

TRILL Working Group	D.M.B. Bond
Internet-Draft	UNH-IOL
Intended status: Standards Track	V.M. Manral
Expires: September 12, 2011	IP Infusion Inc.
	March 11, 2011

RBridges: Operations, Administration, and Maintenance (OAM) Support
draft-bond-trill-rbridge-oam-01

Abstract

The IETF has standardized RBridges, devices that implement the TRILL protocol, a solution for transparent shortest-path frame routing in multi-hop networks with arbitrary topologies, using a link-state routing protocol technology and encapsulation with a hop-count. As RBridges are deployed in real-world situations, operators will need tools for debugging problems that arise. This document specifies a set of RBridge features for operations, administration, and maintenance purposes in RBridge campuses. The features specified in this document include tools for traceroute, ping, and error reporting.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on September 12, 2011.

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](#)
- *1.1. [Requirements Language](#)
- *2. [Acronyms](#)
- *3. [TRILL OAM Message](#)
- *4. [RBridge Tools](#)
- *4.1. [Application RBridge Tools](#)
- *4.1.1. [Hop Count Traceroute](#)
- *4.1.1.1. [Multi-Destination Targets](#)
- *4.1.1.2. [Hop Count Traceroute Example](#)
- *4.1.2. [RBridge Ping](#)
- *4.1.2.1. [Ping Example](#)
- *4.2. [Error Reporting](#)
- *4.2.1. [Hop Count Zero Error](#)
- *4.2.2. [MTU Error](#)
- *5. [TRILL OAM Message Format](#)
- *5.1. [Protocol Code Values](#)
- *5.2. [Protocol Codes Formats](#)
- *5.2.1. [Protocol Application Codes Formats](#)
- *5.2.1.1. [Echo Request](#)
- *5.2.1.2. [Echo Reply](#)
- *5.2.2. [Error Notification Format](#)
- *5.2.2.1. [Error Specifiers](#)
- *5.3. [Type, Length, Value \(TLV\) Encodings](#)
- *5.3.1. [TLV Types](#)
- *5.3.1.1. [Next Hop Nickname](#)

- *5.3.1.2. [Incoming Port ID](#)
- *5.3.1.3. [Outgoing Port ID](#)
- *5.3.1.4. [Outgoing Port MTU](#)
- *6. [Acknowledgments](#)
- *7. [IANA Considerations](#)
- *8. [Security Considerations](#)
- *9. [References](#)
 - *9.1. [Normative References](#)
 - *9.2. [Informative References](#)
- *Appendix A. [Revision History](#)
 - *Appendix A.1. [Changes from -00 to -01](#)
- *[Authors' Addresses](#)

1. Introduction

The IETF has standardized RBridges, devices that implement the TRILL protocol, a solution for transparent shortest-path frame routing in multi-hop networks with arbitrary topologies, using a link-state routing protocol technology and encapsulation with a hop-count ([RFCtrill](#) [*I-D.ietf-trill-rbridge-protocol*]). As RBridges are deployed, operators will face problems that require tools for troubleshooting of connectivity issues in the network. TRILL uses IS-IS for the control plane. IS-IS has a link-state database which contains the information of all links in the TRILL domain and IS-IS has a routing table. This information can be used for trouble shooting purposes. Simply being able to view the link-state database and routing table is insufficient for the requirements of operations, administration, and maintenance (OAM).

In addition, RBridges should support SNMP, as described in [RFCtrill](#) [*I-D.ietf-trill-rbridge-protocol*] and [RBridgeMIB](#) [*I-D.ietf-trill-rbridge-mib*]. SNMP, the routing table, and the link-state database are insufficient as the only OAM tools because while the control plane within an RBridge campus may be functioning successfully the data plane may not be. This motivates the need for OAM tools that allow an operator to test the data plane. Protocols such as IP, MPLS, and IEEE 802.1 have features enabling an operator to exercise the data plane ([RFC 4443](#) [*RFC4443*], [RFC 0792](#) [*RFC0792*], [IEEE 802.1ag](#) [*IEEE.802-1ag*]). There is a need for a similar set of tools in TRILL.

Likewise, there is a need for error reporting capabilities inside an RBridge campus. For instance, if a TRILL Inner.VLAN tag has an illegal value there should be a way for devices to report this error. This would allow administrators of an RBridge campus to quickly locate a problem device in the network. This document specifies a set of RBridge features for operations, administration, and maintenance purposes in RBridge campuses along with a frame format. The features specified in this document include tools for traceroute, ping, and error reporting. [Section 3](#) of this document specifies the general usage of a defined message format. [Section 4](#) specifies some additional applications of the message format. [Section 5](#) specifies the format of the messages on the wire.

1.1. Requirements Language

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 [RFC 2119](#) [RFC2119].

