

Internet Engineering Task Force (IETF)
Request for Comments: 6439
Updates: 6325
Category: Standards Track
ISSN: 2070-1721

R. Perlman
Intel Labs
D. Eastlake 3rd
Y. Li
Huawei Technologies
A. Banerjee
Cisco Systems
F. Hu
ZTE Corporation
November 2011

Routing Bridges (RBridges): Appointed Forwarders

Abstract

The IETF TRILL (TRAnsparent Interconnection of Lots of Links) protocol provides least cost pair-wise data forwarding without configuration in multi-hop networks with arbitrary topology, safe forwarding even during periods of temporary loops, and support for multipathing of both unicast and multicast traffic. TRILL accomplishes this by using IS-IS (Intermediate System to Intermediate System) link state routing and by encapsulating traffic using a header that includes a hop count. Devices that implement TRILL are called "RBridges" (Routing Bridges).

TRILL supports multi-access LAN (Local Area Network) links that can have multiple end stations and RBridges attached. Where multiple RBridges are attached to a link, native traffic to and from end stations on that link is handled by a subset of those RBridges called "Appointed Forwarders", with the intent that native traffic in each VLAN (Virtual LAN) be handled by at most one RBridge. The purpose of this document is to improve the documentation of the Appointed Forwarder mechanism; thus, it updates RFC 6325.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6439>.

Copyright Notice

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.

Table of Contents

1. Introduction	2
1.1. Terminology and Acronyms	3
2. Appointed Forwarders and Their Appointment	4
2.1. Appointment Effects of DRB Elections	5
2.2. Appointment and Removal by the DRB	5
2.2.1. Processing Forwarder Appointments	6
2.2.2. Frequency of Appointments	7
2.2.3. Appointed Forwarders Limit	8
2.3. Local Configuration Action Appointment Effects	8
2.4. VLAN Mapping within a Link	9
3. The Inhibition Mechanism	9
4. Inhibited Appointed Forwarder Behavior	11
5. Multiple Ports on the Same Link	12
6. Security Considerations	12
7. Acknowledgements	13
8. References	13
8.1. Normative References	13
8.2. Informative References	13
Appendix. VLAN Inhibition Example	14

1. Introduction

The IETF TRILL (TRAnsparent Interconnection of Lots of Links) protocol [RFC6325] provides optimal pair-wise data frame forwarding without configuration in multi-hop networks with arbitrary topology, safe forwarding even during periods of temporary loops, and support for multipathing of both unicast and multicast traffic. TRILL accomplishes this by using IS-IS (Intermediate System to Intermediate System) [IS-IS] [RFC1195] link state routing and encapsulating traffic using a header that includes a hop count. The design

supports VLANs (Virtual Local Area Networks) and optimization of the distribution of multi-destination frames based on VLANs and IP-derived multicast groups. Devices that implement TRILL are called "RBridges" (Routing Bridges).

Section 2 of [RFC6327] explains the environment for which the TRILL protocol is designed and the differences between that environment and the typical Layer 3 routing environment.

TRILL supports multi-access LAN (Local Area Network) links that can have multiple end stations and RBridges attached. Where multiple RBridges are attached to a link, native traffic to and from end stations on that link is handled by a subset of those RBridges called "Appointed Forwarders", with the intent that native traffic in each VLAN be handled by at most one RBridge. An RBridge can be Appointed Forwarder for many VLANs.

The purpose of this document is to improve the documentation of the Appointed Forwarder mechanism; thus, it updates RFC 6325. It includes reference implementation details. Alternative implementations that interoperate on the wire are permitted.

The Appointed Forwarder mechanism is irrelevant to any link on which end station service is not offered. This includes links configured as point-to-point IS-IS links and any link with all RBridge ports on that link configured as trunk ports. (In TRILL, configuration of a port as a "trunk port" just means that no end station service will be provided. It does not imply that all VLANs are enabled on that port.)

The Appointed Forwarder mechanism has no effect on the formation of adjacencies, the election of the Designated RBridge (DRB) for a link, MTU matching, or pseudonode formation. Those topics are covered in [RFC6327]. Furthermore, Appointed Forwarder status has no effect on the forwarding of TRILL Data frames. It only affects the handling of native frames.

For other aspects of the TRILL base protocol, see [RFC6325] and [RFC6327]. Familiarity with [RFC6325] and [RFC6327] is assumed in this document. In case of conflict between this document and [RFC6325], this document prevails.

