeast1.

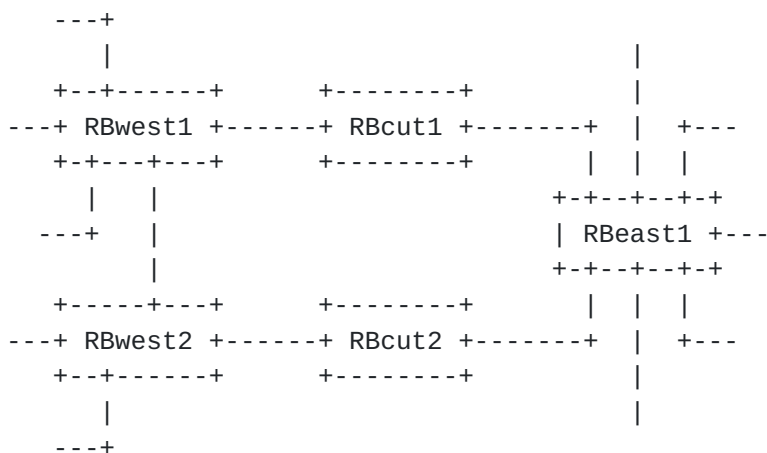


Figure 3.

If all of the mappings at RBcut1 and RBcut2 are symmetric then the VLAN and/or priority of such frames going from west to west via east might get mapped twice but the second mapping would restore them to their original value. Symmetric means, for example, that if RB1 is translating from "VLAN A" to "VLAN B" when forwarding from East to West, it will translate tag "VLAN B" to tag "VLAN A" when forwarding from West to East (see Figure 2).

However, assume that RBcut1 and RBcut2 are configured with asymmetric mappings. Then multiple cut set transit may cause problems. For example, if VLAN A in west is mapped to VLAN B in east and VLAN B in east is mapped to VLAN C in west, then the above scenario could lead to frames in VLAN A from west to west being unexpectedly mapped to VLAN C causing connectivity between VLANs A and C in west and failure to deliver the frame as intended. Similar considerations apply to

priority mappings. The probability of such situations can be

minimized by providing rich interconnectivity within each region and increasing the cost of links to cut set RBridges, so that frames internal to a regions will be routed internally to that region except in cases of low probability multiple failures. It is generally safest to configure symmetric mappings.

2.2 Native Frame Considerations

If the processing model described in [\[RFC6325\]](#) is followed, then no special handling is necessary for the case where a cut set RBridge receives or transmits a native data frame, that is, where the cut set RBridge is also an ingress or egress RBridge. In particular, the processing model used in [\[RFC6325\]](#) provides that an ingressed native frame is always encapsulated, even if it is to be immediately decapsulated and delivered out a different port of the same RBridge in native form. (Of course, implementers are free to handle this in other ways provided the external behavior is the same.) Thus, following this processing model, no changes are needed in an implementation model of VLAN and priority mapping described entirely in terms of the manipulation of the Inner.VLAN tag of TRILL encapsulated frames.

On the other hand, if there are no RBridges in a region, say region West in Figure 1, then all frames will arrive from that region at the cut set RBridges as unencapsulated native frames and all native frames sent into that region will be unencapsulated. Under these limited circumstances, traditional bridge port VLAN and priority mapping could work to assist in performing the inter-regional mappings described in this document.

2.3 More than Two Regions

A TRILL campus may have more than two regions. An RBridge is in the cut set between any pair of such regions if and only if it has at least one port in each of the regions. There may be pairs of regions that, because of intervening regions, have no cut set RBridges connected to them both.

Every RBridge that is in any cut set MUST have every port configured as to which region that port is in. Every RBridge port on a link between two or more cut set RBridges, such as that shown between RB2 and RB3 in Figure 1, SHOULD be configured to be in the same region. The mappings performed on TRILL data frames transiting a cut set RBridge that has ports in three or more regions depend only on the region of that frame's input and output ports and are unaffected by

what region any other ports of that RBridge might be connected to.

It is RECOMMENDED that not only should any mappings be symmetric at every cut set RBridge in a campus that implements the VLAN and priority mapping feature but that all cut set RBridges in the campus should be configured so as to be transitively symmetric and similar. That is, the mapping of the VLAN and priority in a frame going from region A to region Z should be independent of the path that frame follows in the campus and symmetric with the mapping to which any frame going from region Z to region A would be subjected.

2.4 Mapping Implementation

If RB1 is configured to believe port X is in "East" and port Y is in "West", and RB1 is configured such that "East/VLAN A ----> West/VLAN B", then when RB1 forwards data frames from port X to port Y, if the received frame from port X has Inner.VLAN tag VLAN ID equal to VLAN A, then RB1 changes that VLAN ID from VLAN A to VLAN B before it forwards out port Y. Similarly, if priority mapping has been configured, the Inner.VLAN priority field is mapped.