2. Acronyms

- *BPDU - Bridge PDU
- *CHbH - Critical Hop-by-Hop
- *CItE - Critical Ingress-to-Egress
- *DA - Destination Address
- *DR - Designated Router
- *DRB - Designated RBridge
- *ES - End Station
- *ESa - End Station A
- *ESb - End Station B
- *ECMP - Equal-Cost Multi-Path
- *ESADI - End Station Address Distribution Instance
- *FCS - Frame Check Sequence
- *ID - Identification
- *IEEE - Institute of Electrical and Electronics Engineers
- *IETF - Internet Engineering Task Force

*IP - Internet Protocol

*IS-IS - Intermediate System to Intermediate System

*MAC - Media Access Control

*MPLS - Multiprotocol Label Switching

*MTU - Maximum Transmission Unit

*OAM - Operations, Administration, and Maintenance

*P2P - Point-to-point

*PDU - Protocol Data Unit

*RBridge - Routing Bridge

*SA - Source Address

*SNMP - Simple Network Management Protocol

*TLV - Type, Length, Value

*TRILL - TRansparent Interconnection of Lots of Links

*VLAN - Virtual Local Area Network

3. TRILL OAM Message

To facilitate message passing as needed by the OAM requirements, the TRILL OAM Channel ([\[I-D.eastlake-trill-rbridge-channel\]](#)) is utilized. The TRILL Header extended flag MAY be set if so desired. There are two types of TRILL OAM messages defined in this document carried within the TRILL OAM Channel: application and error notification. Frames with an error notification MUST NOT be generated in response to frames with an error notification. Implementations SHOULD rate limit the origination of error notifications. Whereas unknown unicast frames are sent as multi-destination messages, sending unknown unicast frames with an error can lead to an amplification attack. As such special care and rate limiting are necessary for error notifications. The specification of rate limiting is beyond the scope of this document. An RBridge SHOULD maintain counters for each type of error generated. Error notification messages contain the error-causing frame or the initial part thereof after its OAM message. The following are two figures showing application and error notification message structure. [Section 5](#) goes into the details of these formats.

```

+-----+
|   Outer Link Header   |
+-----+
|     TRILL Header     |
+-----+
|   Inner Link Header   |
+-----+
| TRILL OAM Channel Header |
+-----+
| OAM Protocol Spec. Payload |
+-----+

```

Application Frame

```

+-----+
|   Outer Link Header   |
+-----+
|     TRILL Header     |
+-----+
|   Inner Link Header   |
+-----+
| TRILL OAM Channel Header |
+-----+
| OAM Protocol Specific Payload |
+-----+
| Offending Frame TRILL Header |
+-----+
| Offending Frame Inner Link Header |
+-----+
|   Offending Frame Payload   |
+-----+

```

Error Notification Frame

Frames with the TRILL OAM message generated in response to another TRILL data frame MUST have fields set as follows unless otherwise specified:

Frame Type	Field	Value
Application or Error	Inner.MacSA	If the Inner.MacDA of the received frame is one of the MAC addresses of the RBridge generating the frame, the value MUST be that MAC address. Otherwise, it MUST be one of the RBridge's MAC addresses.

Frame Type	Field	Value
Application or Error	Inner.MacDA	The value SHOULD be All-OAM-RBridges . The Inner.MacDA MAY be other values as specified in subsequent sections.
Application or Error	Inner.VLAN ID	The value MUST be one of the VLANs the egress RBridge advertises connectivity on. If the frame is generated in response to another frame it MUST be copied from the received frame.
Application or Error	Ingress RBridge nickname	If the egress RBridge nickname of the received frame is a nickname of the RBridge generating the frame, then the value MUST be that nickname. Otherwise, it MUST be one of the RBridge's nicknames.
Application or Error	Egress RBridge nickname	The value MUST be the ingress RBridge nickname of the received frame. If the ingress RBridge nickname received is unknown or reserved the frame MUST be generated on the port the frame was received on with an Outer.MacDA and egress RBridge nickname of the RBridge that transmitted the invalid frame.
Error	Offending Encapsulated Frame	The value MUST be N bytes of the frame which had the error where N is the minimum of the frame size and the number of bytes that would bring the resulting error frame up to 1470 bytes. This MUST include the TRILL header and MUST NOT include the link-layer header.
Error	M Bit	The value MUST be zero.
Application or Error	Inner.Priority	The value SHOULD be one less than the priority of the received frame, but not less than the lowest priority. Defaults to zero for sent frames.

Frame Field Values

RBridge campuses do not, in general, guarantee lossless transport of frames so a frame containing a TRILL OAM Message, possibly generated in response to some other frame, might be lost.

4. RBridge Tools

This section specifies a number of RBridge OAM tools. For classification purposes they are divided into two sections, applications and error tools.

4.1. Application RBridge Tools

4.1.1. Hop Count Traceroute

The ability to trace the path the data takes through the network is an invaluable debugging tool. RBridge traceroute provides this functionality through use of the TRILL OAM message (See [Section 3](#)). In a hop-count traceroute, the originating RBridge starts by transmitting one TRILL data frame with a TRILL OAM message. This message contains a protocol code of an echo request. (See [Section 5.2.1.1](#)) The ingress RBridge MUST be the RBridge originating the frame.

When a traceroute is initiated, it is either targeting a known unicast target or a multi-destination target as specified by the operator. If the hop-count traceroute is for a known unicast target, the egress RBridge is the destination RBridge to which connectivity will be checked and the M bit MUST be zero. Otherwise, if the hop-count traceroute is for a multi-destination target, the egress RBridge is the distribution tree nickname for the traceroute. Multi-destination targets are handled the same as known unicast targets but require a small amount of additional logic as specified in [Section 4.1.1.1](#).

The first echo request frame transmitted MUST have a hop-count of one. The RBridge will continue transmitting these echo requests, incrementing the hop-count by one each time until a hop-count error notification is received from the destination. Each of these requests in turn will generate a hop-count error notification until the egress RBridge is reached. If a transit RBridge decrements the hop-count by more than one it may transmit multiple hop-count error notifications. The purpose of the traceroute is to confirm connectivity of the data plane, and therefore options such as a flow ID or a security option MAY be included. If an RBridge supports equal-cost multi-pathing (ECMP) or load balancing, the RBridge SHOULD allow operators to specify which flow the traceroute is assigned to. There is no need for all RBridges to use the same assignment method. Being able to specify the flow allows operators to test the path taken by data through the data plane. The purpose of the frame is to mimic a data frame that follows the same path through the data plane that a 'real' data frame would.