1.1. Terminology and Acronyms

This document uses the acronyms defined in [RFC6325].

A "trunk port" is a port configured with the "end station service disable" bit on, as described in Section 4.9.1 of [RFC6325].

In this document, the term "link" means "bridged LAN", that is to say some combination of physical links with zero or more bridges, hubs, repeaters, or the like.

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 [RFC2119].

2. Appointed Forwarders and Their Appointment

The Appointed Forwarder on a link for VLAN-x is the RBridge that ingresses native frames from the link and egresses native frames to the link in VLAN-x. By default, the DRB (Designated RBridge) on a link is in charge of native traffic for all VLANs on the link. The DRB may, if it wishes, act as Appointed Forwarder for any VLAN and it may appoint other RBridges that have ports on the link as Appointed Forwarder for one or more VLANs.

It is important that there not be two Appointed Forwarders on a link that are ingressing and egressing native frames for the same VLAN at the same time. Should this occur, it could form a loop where frames are not protected by a TRILL Hop Count for part of the loop. (Such a condition can even occur through two Appointed Forwarders for two different VLANs, VLAN-x and VLAN-y, if ports or bridges inside the link are configured to map frames between VLAN-x and VLAN-y as discussed in Section 2.4.) While TRILL tries to avoid such situations, for loop safety there is also an "inhibition" mechanism (see Section 3) that can cause an RBridge that is an Appointed Forwarder to not ingress or egress native frames.

As discussed in Section 5, an RBridge may have multiple ports on a link. As discussed in [RFC6327], if there are multiple ports with the same Media Access Control (MAC) address on a link, all but one will be suspended. The case of multiple ports on a link for one RBridge and the case of multiple ports with the same MAC address on a link and combinations of these cases are fully accommodated; however, multiple ports on a link for one RBridge is expected to be a rare condition and duplicate MAC addresses are not recommended by either TRILL or IEEE 802.1 standards.

Appointed Forwarder status has no effect on the forwarding of TRILL Data frames. It only affects the handling of native frames.

There are three mechanisms by which an RBridge can be appointed or un-appointed as Appointed Forwarder: as a result of the DRB elections [RFC6327] as discussed in Section 2.1, as a result of action by the DRB as discussed in Section 2.2, as a result of a local configuration action as discussed in Section 2.3.

2.1. Appointment Effects of DRB Elections

When an RBridge believes that it has become the DRB on a link, by default, it can act as Appointed Forwarder for any VLANs on that link that it chooses as long as its port is not configured as a trunk port and has that VLAN enabled (or at least one of its ports meets these criteria, if it has more than one port on the link).

An RBridge loses all Appointed Forwarder status when:

1. it decides that it has lost the status of being the DRB for a link; or
2. it observes a change in the RBridge that is the DRB for the link without itself becoming the DRB.

In the rare corner case where an RBridge has more than one port on a link, one of which was previously the DRB election winner but has just lost the DRB election to a different port of the same RBridge (possibly due to management configuration of port priorities), there is no change in which RBridge is the DRB. Therefore, neither of the above points applies and there is no change in Appointed Forwarder status.

2.2. Appointment and Removal by the DRB

The DRB may appoint other RBridges on the link through inclusion of one or more Appointed Forwarders sub-TLVs [RFC6326] in a TRILL Hello it sends on the Designated VLAN out the port that won the DRB election. When the DRB sends any appointments in a TRILL Hello, it must send all appointments for that link in that Hello. Any previous appointment not included is implicitly revoked.

Although the DRB does not need to announce the VLANs for which it has chosen to act as Appointed Forwarder by sending appoints for itself, if the DRB wishes to revoke all appointments for RBridges other than itself on the link, it is recommended that it send a TRILL Hello with an appointment for itself for some VLAN.

The DRB MUST NOT send any appointments on a link unless its DRB inhibition timer (see Section 3) for that link is expired.

How the DRB decides what other RBridges on the link, if any, to appoint forwarder for which VLANs is beyond the scope of this document.

2.2.1. Processing Forwarder Appointments

When a non-DRB RBridge that can offer end station service on a link receives a TRILL Hello that is not discarded for one of the reasons given in [RFC6327], it checks the source MAC address and the Port ID and System ID in the Hello to determine if it is from the winning DRB port. If it is not from that port, any Appointed Forwarder sub-TLVs in the Hello are ignored, and there is no change in the receiving RBridge's Appointed Forwarder status. Also, if no Appointed Forwarder sub-TLVs are present in the TRILL Hello, there is no change in the receiver's Appointed Forwarder status.