This mapping is performed whether RB1 is the appointed forwarder on port X for VLAN A and the frame arrives unencapsulated, or whether the frame has arrived already encapsulated as a TRILL Data frame. Likewise, RB1 performs the same VLAN and priority mapping, depending on input and output port, whether the frame is to a known unicast address or is multi-destination.

RBridges may implement campus region VLAN and priority mapping in any way desired so long as the externally visible behavior matches this specification. Two example models of internal processing are described below.

In the forwarding-oriented model, VLAN and priority mappings occur once as part of the inter-port forwarding process and depend on the ordered pair on input-port-region and output-port-region.

If the port-oriented model, VLAN and priority mappings occur once or twice associated with input and/or output ports. For example, for VLANs, each input port of a cut set RBridge could (after encapsulation in the case of a native frame) map the Inner.VLAN to a value in an RBridge specific generic VLAN space, with the mapping dependent only on the region to which that input port was assigned. Then, the output port through which the frame was sent would map from that general VLAN space to a specific VLAN in the Inner.VLAN with the mapping depending only on the region to which the output port was assigned. Either mapping could be the mapping that did not change the VLAN ID and/or priority. A similar model could be used for priority mapping with similar considerations.

These two processing models are logically interchangeable.

3. End Node Address Learning Between Regions

RBridges by default learn end node MAC addresses and VLANs from the observation of ingressed native frames and the decapsulation of native frame at egress, as described in [[RFC6325](#)]. This process requires no modification at internal RBridges to accommodate VLAN mapping as described herein as the VLAN will be appropriate for the region where it is observed.

For a cut set RBridge, each port is specified to be in a particular region. For such an RBridge, the VLAN portion of the addresses learned at a port providing direct end station service will be that VLAN in the region to which the cut set RBridge has assigned the port. Care must be taken within a cut set RBridge when using such learned information. For example, if a native frame is received in VLAN X from region Y destined for MAC address Z, then address Z can be looked up in the address information learned for other regions only after applying any mapping for VLAN X to that region.

TRILL also allows RBridges to optionally advertise attached end nodes. This end node advertisement uses the TRILL ESADI (End System Address Distribution Information) protocol. Because TRILL ESADI frames do not include the VLAN to which they are applicable anywhere except in their Inner.VLAN tag and ESADI frames are forwarded just like ordinary multi-destination TRILL Data frames, the VLAN mapping described above works for ESADI learning. Because of this, any future ESADI extensions MUST NOT require VLAN ID fields inside the ESADI frame payload.

4. Cut Set Attraction of VLANs and Multicast

The above described mechanisms are all that is required for VLAN and priority mapping of frames sent to known unicast addresses. However, to correctly handle multi-destination traffic, additional steps are required. In particular, unless cut-set RBridges take additional action, multi-destination frames that they need to forward from one region to another might not reach the cut set RBridge due to the optional pruning of distribution trees by internal RBridges.

If RB1 is configured to translate VLAN A in East to VLAN B in West, then RB1 MUST report, in its LSP, that it is interested in both VLAN A and VLAN B data, even if RB1 is not appointed forwarder for either or both VLAN A or VLAN B. If it did not do this, a multi-destination frame in VLAN A in East might be pruned before reaching RB1 and not mapped to VLAN B and forwarded to West as it should.

If RB1 is configured to translate VLAN A in East to VLAN B in West, then RB1 MUST take steps to ensure that a multicast packet for group G in VLAN A will not be filtered inside the East region. To solve this problem RB1 MUST report that it is connected in VLAN A to an IPv4 and IPv6 multicast router so it will get all multi-cast traffic in VLAN A and can forward appropriate multicast frames mapped to VLAN B. While this increases traffic to cut set RBridges, it does so to an extent no worse than an RBridge connected to an actual Layer 3 multicast router or routers.

Because all the regions operate as a single TRILL campus with a unified IS-IS link state database, it is not possible to confine the above required announcements to particular regions.

Cut set RBridges and the links connecting them to the rest of the network should be appropriately engineered for any additional traffic load these requirements impose.

5. Advertisement of VLAN and Priority Mappings

To help detect misconfiguration, a cut set RBridge RB1 MAY advertise its VLAN and priority mappings in its LSP. To enable this, a 16-bit unsigned ID is assigned to each of the regions by manual configuration. All cut set RBridges SHOULD be configured with the same IDs for the regions but means of accomplishing this are outside of the scope of this document. So, in our example Figure 1, if "East" is "1" and "West" is "2", and VLAN A in East is mapped to VLAN B in West, and vice versa to be symmetric, the LSP would report a set of mappings, including:

{VLAN: (1:A,2:B), (2:B,1:A)}

Illegal VLAN IDs (0x000 or 0xFFFF) should never appear as a VLAN ID in an LSP advertising VLAN mappings but if they do, the mapping where they appear are ignored for consistency checking.