The echo request MAY have an arbitrary 32-bit unsigned integer sequence number to assist in matching reply messages to the request. This is important for the hop-count traceroute since replies may return to the ingress RBridge in a different order than their matching requests were sent.

The Inner.VLAN, Inner.MacSA, and Inner.MacDA SHOULD default to the values specified in [Table 1](#). RBridges SHOULD provide an option to change these values to assign the TRILL data frame to a flow.

The replying RBridge MUST include its 16-bit port ID from the port on which the hop-count error generating frame was received in the incoming port field of the reply. It MUST also include its 16-bit port ID from the port on which the frame would be forwarded if the frame did not have a hop-count error. A port ID of 0xFFFF indicates the frame was

consumed by the RBridge itself. Finally the reply MUST include the 16-bit nickname of the next hop RBridge the frame would have been sent to if there were no error. If the request is a multi-destination frame, this field MUST be set to the nickname of the RBridge the frame was received from. This is the previous hop RBridge. This is to facilitate knowledge of a more precise path through the campus as seen in [RFC 5837 \[RFC5837\]](#).

The advantage of this traceroute method is the transit RBridges do not have to do any special processing of the frames until a hop-count error is detected, a condition they are required to detect by the TRILL base protocol. The disadvantage is the request-originating RBridge needs to transmit as many frames as there are hops between itself and the destination RBridge.


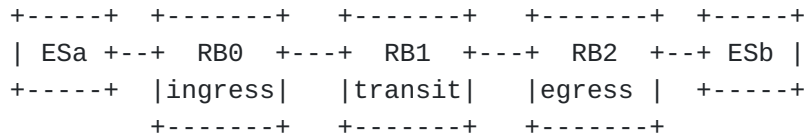
The end stations are not involved in this process. RBridge traceroutes are from RBridge to RBridge. While the frames sent may emulate data sent from ESa to ESb, the end stations are not, in fact, involved.

4.1.1.1. Multi-Destination Targets

For multi-destination targets at each branch in the tree the tagged frame will be replicated causing each RBridge in the tree, possibly pruned by VLAN and/or multicast group, to send a response to the echo request. If all RBridges in the possibly pruned distribution tree support the echo request message, then the ingressing RBridge will receive an echo reply from each of them. This is in contrast to a known unicast tagged frame where only the RBridges along the path from ingress to egress transmit the error notification. The ingressing RBridge can compile all of these replies, using the parent pointers located in the nexthop nickname field, into an output of the tree the traffic traversed. In the case that a non-valid distribution tree nickname is specified the traceroute frames SHOULD still be generated. The traceroute application MUST report any errors received, such as an invalid distribution tree nickname, caused by the hop-count traceroute frames. RBridges receiving a multicast destination echo request MUST NOT transmit an echo reply if the multi-destination bit is set. Echo requests that are not used with the hop-count traceroute come from the ping tool, and pings are not valid to multi-destination traffic. In a hop-count traceroute devices will already be transmitting a hop-count error notification and so there is no reason to transmit a double set of replies. A multi-destination hop-count traceroute does not stop when an echo reply is received. It stops when the transmitted hopcount reaches 0x3F.

4.1.1.2. Hop Count Traceroute Example

[Figure 3](#) contains a campus with three RBridges. Consider a hop-count traceroute from RB0 to RB2.



Hop Count Traceroute Example Topology

In this diagram RB0 transmits frame (1) destined to RB2. This frame contains the echo request message and a hop-count of 0. When RB1 receives this frame it drops it and transmits a hop-count-exceeded message, (2), to RB0. RB0 then transmits a frame, (3), with a hop-count of 1. RB1 decrements this hop-count by 1 to 0 and forwards it to RB2. RB2 drops frame (3) and transmits a hop-count-exceeded message, (4), to RB0. The traceroute is now complete.

Below are some select fields for the frames:

Frame #	Ingress RBridge	Egress RBridge	TRILL OAM Protocol	Sequence Number	Hop Count
(1)	RB0	RB2	Echo Request	1	1
(2)	RB1	RB0	Hop Count Error	1	N/A
(3) @ RB1	RB0	RB2	Echo Request	2	2
(3) @ RB2	RB0	RB2	Echo Request	2	1
(4) @ RB1	RB2	RB0	Hop Count Error	2	N/A
(4) @ RB0	RB2	RB0	Hop Count Error	2	N/A

Hop Count Traceroute Example Frames

For example, if the nicknames for RB0, RB1, and RB2 are 0x0001, 0x0002, and 0x0003 respectively, the console output from such a trace might be:
Hop Count Tracing

RBridge	Incoming Port Id	Outgoing Port Id	RBridge	Nexthop Nickname
0x0001	0xFFFF	0x0001	0x0002	
0x0002	0x0000	0x0001	0x0003	

RBridge	Incoming Port Id	Outgoing Port Id	RBridge	Nextthop	Nickname
0x0003	0x0000	0xFFFF	0x0000		

Hop Count Traceroute Example Output

In this example, the first line of output is generated from local information, no hop-count frames are sent to generate it.