However, if the TRILL Hello is from the winning DRB port and the Hello includes one or more Appointed Forwarder sub-TLVs, then the receiving RBridge becomes appointed for the VLANs that are both listed for it in the Hello and are enabled on the receiving port. (If the appointment includes VLAN IDs 0x000 or 0xFFFF, they are ignored, but any other VLAN IDs are still effective.) If the receiver was Appointed Forwarder for any other VLANs, its Appointed Forwarder status for such other VLANs is revoked. For example, if none of these sub-TLVs in a Hello appoints the receiving RBridge, then it loses all Appointed Forwarder status and is no longer Appointed Forwarder for any VLAN on the port where the Hello was received.

The handling of one or more Appointed Forwarder sub-TLVs in a Hello from the winning port that appoints the receiving RBridge is as follows. An appointment in an Appointed Forwarder sub-TLV is for a specific RBridge and a contiguous interval of VLAN IDs; however, as stated above, it actually appoints that RBridge forwarder only for the VLAN(s) in that range that are enabled on one or more ports that RBridge has on the link (ignoring any ports configured as trunk ports or as IS-IS point-to-point ports). If the RBridge was Appointed Forwarder for any additional VLANs beyond the VLANs for which it was being appointed, it loses Appointed Forwarder status for such additional VLANs.

There is no reason for an RBridge to remember that it received a valid appointment message for a VLAN that was ineffective because the VLAN was not enabled on the port where the message was received or because the port was a trunk or point-to-point port. It does not become Appointed Forwarder for such a VLAN just because that VLAN is later enabled or the port later reconfigured.

It should be straightforward for the DRB to send, within one Hello, the appointments for several dozen VLAN IDs or several dozen blocks of contiguous VLAN IDs. Should the VLANs the DRB wishes to appoint be inconveniently distributed, for example, the proverbial case where the DRB RB1 wishes to appoint RB2 forwarder for all even-numbered VLANs and appoint RB3 forwarder for all odd-numbered VLANs, the following method may be used. The network manager normally controls what VLANs are enabled on RBridge port. Thus, the network manager can appoint an RBridge forwarder for an arbitrary set of scattered VLANs by enabling only those VLANs on the relevant port (or ports) and then having the DRB send an appointment that appears to appoint the target RBridge forwarder for all VLANs. However, for proper operation and inter-RBridge communication, the Designated VLAN for a link SHOULD be enabled on all RBridge ports on that link, and it may not be desired to appoint the RBridge forwarder for the Designated VLAN. Thus, in the general case, it would require two appointments, although it would still only require one appointment if the Designated VLAN were an extreme low or high value such as VLAN 0xFFFE or the default VLAN 1.

For example, assume the DRB wants RB2 to be Appointed Forwarder for all even-numbered VLANs and the Designated VLAN for the link is VLAN 101. The network manager could cause all even-numbered VLANs plus VLAN 101 to be enabled on the relevant port of RB2 and then, with the desired effect, cause the DRB to send appointments to RB2 appointing it forwarder for all VLANs from 1 through 100 and from 102 through 4,094.

Should the network manager have misconfigured the enabled VLANs and Appointed Forwarders, resulting in two RBridges believing they are Appointed Forwarders for the same VLAN, then item 4 in Section 3 will cause one or more of the RBridges to be inhibited for that VLAN.

2.2.2. Frequency of Appointments

It is not necessary for the DRB to include the forwarder appointments in every TRILL Hello that it sends on the Designated VLAN for a link. For loop safety, every RBridge is required to indicate, in every TRILL Hello it sends in VLAN-x on a link, whether it is an Appointed Forwarder for VLAN-x for that link (see item 4 in Section 3). It is also RECOMMENDED that the DRB have all VLANs for which end station service will be offered on the link as well as the Designated VLAN, enabled. Thus, the DRB will generally be informed by other RBridges on the link of the VLANs for which they believe they are Appointed Forwarder. If this matches the appointments the DRB wishes to make, it is not required to re-send its forwarder appointments; however, for robustness, especially in cases such as VLAN misconfigurations in

a bridged LAN link, it is RECOMMENDED that the DRB send its forwarder appointments on the Designated VLAN at least once per its Holding Time on the port that won the DRB election.

