

TRILL Working Group
INTERNET-DRAFT
Expires: August 2007

Donald Eastlake 3rd
Motorola
February 2007

Interaction of RBridges and 802 Protocols

<[draft-eastlake-trill-802-protocols-00.txt](#)>

Status of This Document

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of 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>

Abstract

This is a working document discussing the relationship of RBridges, that is, devices implementing the protocol being developed by the IETF TRILL working group, and various IEEE 802.1 and 802.3 protocols.

INTERNET-DRAFT

RBridges & 802 Protocols

Table of Contents

Status of This Document.....	1
Abstract.....	1
Table of Contents.....	2
1 . Introduction.....	3
1.1 802 Document Designation Conventions.....	3
1.2 IEEE 802 Standards Prcoess.....	3
2 . IEEE 802.1 Protocols.....	6
2.1 802.1D MAC Bridges.....	6
2.2 802.1Q Virtual Bridged Local Area Networks.....	6
2.2.1 802.1ad Provder Bridges.....	6
2.2.2 802.1ag Connectivity Fault Management.....	7
2.2.3 802.1ah Provder Backbone Bridges.....	7
2.2.5 802.1au Congestion Management.....	7
2.3 802.1X Port Based Network Access Control.....	8
2.3.1 802.1af Authenticated Key Agreement.....	8
2.4 802.1AB Connectivity Discovery.....	8
2.5 802.1AE MAC Security.....	9
3 . IEEE 802.3 Protocols.....	10
3.1 802.3as.....	10
4 . Multicast Addresses.....	12
5 . IANA Considerations.....	13
6 . Security Considerations.....	13
7 . References.....	13
7.1 Normative References.....	13
7.2 Informative References.....	13
Copyright, Disclaimer, and Additional IPR Provisions.....	15
Author's Address.....	16
Expiration and File Name.....	16

INTERNET-DRAFT

RBridges & 802 Protocols

1. Introduction

The IETF TRILL Working Group is developing the protocol for RBridges. RBridges are somewhere between being routers and being bridges. Simplifying a lot, RBridges are like Bridges in that they act transparently, form a broadcast domain, and forward primarily based on layer 2 MAC addresses. But they are like routers in that they swap the outer MAC addresses on each RBridge hop, incorporate a hop count, use a link state algorithm traditionally associated with routing, and deal with some layer 3 issues such as optimizing IP multicast and ARP/ND.

The purpose of this document is to briefly survey the relationship of RBridges to some of the myriad 802.1 and 802.3 standards, amendments to those standards, and work in progress to amend those standards which might be relevant. In its present form, it is only intended for internal use in the TRILL Working Group. However, it is possible that some material in this document could end up being appended to the RBridges protocol specification and/or some later version of this document might be published as an informational RFC.

Comments and suggests for improvement are welcome

1.1 802 Document Designation Conventions

IEEE 802 document designations ending with an upper case letter or letters, such as 802.1D, are stand alone standards. Those ending with a lower case letter or letters, such as 802.1ad, are amendments which are intended to be incorporated into a later rollup of the base

standard document they amend. If a hyphen and a year are appended to a document, such as 802.1D-2004, it indicates an adopted standard whose final approval occurred in the year given.

[1.2](#) IEEE 802 Standards Process

Below is a high level description of the IEEE 802 standards process intended to make comments on the standards status of particular 802 efforts more understandable from an IETF perspective. Some details and possible branches are omitted.

The first official stage is when a PAR and 5 Criterion document are drafted and approved by an 802 working group, such as 802.1. A PAR, or Project Authorization Request, is similar to a Charter and the 5 Criterion is a document justifying the need for and reasonableness of expecting success for the effort. In some cases a Study Group, somewhat akin to a BoF, is formed to explore the topic and draft the

PAR and 5 Criterion if the Study Group believes an effort should be chartered.

These documents must then be approved by the 802 Executive Committee.

After that the PAR only goes up to the New Standards Committee (NESCO) of the IEEE Standards Board and after NESCOM approval the PAR must finally be approved by the entire Standards Board (which normally meets the day after NESCOM meetings).

With an approved PAR, similar to an approved working group Charter in the IETF, the 802 working group holding the PAR can proceed with the work itself but, in 802.1 and 802.3, it usually creates a subsidiary Task Group for the PAR or assigns it to an existing Task Group such as the Interworking Task Group in 802.1 that handles bridging.

The body working on the PAR comes up with a Draft, initially version 0.01, perhaps equivalent to a -00 protocol draft in the IETF. It usually continues to work on the this Draft through multiple revisions (0.02, 0.03, ...) until it believes it is ready for Working Group Letter Ballot at which point the latest draft is re-numbered 1.0.