4.1.2. RBridge Ping

Ping is a tool for verifying RBridge connectivity. As with an RBridge traceroute, the ping-originating RBridge transmits one or more TRILL data frames with a TRILL OAM message. This message contains the code of an echo request (See [Section 5.2.1.1](#)). The ingress RBridge MUST be the RBridge-originating frame. The egress RBridge is the destination RBridge to which connectivity will be checked. The M bit MUST be zero. As with RBridge traceroute, options such as a flow ID or a security option MAY be included. If an RBridge supports equal-cost multi-pathing (ECMP) or load balancing, the RBridge SHOULD allow operators to specify which flow the ping is assigned to. There is no need for all RBridges to use the same assignment method. This ping traffic, once again, will mimic real traffic through the network, like traceroute traffic as previously specified in [Section 4.1.1](#).

The echo request MAY have an arbitrary 32-bit unsigned integer sequence number to assist in matching reply messages to the request. In most circumstances, a single echo request is needed to complete the ping but it might be desirable for a single RBridge to ping multiple egress RBridges, or trace differing flows simultaneously. Assigning differing sequence numbers to each frame aids in matching which trace the reply belongs to.

The Inner.VLAN, Inner.MacSA, and Inner.MacDA SHOULD default to the values specified in [Table 1](#). RBridges SHOULD provide the ability to change these values as to assign the TRILL data frame to a flow. The payload of the frame is arbitrary and MAY contain any value. This value can have an influence on which flow the frame is assigned to.

RBridges implementing ping MAY issue a reply in response to this request. See [Section 8](#) for reasons on some RBridges are allowed to choose not to respond to a request. If an RBridge chooses to respond to the request, the reply MUST consist of one TRILL data frame per request with an OAM message containing the protocol code of an echo reply. The echo reply MUST have the same sequence number as the request being matched.

For the echo reply the ingress RBridge field MUST be the reply-originating RBridge's nickname. The egress RBridge MUST be the request-originating RBridge's nickname. The Inner.VLAN, Inner.MacSA, and Inner.MacDA SHOULD default to the values specified in [Table 1](#). The Outer.VLAN ID MUST be preserved. The M bit MUST be zero.

The reply-originating RBridge MUST include its 16-bit port ID from the port on which the request was received in the incoming port field of

the reply. It MUST also include its 16-bit port ID from the port on which the frame is forwarded. A port ID of 0xFFFF indicates the frame was consumed by the RBridge itself. The nickname field in the generated frame MUST be set to all zeros on transmission and ignored on reception.



The Internal Hop Count field of the reply MUST be set to zero. The ping functionality does not use the Internal Hop Count field of the reply. (See [Section 5.2.1.2](#))

The reply frame need not follow the same path though the campus. The reply messages are not meant to test the data plane.

End stations are not involved in this the ping process. RBridge pings are from RBridge to RBridge. While the frames sent may emulate data sent from ESa to ESb, the end stations are not, in fact, involved. The transmitting RBridge MUST wait for a reply frame until a time-out occurs. At that time, the RBridge MUST assume the frame was lost, and this MUST be indicated to the operator. The length of this time-out is not specified in this document.

4.1.2.1. Ping Example

[Figure 4](#) contains a campus with three RBridges. Consider a ping from RB0 to RB2.



Ping Example Topology

In this diagram RB0 transmits frame (1) destined to RB2. This frame contains the echo request message. When RB1 receives this frame it forwards it to RB2. When RB2 receives this frame it transmits and echo reply frame (2) destined to RB0. RB1 receives this frame and forwards it to RB0.

Below are some select fields for the frames:

Frame #	Ingress RBridge	Egress RBridge	TRILL OAM Protocol	Sequence Number
(1)	RB0	RB2	Echo Request	1
(2)	RB2	RB0	Echo Reply	1

Ping Example Frames

For example, if the nicknames for RB0, RB1, and RB2 are 0x0001, 0x0002, and 0x0003 respectively, the console output from such a ping might be:

Pinging

```
... from 0x0001 to 0x0003... 0x0003 is alive
... from 0x0001 to 0x0003... 0x0003 is alive
... from 0x0001 to 0x0003... 0x0003 is alive
```

Ping Example Output

In this example, the ping was repeated three times with the sequence number being changed each time.

[4.2. Error Reporting](#)

Errors can occur through the reception of TRILL data frames. For this purpose, the error notification format is specified. These are generated due to various events as specified subsequently. When a TRILL data frame is received with an error, an error notification frame MAY be generated. See [Section 8](#) for reasons on some RBridges are allowed to choose not to respond to a request. The generated reply MUST contain the error notification. The sub-code MUST contain a code specifying the error encountered. The valid values are specified in [Section 5.2.2.1](#). Two of these sub-codes contain TLVs with additional information. The error notification also contains a 3 bit error type field which describes the error.

This frame has a TRILL header and it contains, as its payload, the frame received with the error. If the size of the received frame would cause the generated frame to exceed 1470 bytes, the payload MUST be truncated to the 1470 bytes. The payload MUST include the TRILL header of the received frame and MUST NOT include the link-layer header. The generated reply MUST contain the error notification message specific to the error.

When the original ingress RBridge receives the error frame, at a minimum, the RBridge SHOULD update a counter specifying the number of error frames received for the causing error. The encapsulated frame MUST NOT be decapsulated and transmitted. The RBridge SHOULD also keep a set of counters for errors reported by other RBridges.

The two sub-codes that contain TLVs with additional information are described below. All other sub-codes specified in this document do not contain TLVs.