2.2.3. Appointed Forwarders Limit

The mechanism of DRB forwarder appointment and the limited length of TRILL Hellos impose a limit on the number of RBridges on a link that can be Appointed Forwarders. To obtain a conservative estimate, assume that no more than 1000 bytes are available in a TRILL Hello for such appointments. Assume it is desired to appoint various RBridges on a link forwarder for arbitrary non-intersecting sets of VLANs. Using the technique discussed above would generally require two appointments, or 12 bytes, per RBridge. With allowance for sub-TLV and TLV overhead, appointments for 83 RBridges would fit in under 1000 bytes. Including the DRB, this implies a link with 84 or more RBridges attached. Links with more than a handful of RBridges attached are expected to be rare.

Note: If the Designated VLAN were an extreme low or high value, such as VLAN 1, which is the default and may be a common value in practice, only 6 bytes per RBridge would be required. This would permit twice as many different Appointed Forwarder RBridges than indicated by the general analysis above or, alternatively, would take only half as much space to appoint the same number of Appointed Forwarders.

Unnecessary changes in Appointed Forwarders SHOULD NOT be made as they may result in transient lack of end station service. Large numbers of Appointed Forwarders on a link (in excess of 65) are NOT RECOMMENDED due to the complexity of their establishment and maintenance.

2.3. Local Configuration Action Appointment Effects

Disabling VLAN-x at an RBridge port cancels any Appointed Forwarder status that RBridge has for VLAN-x unless VLAN-x is enabled on some other port that the RBridge has connected to the same link. Configuring a port as a trunk port or point-to-point port revokes any Appointed Forwarder status that depends on enabled VLANs at that port.

Causing a port to no longer be configured as a trunk or point-to-point port or enabling VLAN-x on a port does not, in itself, cause the RBridge to become an Appointed Forwarder for the link that port is on. However, such actions can allow the port's RBridge to become Appointed Forwarder by choice if it is the DRB or by appointment, if it is not the DRB on the link.

2.4. VLAN Mapping within a Link

TRILL Hellos include a field that is set to the VLAN in which they are sent. If they arrive on a different VLAN, then VLAN mapping is occurring within the link. (Such VLAN mapping within a link between RBridges should not be confused with VLAN mapping inside an RBridge [VLANMAP]). VLAN mapping between VLAN-x and VLAN-y can lead to a loop if the Appointed Forwarders for the VLANs are different. If such mapping within a link was allowed and occurred on two or more links so that there was a cycle of VLAN mappings, a broadcast frame, for example, would loop forever.

To prevent this potential problem, if the DRB on a link detects VLAN mapping by receiving a Hello in VLAN-x that was sent on VLAN-y, it MUST make or revoke appointments so as to assure that the same RBridge (possibly the DRB) is the Appointed Forwarder on the link for both VLAN-x and VLAN-y.

3. The Inhibition Mechanism

An RBridge has, for every link on which it can offer end station service (that is every link for which it can act as an Appointed Forwarder), the following timers denominated in seconds:

- a DRB inhibition timer,
- a root change inhibition timer, and
- up to 4,094 VLAN inhibition timers, one for each legal VLAN ID.

The DRB and root change inhibition timers MUST be implemented.

The loss of native traffic due to inhibition will be minimized by logically implementing a VLAN inhibition timer per each VLAN for which end station service will ever be offered by the RBridge on the link; this SHOULD be done. (See the Appendix for an example motivating VLAN inhibition timers.) However, if implementation limitations make a full set of such timers impractical, the VLAN inhibition timers for more than one VLAN can, with care, be merged into one timer. In particular, an RBridge MUST NOT merge the VLAN inhibition timers together for two VLANs if it is the Appointer Forwarder for one and not for the other, as this can lead to unnecessary indefinitely prolonged inhibition. In the limit, there will be safe operations, albeit with more native frame loss than would otherwise be required, even if only two VLAN inhibition timers are provided: one for VLANs for which the RBridge is the Appointed Forwarder and one for all other VLANs. At least two VLAN inhibition

timers MUST be implemented. Where a VLAN inhibition timer represents more than one VLAN, an update or test that would have been done to the timer for any of the VLANs is performed on the merged timer.