A Letter Ballot is its own process which involves the members of the ballot pool submitting yes, no, or abstain votes. At least 50% of the pool must vote yes or no and 75% of those voting yes or no must vote yes to pass Letter Ballot. Votes can be accompanied by comments and all valid no votes must be accompanied by comments including a description of what change would satisfy the commenter.

If a Draft fails Letter Ballot, the group works on it some more until it again thinks it is ready and then reballots it with the next highest integer version number. Thus, if you see, say, as draft D3.5 you know it has gone through 3 Letter Ballots, the most recent on D3.0, and been modified 5 times since that Letter Ballot, but you don't know if it passed any of these Letter Ballots.

But getting 75% in a letter ballot is not the end of the story. All comments that were submitted with no votes then must be resolved (rejected, accepted, or accepted but with a counter proposal for a satisfying change in the draft) by a 3/4 vote. And when all comments have been resolved, the draft is "recirculated" to the same voting pool. This iterates until certain criterion are met and working group votes to advance to the next step. This does not usually occur until recirculation approval is over 95%.

After getting through Working Group Letter Ballot, the draft then advances to "Sponsor" Letter Ballot where the balloting pool is a group selected by the IEEE to balance the constituencies of interest in the standard. The whole Letter Ballot process then repeats for

this new pool, although the draft numbers don't reset but just keep going up.

Finally, when the standard meets sufficient criteria in Sponsor Letter Ballot, it is forwarded for approval to the Standards Revision Committee (REVCOM) of the Standards Board and then to the whole Standards Board.

INTERNET-DRAFT

RBridges & 802 Protocols

[2.](#) IEEE 802.1 Protocols

802.1 stand alone protocols are listed as second level headings below and amendments to those standards appear as third level headings. in each case, they are ordered by standards labeling order (A to Z then AA, AB, ...).

Note: 802-2001 ("IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture") and its two adopted amendments, 802a-2003 and 802b-2004, define the basic 48 bit address formats, OUIs (organizationally unique identifiers), Ethertypes, and similar basic architectural quantities and formats.

[2.1](#) 802.1D MAC Bridges

- "IEEE Standard for Local and metropolitan area networks / Media Access Control (MAC) Bridges"

802.1D-2004 is the basic 802 bridging description and includes spanning tree. The Rbridges charter implies that unconfigured (plug-and-play) Rbridges provide essentially the same service and that for almost all purposes, a network can be converted from 802.1D bridges to Rbridges by incremental replacement.

Of necessity, 802.1D must be fully accounted for in the specification of Rbridges.

[2.2](#) 802.1Q Virtual Bridged Local Area Networks

- "IEEE Standard for Local and metropolitan area networks / Virtual Bridged Local Area Networks"

VLANs are added to bridging by 802.1Q-2005. VLAN support is included in Rbridges.

[2.2.1](#) 802.1ad Provider Bridges

- "IEEE Standard for Local and metropolitan area networks / Virtual Bridged Local Area Networks / Amendment 4: Provider Bridges" amendment to 802.1Q

To support arbitrary VLANs in a bridged service provider cloud which is providing paths between arbitrary customer clouds, 802.1ad provides for provider VLAN tags that have the identical structure to

customer VLAN tags but a different Ethertype. Presumably the different Ethertype is for robustness or error detection in case of misconfiguration or the like.

This is aimed at solving a problem which is not currently being considered in the RBridge specification.

[2.2.2](#) 802.1ag Connectivity Fault Management

- "Draft Standard for Local and Metropolitan Area Networks / Virtual Bridged Local Area Networks / Amendment 5: Connectivity Fault Management" amendment to 802.1Q

TBD

[2.2.3](#) 802.1ah Provider Backbone Bridges

- "Draft Standard for Local and Metropolitan Area Networks / Virtual Bridged Local Area Networks / Amendment 6: Provider Backbone Bridges" amendment to 802.1Q

The problem with expanding to a two level system with customer and provider bridges just using 802.1ad is presumably that the provider bridges see too many different customer MAC addresses. 802.1ah encapsulates the ultimate MAC addresses so the provider backbone bridges only see the MAC addresses of the devices at the border between the customer clouds and the provider cloud.

This is aimed at solving provider services problems which are not currently being considered in the RBridge specification.

[2.2.5](#) 802.1au Congestion Management

- "Draft Standard for Local and Metropolitan Area Networks / Virtual Bridged Local Area Networks / Amendment 7: Congestion Management" amendment to 802.1Q

The 802.1au draft is in a very early stage (version D0.01, equivalent to an IETF -00 draft).

Congestion Management (BCN, "Backward Congestion Notification") is a method of reducing or eliminating dropped frames within a subnet of bounded delay bandwidth product. The idea is that 802.1au aware devices which experience congestion based on frame queues backing up

INTERNET-DRAFT

RBridges & 802 Protocols

send messages back to frame sources, generally identified by MAC addresses and within MAC address by flow, specifying the maximum rate at which that source should transmit frames. Future frames, from that source, which are tagged to indicate that they conform to such a rate limit are given higher priority.