[4.2.1. Hop Count Zero Error](#)

When a TRILL data frame is received with a hop-count of zero, an error notification frame MAY be generated. The generated reply MUST contain the hop-count zero error sub-code. If the received frame has the echo


request message, the hop-count zero error notification MUST have a sequence number matching the echo request. Otherwise, the sequence number MUST be set to zero. The incoming port ID MUST be the port ID the received frame arrived on. The outgoing port ID MUST be the port ID of the port the received frame would have been forwarded onto if the hop-count was not zero. Finally, the error notification MUST include the 16-bit nickname of the next hop RBridge the frame would have been sent to. If the request is a multi-destination frame, this field MUST be set to all zeros on transmission and ignored on reception. If the RBridge transmitting the request is the egress RBridge, this field MUST be set to 0x0000.

4.2.2. MTU Error

When a TRILL data frame is received with a payload that would exceed the MTU of the port the frame would otherwise be forwarded to, an error notification frame MAY be generated. The generated reply MUST contain the MTU error sub-code. The outgoing port MTU field MUST have the MTU of the port the frame would have otherwise been transmitted on. The incoming port ID MUST be the port ID the received frame arrived on. The outgoing port ID MUST be the port ID of the port the received frame would have been forwarded onto if the frame size was not too large. Finally, the error notification message MUST include the 16-bit nickname of the next hop RBridge the frame would have been sent to. If this is a multi-destination frame this field MUST be set to all zeros on transmission and ignored on reception. If the RBridge transmitting the request is the egress RBridge, this field MUST be set to 0x0000.

5. TRILL OAM Message Format

This section specifies the format of the TRILL OAM message on the wire beyond the ethertype as encoded in the OAM Channel



TRILL OAM Message Common Initial Part

The message fields and flags are as follows:

*Version, TRILL OAM Protocol, SL, MH, Flags, and ERR: The usage is specified in [\[I-D.eastlake-trill-rbridge-channel\]](#). The SL bit SHOULD be 0. The MH bit MUST be 1. The version must be 0. ERR MUST be all zeros. The TRILL OAM Protocol is further specified by the tool type.

*Sequence Number: This field is used to sequence frames for certain tools. Not all tools utilize the sequence number field.

5.1. Protocol Code Values

The protocol code values which specify the tool type are:

*0x004 (Suggested): Echo Request, See [Section 5.2.1.1](#)


*0x005 (Suggested): Echo Reply, See [Section 5.2.1.2](#)

*0x006 (Suggested): Error Notification, See [Section 5.2.2](#)

5.2. Protocol Codes Formats

5.2.1. Protocol Application Codes Formats

5.2.1.1. Echo Request



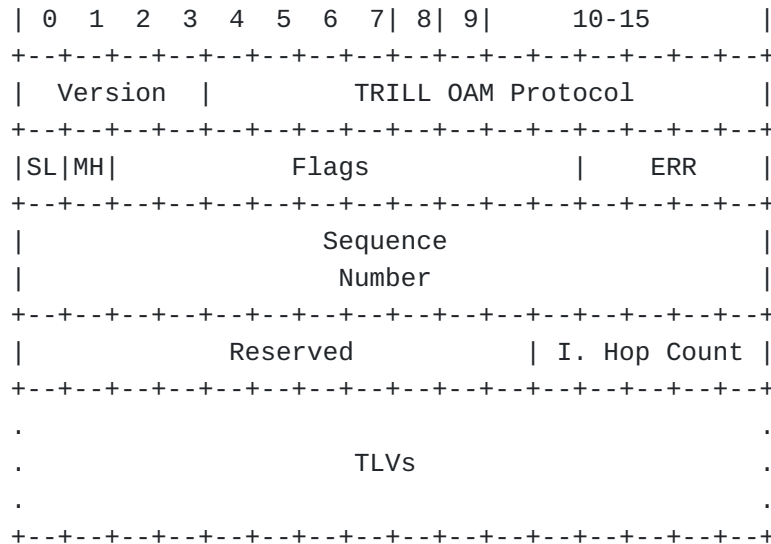
Echo Request

This message is used by ingress RBridges to request an echo reply from the egress RBridge. Further uses are specified in [Section 4.1.1](#) and [Section 4.1.2](#)

*Version, TRILL OAM Protocol, SL, MH, Flags, and ERR: The usage is specified in [\[I-D.eastlake-trill-rbridge-channel\]](#). The SL bit SHOULD be 0. The MH bit MUST be 1. The version must be 0. ERR MUST be all zeros. TRILL OAM Protocol MUST be 0x004 (Suggested).

*Sequence Number: An arbitrary 32-bit unsigned integer used to aid in matching reply messages to echo requests. MAY be zero.

5.2.1.2. Echo Reply



Echo Reply Format

This message is used by egress RBridges to reply to an echo request from the ingress RBridge. Further uses are specified in [Section 4.1.1](#) and [Section 4.1.2](#).

*Version, TRILL OAM Protocol, SL, MH, Flags, and ERR: The usage is specified in [\[I-D.eastlake-trill-rbridge-channel\]](#). The SL bit SHOULD be 0. The MH bit MUST be 1. The version must be 0. ERR MUST be all zeros. TRILL OAM Protocol MUST be 0x005 (Suggested).