These timers are set as follows:

1. On booting or management reset, each port will have its own set of timers, even if two or more such ports are on the same link, because the RBridge will not have had a chance to learn that yet. All inhibition timers are set to expired except the DRB inhibition timer that is set in accordance with item 2 below. The DRB inhibition timer is handled differently because each port will initially believe it is the DRB.
2. When an RBridge decides that it has become the DRB on a link, including when it is first booted or reset by management, it sets the DRB inhibition timer to the Holding Time of its port on that link that won the DRB election.
3. When an RBridge decides that it has lost DRB status on a link, it sets the DRB inhibition timer to expired.

Note: In the rare corner case where one port of an RBridge was the DRB election winner, but later lost the DRB election to a different port of the same RBridge on that link (perhaps due to management configuration of port priority), neither 2 nor 3 above applies, and the DRB timer is not changed.

4. When an RBridge RB1 receives a TRILL Hello asserting that the sender is the Appointed Forwarder that either (1) arrives on VLAN-x or (2) was sent on VLAN-x as indicated inside the Hello, then RB1 sets its VLAN-x inhibition timer for the link to the maximum of that timer's existing value and the Holding Time in the received Hello. An RBridge MUST maintain VLAN inhibition timers for a link to which it connects if it can offer end station service on that link even if it is not currently Appointed Forwarder for any VLAN on that link.
5. When an RBridge RB1 enables VLAN-x on a port connecting to a link and VLAN-x was previously not enabled on any of RB1's ports on that link, it sets its VLAN inhibition timer for VLAN-x for that link to its Holding Time for that port. This is done even if the port is configured as a trunk or point-to-point port as long as there is some chance it might later be configured not to be a trunk or point-to-point port.

6. When an RBridge detects a change in the common spanning tree root bridge on a port, it sets its root change inhibition timer for the link to an amount of time that defaults to 30 seconds and is configurable to any value from 30 down to zero seconds. This condition will not occur unless the RBridge is receiving Bridge PDU (BPDUs) on the port from an attached bridged LAN. It is safe to configure this inhibition time to the settling time of an attached bridged LAN. For example, if it is known that Rapid Spanning Tree Protocol (RSTP [802.1Q]) is running throughout the attached bridged LAN, it should be safe to configure this inhibition time to 7 seconds or, if the attached bridges have been configured to have a minimum Bridge Hello Timer, safe to configure it to 4 seconds. Note that, while an RBridge could determine what version of spanning tree is running on the physical link between it and any directly connected bridge by examination of the BPDUs it receives, it could not tell if inter-bridge links beyond those directly connected bridges were running classic Spanning Tree Protocol (STP), which might require the root change inhibition timer to be set to 30 seconds for safety.
7. When an RBridge decides that one of its ports (or a set of its ports) P1 is on the same link as another of its ports (or set of its ports) P2, then the inhibition timers are merged to a single set of inhibition timers by using the maximum value of the corresponding timers.
8. When an RBridge decides that a set of its ports that it had been treating as being on the same link are no longer on the same link, those ports will necessarily be on two or more links (one link per port in the limit). This is handled by cloning a copy of the timers for each of the two or more links to which the RBridge has decided these ports connect.

4. Inhibited Appointed Forwarder Behavior

An Appointed Forwarder for a link is inhibited for VLAN-x if:

1. its DRB inhibition timer for that link is not expired, or
2. its root change inhibition timer for that link is not expired, or
3. its VLAN inhibition timer for that link for VLAN-x is not expired.

If a VLAN-x Appointed Forwarder for a link is inhibited and receives a TRILL Data frame whose encapsulated frame is in VLAN-x and would normally be egressed to that link, it decapsulates the native frame

as usual. However, it does not output it to or queue it for that link, although, if appropriate (for example, the frame is multi-destination), it may output it to or queue it for other links.

If a VLAN-x Appointed Forwarder for a link is inhibited and receives a native frame in VLAN-x that would normally be ingressed from that link, the native frame is ignored except for address learning.

An RBridge with one or more unexpired inhibition timers, possibly including an unexpired inhibition timer for VLAN-x, is still required to indicate in TRILL Hellos it sends on VLAN-x whether or not it is Appointed Forwarder for VLAN-x for the port on which it sends the Hello.

Inhibition has no effect on the receipt or forwarding of TRILL Data frames.

5. Multiple Ports on the Same Link

An RBridge may have multiple ports on the same link. Some of these ports may be suspended due to MAC address duplication as described in [RFC6327]. Suspended ports never ingress or egress native frames.