The actual encoding of this information and the Type or sub-Type values for any new TLV or sub-TLV data elements are specified in a separate document

6. IANA Considerations

This document requires no IANA actions. RFC Editor: Please delete this section before publication.

7. Security Considerations

See [[RFC6325](#)] for general RBridge Security Considerations.

If cut set RBridges have misconfigured VLAN mappings, VLANs may be inadvertently partitioned or inadvertently merged and frames may be delivered in the wrong VLAN, which could violate security policies. However, misconfiguration of VLAN or priority mappings cannot cause loops because mappings of VLANs and/or priority have no effect on unicast frame routing, shortest path calculations, distribution tree construction or selection, or the like.

8. Normative References

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

[RFC6325] - Perlman, R., Eastlake 3rd, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification", [RFC 6325](#), July 2011.

9. Informative References

None.

Appendix Z: Change Summary

RFC Editor Note: Please delete this [Appendix Z](#) on publication.

Changes from -00 to -01

1. Because RBridges cannot tell what cloud other RBridges are in, drop the "optimized" option for advertising multicast listeners and require the advertisement of multicast router connectivity.
2. Specify that the cloud connectivity must be specified for all cut set RBridges and that cloud IDs are manually configured and are 16 bit.
3. Expand rules for VLAN ID mapping/handling at a cut set RBridge so as to drop frames that are for a VLAN ID to which another VLAN ID is being mapped. (See [Section 3](#).)
4. Add mention of "VLAN ID translation", the 802.1 name for VLAN mapping.
5. Minor editing changes.

Changes from -01 to -02

1. Remove previous confused text about VLAN mapping (point 3 in changes from -00 to -01).
2. Add text allowing mapping to zero to indicate frames should be dropped. Add text and diagram explaining that this can lead to VLAN partition.
3. Add normative reference to [draft-ietf-isis-layer2](#).
4. Minor editing changes.

Changes from -02 to -03

This was a substantial re-write of the draft but there was no fundamental conceptual change in the mapping mechanism.

1. Replace "cloud" with "region".
2. Introductory material was re-written to primarily reference

RBridge campuses and reduce references to 802.1 bridges.

3. Mapping of priority was added to mapping of VLANs.
4. Two different models are now described for implementation of mappings, one in the forwarding mechanism as before and one associated with the RBridge ports.
5. Add the specification of the TRILL GenApp TLV. Switch to using TRILL GenApp TLV sub-TLVs to advertise VLAN and priority mappings. Add specification of those sub-TLVs. Remove reference to [draft-ietf-isis-layer2](#).
6. The IANA considerations section calls for the allocation of a GenApp TLV code for TRILL and provides for sub-TLVs under that code where the LSP advertisement of VLAN and priority mappings was moved. Set up IANA registry for TRILL GenApp sub-TLVs.
7. Numerous minor editing changes.

Changes from -03 to -04

1. Because distribution trees for multi-destination frames may cause frames to cross region boundaries multiple times even to get between RBridges within a single regions, remove facilities for dropping frames at region boundaries.
2. Due to questions about the timing of the approval of the IS-IS GenApp draft, move VLAN/priority mapping informational advertisement code points and data structures to a separate draft.
3. Numerous minor editing changes.

Changes from -04 to -05

Increment version and update dates. Update author info. One or two minor editorial changes.

Changes from -05 to -06

Update draft reference to [[RFC6325](#)]. Increment version and update dates.

Changes from -06 to -07

Update author information, increment version, and update dates.

Changes from -07 to -08

Minor editorial changes, update author information, increment version, and update dates.

Authors' Addresses

Radia Perlman
Intel Labs
2200 Mission College Blvd.
Santa Clara, CA 95054-1549 USA

Phone: +1-408-765-8080
Email: Radia@alum.mit.edu

Anil Rijhsinghani
HP Networking
350 Campus Drive
Marlboro, MA 01752-3082 USA

Phone: +1-508-323-1251
Email: anil.rijhsinghani@hp.com

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

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

Ayan Banerjee
Cumulus Networks
1089 West Evelyn Avenue
Sunnyvale, CA 94086 USA

Email: ayabaner@gmail.com

Dinesh G. Dutt
Cumulus Networks
1089 West Evelyn Avenue
Sunnyvale, CA 94086 USA

Email: ddutt.ietf@hobbesdutt.com

Copyright, Disclaimer, and Additional IPR Provisions

Copyright (c) 2013 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.