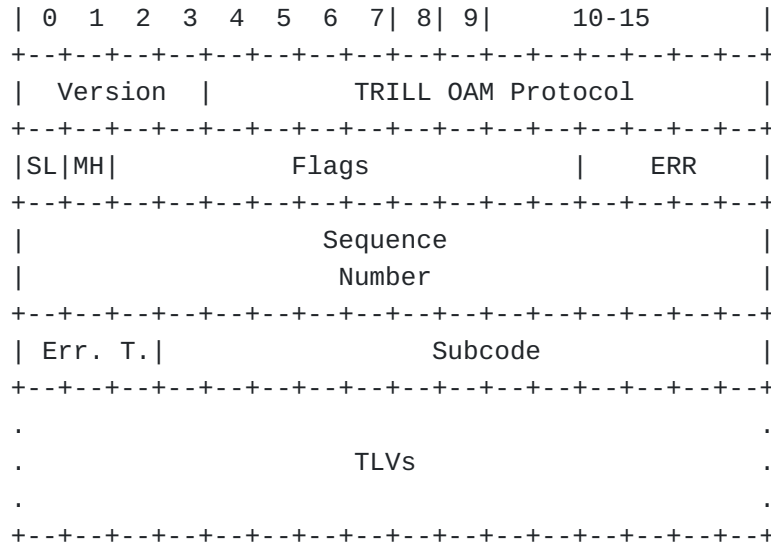
*Reserved: A reserved field. Set to zero on transmission and ignored on reception.

*Internal Hop Count: If the request being replied to was an echo request, this value MUST be zero on transmission and ignored on reception. If the request being replied to was a respond request, this value is a copy of the TRILL Hop Count value in the request. The reserved and internal hop-count fields combined occupy the subcode field of the TRILL OAM message.

*Sequence Number: A 32-bit unsigned integer used to aid in matching reply messages to echo requests. This MUST match the request being replied to.

*TLVs: A set of type, length, value encoded fields as specified in [Section 5.3](#). The next hop nickname, outgoing port ID, and incoming port ID TLVs are required.

5.2.2. Error Notification Format



Error Format

This message is used by RBridges to signal that an error has occurred.

*Version, TRILL OAM Protocol, SL, MH, Flags, and ERR: The usage is specified in [\[I-D.eastlake-trill-rbridge-channel\]](#). The SL bit SHOULD be 0. The MH bit MUST be 1. The version must be 0. ERR MUST be all zeros. TRILL OAM Protocol MUST be 0x006 (Suggested).

*Sequence Number: For all sub-codes except for the hop count error this field is unused. It is set to zero on transmission and ignored on reception. For the hop count error this is a 32-bit unsigned integer used to aid in matching reply messages to echo requests requests. If the frame whose hop-count dropped to zero contains the echo request message (See [Section 5.2.1.1](#)), this MUST match the sequence number echo request found in that message. If this is not in reply to a request, then the sequence number MUST be set to zero.

*Error Type: MUST be a specifier of the error type describing the error. The values are: 0 (Permanent Error), 1 (Transient Error), 2 (Warning), 3 (Comment). Values 4 through 7 are available for allocation by IETF Review.

*Subcode: MUST be a specifier of the error discovered in the frame. The valid values are specified in [Section 5.2.2.1](#)

*TLVs: A set of type, length, value encoded fields as specified in [Section 5.3](#). For next hop errors the next hop nickname, outgoing port ID, and incoming port ID TLVs MUST be present. For MTU errors the outgoing port MTU, next hop nickname, outgoing port ID, and incoming port ID TLVs MUST be present. For all other errors the TLVs are not used and the length of this section is set to zero.

5.2.2.1. Error Specifiers

The sub-code values fall into three categories: errors, warnings, and comments. All sub-codes represent something out of the ordinary that has gone wrong, but certain ones are more important than others. Sub-codes that are classified as errors are the most severe with warning sub-codes being slightly less severe. These are enabled by default. Sub-codes classified as comments are minor and are disabled by default. They may be useful for operators debugging a network. All error generations are optional and therefore MAY be generated or not generated depending on security and implementation constraints. The error specifiers sub-code values are:
Sub-codes

*0: Unknown Error: Indicates an error has occurred.

*1: Corrupt Frame: Frame received with invalid FCS or that was not an 8-bit multiple in length. It could be impossible for a device to signal this if the low-level port hardware hides this from the software.

*2: Invalid Outer.MacDA: Indicates the MAC Address is a multicast address and the M bit is zero, the MAC Address is not a multicast address and the M bit is one, or the M bit is zero and the frame carried is an ESADI frame.

*3: Illegal Outer.VLAN: Indicates the Outer.VLAN ID is 0xFFFF.

*4: Invalid Outer.VLAN: Indicates the Outer.VLAN ID was not the designated VLAN ID.

*5: Unknown TRILL Version: Indicates the TRILL Version is unknown.

*6: Op-Length Exceeds Frame Length: Indicates the Op-Length says the options field extends beyond the end of the received frame length.

*8: Unknown Egress RBridge: Indicates the Egress RBridge in a received frame is unknown.

*9: Unknown Ingress RBridge: Indicates the Ingress RBridge in a received frame is unknown.

*10: Unsupported Critical Hop-by-hop Option: Indicates an unsupported critical hop-by-hop option was received.

*11: Unsupported Critical Ingress-to-Egress Option: Indicates an unsupported critical ingress-to-egress option was received.

*12-84: Available for allocation by IETF Review

*85: Reserved for Private Experimentation

Warning Sub-codes

*86: Illegal Inner.VLAN: Indicates the Inner.VLAN ID is 0xFFFF.

*87: Inner/Outer VLAN Priority Mismatch: Indicates the priority values in the inner and outer VLANs do not match.

*88: P2P Hello on TRILL Hello Link: Indicates a P2P Hello was received on a TRILL Hello Link.

*89: TRILL Hello on P2P Hello Link: Indicates a TRILL Hello was received on a P2P Hello Link.

*90: No Adjacency: Indicates a TRILL data frame was sent from an RBridge the receiving RBridge is not adjacent to.

*91: Encapsulated BPDUs/VRP Frame: A TRILL Frame containing a BPDUs or VRP frame was received.