If an RBridge has one or more non-suspended ports on a link and those ports offer end station service, that is, those ports are not configured as point-to-point or trunk ports, then that RBridge is eligible to be an Appointed Forwarder for that link. It can become Appointed Forwarder either by its choice, because it is the DRB, or by appointment by the DRB as described in Sections 2.1 and 2.2.

If an RBridge that is the Appointed Forwarder for VLAN-x on a link has multiple non-suspended ports on that link, it may load share the task of ingressing and egressing VLAN-x native frames across those ports however it chooses, as long as there is no case in which a frame it egresses onto the link from one port can be ingressed on another of its ports, creating a loop. If the RBridge is the Appointed Forwarder for multiple VLANs, a straightforward thing to do would be to partition those VLANs among the ports it has on the link.

6. Security Considerations

This memo provides improved documentation of the TRILL Appointed Forwarder mechanism. It does not change the security considerations of the TRILL base protocol. See Section 6 of [RFC6325].

7. Acknowledgements

The authors of [RFC6325] and [RFC6327], those listed in the Acknowledgements section of [RFC6325] and [RFC6327], and Ron Bonica, Stewart Bryant, Linda Dunbar, Les Ginsberg, Erik Nordmark, Dan Romascanu, and Mike Shand are hereby thanked for their contributions.

8. References

Normative and Informative references for this document are listed below.

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.
- [IS-IS] ISO/IEC 10589:2002, Second Edition, "Intermediate System to Intermediate System Intra-Domain Routeing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)", 2002.
- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, December 1990.
- [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.
- [RFC6326] Eastlake, D., Banerjee, A., Dutt, D., Perlman, R., and A. Ghanwani, "Transparent Interconnection of Lots of Links (TRILL) Use of IS-IS", RFC 6326, July 2011.
- [RFC6327] Eastlake 3rd, D., Perlman, R., Ghanwani, A., Dutt, D., and V. Manral, "Routing Bridges (RBridges): Adjacency", RFC 6327, July 2011.

8.2. Informative References

- [VLANMAP] Perlman, R., Dutt, D., Banerjee, A., Rijhsinghani, A., and D. Eastlake, "RBridges: Campus VLAN and Priority Regions", Work in Progress, October 2011.

Appendix. VLAN Inhibition Example

The per-VLAN inhibition timers (or the equivalent) are needed to be loop safe in the case of misconfigured bridges on a link.

For a simple example, assume that RB1 and RB2 are the only RBridges on the link, that RB1 is higher priority to be the DRB, and that they both want VLAN 1 (the default) to be the Designated VLAN. However, there is a bridge between them configured so that RB1 can see all the frames sent by RB2 but none of the frames from RB1 can get through to RB2.

Both will think they are the DRB. RB1 because it is higher priority even though it sees the Hellos from RB2, and RB2 because it doesn't see the Hellos from RB1 and therefore thinks it is highest priority.

Say RB1 chooses to act as Appointed Forwarder for VLANs 2 and 3 while RB2 chooses to act as Appointed Forwarder for VLANs 3 and 4. There is no problem with VLANs 2 and 4 but if you do not do something about it, you could have a loop involving VLAN 3. RB1 will see the Hellos RB2 issues on VLAN 3 declaring itself Appointed Forwarder, so RB1 will be inhibited on VLAN 3. RB2 does not see the Hellos issued by RB1 on VLAN 3, so RB2 will become uninhibited and will handle VLAN 3 native traffic.

However, this situation may change. RB2 might crash, the bridge might crash, or RB2 might be reconfigured so it no longer tried to act as Appointed Forwarder for VLAN 3, or other issues may occur. So, RB1 has to maintain a VLAN 3 inhibition timer, and if it sees no Hellos from any other RBridge on the link claiming to be Appointed Forwarder for VLAN 3 in a long enough time, then RB1 becomes uninhibited for that VLAN on the port in question and can handle end station traffic in VLAN 3.

Authors' Addresses

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

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

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

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

Yizhou Li
Huawei Technologies
101 Software Avenue,
Nanjing 210012, China

Phone: +86-25-56622310
EMail: liyizhou@huawei.com

Ayan Banerjee
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134 USA

Phone: +1-408-333-7149
EMail: ayabaner@cisco.com

Fangwei Hu
ZTE Corporation
889 Bibo Road
Shanghai 201203
China

Phone: +86-21-68896273
EMail: hu.fangwei@zte.com.cn