Bridges have to be congestion aware to support BCN and so would Rbridges. This has been discussed on the TRILL working group mailing list and at TRILL meetings.

[2.3](#) 802.1X Port Based Network Access Control

- "IEEE Standard for Local and metropolitan area networks / Virtual Bridged Local Area Networks / Amendment 4: Provider Bridges"

IEEE 802.1X-2004 is a standard for authenticating devices before permitting access. It is logically unrelated to bridging but is listed here because the 802.1af amendment to 802.1X is relevant to 802.1AE.

[2.3.1](#) 802.1af Authenticated Key Agreement

- "Draft Standard for Local and Metropolitan Area Networks Port-Based Network Access Control / Amendment 1: Authenticated Key Agreement for Media Access Control (MAC) Security" amendment to 802.1X

The 802.1af (sometimes called MacKey) amendment to 802.1X extends it to provided keying for use by 802.1AE (see [section 2.5](#) below).

[2.4](#) 802.1AB Connectivity Discovery

- "IEEE Standard for Local and metropolitan area networks / Station and Media Access Control / Connectivity Discovery"

This standard specifies frames that can be emitted by devices to announce themselves to other devices on the same link. Possibly the

RBridge specification could contain advance as to the 802.1AB announcements sent by RBridges which choose emit such announcements.

[2.5](#) 802.1AE MAC Security

- "IEEE Standard for Local and metropolitan area networks / Media Access Control (MAC) Security"

802.1AE-2006 (sometimes called MacSec) is a completed standard but has no provisions for keying. One set of such provisions is in 802.1af which, as described above in [section 2.3.1](#), is under development as an amendment to 802.1X.

With one exception, 802.1AE simply provides security over one hop between 802.1AE conformant points of attachment. As such, 802.1AE seems equally applicable to such points of attachment on bridges, routers, Rbridges, and end equipment such as non-routing hosts.

This exception being where a piece of customer equipment is connected to provider backbone equipment as described in 802.1ah.

It is not clear how successful 802.1AE will be. The various wireless protocols under 802 (802.11 (Wi-Fi), 802.15.1 (Blue Tooth), 802.15.4 (ZigBee), 802.16 (WiMax), etc.) have all adopted their own security protocols. In the wired world, connections are increasingly point-to-point and for such links you would want security closer to the physical layer so as to provide protection against traffic analysis and avoid leaving the frame structure/size, MAC addresses, etc., visible as plain text the way 802.1AE does. Note that there was a previous general 802 link security protocol, 802.10, which received little deployment and has been withdrawn.

[3. IEEE 802.3 Protocols](#)

- "IEEE Standard for Information technology / Telecommunications and information exchange between systems / Local and metropolitan area networks / Specific requirements / Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications"

802.3-2005 is a completed link standard. The ideas in Rbridges can be applied to a variety of link protocols so, for the most part, the Rbridge architecture is link independent. Nevertheless, it is first being specified for 802.3.

802.3-2005 Clause 31 "MAC Control" defines a low level facility for control on 802.3 links. Subclause 31.4 "MAC Control Frames" gives some further format details while Annex 31A "MAC Control opcode assignments" lists the current opcodes available. There are actually six but the only one I had ever heard of before looking into this document was opcode 1, PAUSE, which is further explained in Annex 31B "MAC Control PAUSE operation".

I expect that PAUSE will eventually be deprecated in most circumstances because 802.1au, Congestion Notification, is superior.

[3.1](#) 802.3as

- IEEE Standard for Information technology / Telecommunications and information exchange between systems / Local and metropolitan area networks / Specific requirements / Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications / Amendment: Frame format extensions" amendment to 802.3

You may have been wondering why no one seems particularly worried about exceeding MTUs when frames are lengthened by the RBridge encapsulation. In fact, 802.3 and similar hardware has long been able to handle frames that were longer than those specified in the standards to accomodate such things.

Apparently, when Ethernet was first specified, the idea was that data would be limited to 1024 bytes and the maximum frame size was set to 1500 to leave lots of room for things like headers and encapsulation. But vendors crammed in as much data as they could so that, when Q-tags (VLAN) were added, even though it was only 4 more bytes, it was thought necessary to define a higher limit for Q-tagged frames. Now that lots of things are considering additional tags, such as security tags with 802.1AE, BCN related tags with 802.1au, provider things like 802.1ad and 802.1ah, and RBridges, a further extension has been

defined along with more stringent words saying not to use up the extra space with data.

The current 802.3as draft defines three frame types for conformance purposes:

- 1500 octets - basic frames
- 1504 octets - Q-tagged frames
- 1982 octets - envelope frames

[4.](#) Multicast Addresses

A block of 16 multicast MAC addresses is reserved for use by 802.1/802.3 protocols (officially 802.1 and media specific protocols). These are in the range from 0180C2000000 to 0180C200000F in hexadecimal.