*92: Invalid Mutability Flag: Indicates the mutability flag was set on a received CHbH Option.

*93: Invalid TLV Option Length: Indicates the option length field of a TLV option was between 121 and 127.

*94: Options Ordering Error: Indicates the TLV options are ordered incorrectly.

*95: Additional Flag TLV Zero: Indicates a problem in the additional Flag TLV.

*96: Configured Nickname Collision: Indicates an RBridge was detected in the campus with the same nickname (Configured or not).

*97: Multiple DRBs detected.

*98: Multiple appointed forwarders detected.

*99-169: Available for allocation by IETF Review

*170: Reserved for Private Experimentation

Comment Sub-codes

*171: Inner.VLAN C-Bit Set: Indicates the C-Bit in the Inner.VLAN is set.

*172: Unknown Inner.MacDA: Indicates the Inner.MacDA is unknown. This may occur if devices are configured to explicitly register end stations and an unknown Inner.MacDA occurs in a unicast TRILL data frame. This also only applies at egress and could indicate that the Inner.MacDA was a learned address that has timed out.

*173: Unknown Inner.MacSA: Indicates the Inner.MacSA is unknown. This may occur if devices are configured to explicitly register end stations and an unknown Inner.MacSA occurs in a TRILL data frame.

*174: Outer.VLAN C-Bit Set: Indicates the C-Bit in the Outer.VLAN is set for an Ethernet frame.

*175: Invalid Reserved Bits: Indicates the reserved bits are non-zero in a received frame.

*176: Invalid Nickname: Indicates a nickname in the reserved space of 0xFFC0 to 0xFFFF was received that is not implemented at the receiving RBridge.

*177: Unsupported Non-Critical Hop-by-hop Option: Indicates an unsupported non-critical hop-by-hop option was received. While sending a non-critical option to an unsupported device is not an error, this could be used to support identification of devices needing an upgrade.

*178: Unsupported Non-Critical Ingress-to-Egress Option: Indicates an unsupported non-critical ingress-to-egress option was received. While sending a non-critical option to an unsupported device is not an error, this could be used to support identification of devices needing an upgrade.

*179: Performance Exceeded: Indicates a frame was discarded due to performance problems such as a buffer overflow.


*180: Insufficient Hop Count: Indicates a frame was received with a hop-count that was insufficient to reach the destination.

*181-254: Available for allocation by IETF Review

*255: Reserved for Private Experimentation


5.3. Type, Length, Value (TLV) Encodings

To facilitate future interoperable expansion of the data carried in OAM sub-messages some sub-messages use a TLV encoding. These TLV sections consist of a list of type, length, value encoded data where the type signals to the RBridge how to interpret the value, and the length tells the RBridge the length of the value in bytes. The type and length are both 8 bit fields. A length of zero indicates the value is a UTF-8 string with a NULL ('\0') terminating byte. Preceding the list of TLVs is a 16 bit total length field which specifies the total length of all the length fields in octet units.



TLVs Format

Each TLV in the TLV List appears on the wire as such:



TLV Format

The type values are:

- *0: Next Hop Nickname, See [Section 5.3.1.1](#)
- *1: Outgoing Port ID, See [Section 5.3.1.3](#)
- *2: Incoming Port ID, See [Section 5.3.1.2](#)
- *3: Outgoing Port MTU, See [Section 5.3.1.4](#)

*4-253: Available for allocation by IETF Review

*254: Reserved for Private Use

*255: Reserved

5.3.1. TLV Types

5.3.1.1. Next Hop Nickname

```
| 0 1 2 3 4 5 6 7| 8| 9| 10-15 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type = 0x01   |   Length = 0x02   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Next Hop Nickname               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

Next Hop Nickname Format

For traceroutes targeting known unicast destinations, hop-count errors, and MTU errors, this TLV MUST be the 16-bit nickname of the next hop RBridge the frame is being or would have been sent to. If the RBridge transmitting the TLV is the egress RBridge this field MUST be set to 0x0000. For traceroutes targeting multi-destination destinations, e.g. with the TRILL M bit high, this field contains the nickname of the RBridge the frame being responded to is from. For pings, this field MUST be set to all zeros on transmission and ignored on reception. For multi-destination hop-count errors this field contains the nickname of the RBridge the frame with the exceeded hop-count was sent from. For multi-destination MTU error traffic, this field MUST be set to all zeros on transmission and ignored on reception.

5.3.1.2. Incoming Port ID

```
| 0 1 2 3 4 5 6 7| 8| 9| 10-15 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type = 0x02   |   Length = 0x02   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Incoming Port ID               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

Incoming Port ID Format

This TLV MUST be set to the Port ID found in 'The Special VLANs and Flags sub-TLV' for the port the request being replied to was received on. ([\[I-D.ietf-isis-trill\]](#))

5.3.1.3. Outgoing Port ID

```
| 0 1 2 3 4 5 6 7| 8| 9| 10-15 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type = 0x03   |   Length = 0x02   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Outgoing Port ID               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

Outgoing Port ID Format

This TLV MUST be set to the Port ID found in 'The Special VLANs and Flags sub-TLV' for the port the frame is being forwarded on to (or would have been for an echo request/hop-count error). ([\[I-D.ietf-isis-trill\]](#)) If the request was consumed by the replying RBridge, the port ID MUST be 0xFFFF.

5.3.1.4. Outgoing Port MTU

```
| 0 1 2 3 4 5 6 7| 8| 9| 10-15 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type = 0x04   |   Length = 0x02   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Outgoing Port MTU               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

Outgoing Port MTU Format

This TLV MUST be the MTU of the outgoing port specified in the outgoing port ID TLV.