The general behavior of bridges is to flood multicast on a spanning tree so that multicast frames reach all stations. However, for multicast addresses in this block, the general behaviour of bridges,

routers, etc., is specified as discarding the frame unless they "know what to do with it". This usually involves processing the contents of the frame and that processing may cause frames, including frames addressed to the same multicast destination MAC address, to be originated by the device.

The current uses or anticipated future uses for these addresses are shown below:

- 0180C2000000 - All Bridges: Used for BPDUs. Must be recognized by R Bridges due to R Bridge port participation in spanning tree as a leaf.
- 0180C2000001 - 802.3 Clause 31 use, i.e. Full Duplex PAUSE operation
- 0180C2000002 - 802.3 Clause 43 (Link Aggregation) and Clause 57 (OAM) use, aka "Slow Protocols" Multicast address
- 0180C2000003 - 802.1X Port Authenticator Entity (PAE) address
- 0180C2000004-5 - Reserved for future media access specific method standardization
- 0180C2000006-7 - Reserved for future standardization
- 0180C2000008 - All Provider Bridges
- 0180C2000009-C - Reserved for future standardization
- 0180C200000D - Provider Bridge GVRP Address
- 0180C200000E - 802.1AB Link Layer Discovery Protocol address
- 0180C200000F - Reserved for future standardization

None.

6. Security Considerations

TBD.

7. References

Many of the completed 802 standards are available for free download through the "Get 802" program. See <http://standards.ieee.org/getieee802/>.

7.1 Normative References

None?

7.2 Informative References

[802.1] "IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture", IEEE 802-2001, 8 March 2002.

[802.1D] "IEEE Standard for Local and metropolitan area networks / Media Access Control (MAC) Bridges", 802.1D-2004, 9 June 2004.

[802.1Q] "IEEE Standard for Local and metropolitan area networks / Virtual Bridged Local Area Networks", 802.1Q-2005, 19 May 2006.

[802.1X] "IEEE Standard for Local and metropolitan area networks / Virtual Bridged Local Area Networks", 802.1X-2004, 13 December 2004.

[802.1AB] "IEEE Standard for Local and metropolitan area networks / Station and Media Access Control Connectivity Discovery", 802.1AB-2005, 6 May 2005.

[802.1ad] "IEEE Standard for Local and metropolitan area networks / Virtual Bridged Local Area Networks / Amendment 4: Provider Bridges", 802.1ad-2005 (amendment to 802.1Q), 26 May 2006.

[802.1AE] "IEEE Standard for Local and metropolitan area networks / Media Access Control (MAC) Security", 802.1AE-2006, 18 August 2006.

[802.1af] "Draft Standard for Local and Metropolitan Area NetworksPort-Based Network Access Control / Amendment 1: Authenticated Key Agreement for Media Access Control (MAC) Security", Draft D1.2, 20 January 2007.

[802.1ag] "Draft Standard for Local and Metropolitan Area Networks / Virtual Bridged Local Area Networks / Amendment 5: Connectivity Fault Management", Draft D7.1, 7 November 2006.

[802.1ah] "Draft Standard for Local and Metropolitan Area Networks / Virtual Bridged Local Area Networks / Amendment 6: Provider Backbone Bridges", Draft D3.3, 13 December 2006.

[802.1ak] "Draft Standard for Local and Metropolitan Area Networks / Virtual Bridged Local Area Networks / Amendment 07: Multiple Registration Protocol", Draft D6.0, 10 June 2006.

[802.1au] "Draft Standard for Local and Metropolitan Area Networks / Virtual Bridged Local Area Networks / Amendment 7: Congestion Management", Draft D0.01, 29 September 2006.

[802.3] "IEEE Standard for Information technology / Telecommunications and information exchange between systems / Local and metropolitan area networks / Specific requirements / Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications", 802.3-2005, 9 December 2005.

[802.3as] "Draft Amendment of IEEE Standard for Information technology / Telecommunications and information exchange between systems / Local and metropolitan area networks / Specific requirements / Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications / Amendment: Frame format extensions", Draft D3.1, 24 April 2006.

INTERNET-DRAFT

RBridges & 802 Protocols

Copyright, Disclaimer, and Additional IPR Provisions

Copyright (C) The IETF Trust (2007). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

INTERNET-DRAFT

RBridges & 802 Protocols

Author's Address

Author's Addresses

Donald Eastlake 3rd
Motorola Laboratories
111 Locke Drive
Marlborough, MA 01752

Tel: +1-508-786-7554
Email: Donald.Eastlake@motorola.com

Expiration and File Name

This draft expires in August 2007.

Its file name is [draft-eastlake-trill-802-protocols-00.txt](#).