6. Acknowledgments

Many people have contributed to this work, including the following, in alphabetic order: Sam Aldrin, Dinesh Dutt, Donald E. Eastlake 3rd, Anoop Ghanwani, Jeff Laird, and Marc Sklar

7. IANA Considerations

IANA is request to create a new subregistry within the TRILL Parameter registry for "TRILL OAM Message Error Sub-Message Error Specifiers". This subregistry that is initially populated as specified in [Section 5.2.2.1](#). Additional values are allocated by IETF Review [\[RFC5226\]](#). IANA is requested to create a new subregistry within the TRILL Parameter registry for "TRILL Error Reporting Protocol TLV Types" with

initial values as listed in Section 5.3. Additional values are allocated by IETF Review [\[RFC5226\]](#).

This draft also requires action to reserve the TRILL OAM Control Channel protocol codes. IANA is requested to allocate the TRILL OAM Channel protocol codes for as listed in [Section 5.1](#).

8. Security Considerations

The nature of the TRILL OAM Message lends itself to security concerns. By providing information about the topology of a network, attackers can gain greater knowledge of a network in order to exploit the network. Passive attacks such as reading frames with an OAM message could be used to gain such knowledge or active attacks where an attacker mimics an RBridge can be used to probe the network. Authentication, data integrity, protection against replay attacks, and confidentiality for TRILL OAM frames may be provided using a to-be-specified TRILL Security Option. Using such a security option would mitigate both the passive and active attacks.

For instance, data origin authentication could be provided in the future using a security options in the TRILL Header by verifying a hash using shared keys or a mechanism like SEND with CGA. To prevent replay attacks rate limiting, sequence numbers as well as some nonce based mechanism could be provided. Confidentiality for TRILL OAM frames could be provided based on some future security option extension which encrypts TRILL frames.

In a network where one does not wish to configure a security option, the threat of attackers is still present. For this reason, generation of any TRILL OAM Message frames is optional and SHOULD be configurable by an operator on a per RBridge basis. An RBridge MAY have this configurable on a per port basis. For instance, an operator SHOULD be able to disable hop-count traceroute reply messages or error notification message generation per port.

Another security threat is denial of service through use of OAM messages. For this reason, RBridges MUST rate limit the generation of OAM message frames. For multi-destination frames, the frames MAY be discarded silently to prevent any denial of service attacks in case of an errored packet such as an 'options not recognized' error notification.

9. References

9.1. Normative References

- | | |
|-----------------------------------|---|
| [I-D.ietf-trill-rbridge-protocol] | 3rd, D, Dutt, D, Gai, S, Ghanwani, A and R Perlman, " Rbridges: Base Protocol Specification ", Internet-Draft draft-ietf-trill-rbridge-protocol-16, March 2010. |
| [I-D.ietf-isis-trill] | |

- 3rd, D, Banerjee, A, Dutt, D, Perlman, R and A Ghanwani, "[TRILL Use of IS-IS](#)", Internet-Draft draft-ietf-isis-trill-05, February 2011.
- [RFC2119] [Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels"](#), BCP 14, RFC 2119, March 1997.
- [I-D.eastlake-trill-rbridge-channel] 3rd, D, Manral, V, Ward, D, Yizhou, L and S Aldrin, "[RBridges: OAM Channel Support in TRILL](#)", Internet-Draft draft-eastlake-trill-rbridge-channel-00, March 2011.

9.2. Informative References

- [IEEE. 802-1ag] Institute of Electrical and Electronics Engineers, "IEEE Standard for Local and metropolitan area networks / Virtual Bridged Local Area Networks / Connectivity Fault Management", IEEE Standard 802.1Q, December 2007.
- [I-D.ietf-trill-rbridge-mib] Rijhsinghani, A and K Zebrose, "[Definitions of Managed Objects for RBridges](#)", Internet-Draft draft-ietf-trill-rbridge-mib-02, March 2011.
- [RFC0792] Postel, J., "[Internet Control Message Protocol](#)", STD 5, RFC 792, September 1981.
- [RFC4443] Conta, A., Deering, S. and M. Gupta, "[Internet Control Message Protocol \(ICMPv6\) for the Internet Protocol Version 6 \(IPv6\) Specification](#)", RFC 4443, March 2006.
- [RFC5226] Narten, T. and H. Alvestrand, "[Guidelines for Writing an IANA Considerations Section in RFCs](#)", BCP 26, RFC 5226, May 2008.
- [RFC5837] Atlas, A., Bonica, R., Pignataro, C., Shen, N. and JR. Rivers, "[Extending ICMP for Interface and Next-Hop Identification](#)", RFC 5837, April 2010.

Appendix A. Revision History

RFC Editor: Please delete this appendix before publication.

Appendix A.1. Changes from -00 to -01

- *Reworked the document to use the OAM Channel rather than an OAM option.
- *Changed the frame formats to work within the OAM Channel.
- *Numerous minor typo corrections and wording clarifications.
- *Removed the route-respond traceroute.
- *Combined all the error notifications into one OAM Channel.

Authors' Addresses

David Michael Bond Bond University of New Hampshire InterOperability
Laboratory 121 Technology Drive Suite #2
Durham, New Hampshire 03824 US Phone: +1-603-339-7575 EMail:
david.bond@iol.unh.edu URI: <http://mokon.net>

Vishwas Manral Manral IP Infusion Inc. 1188 E. Arques Ave.
Sunnyvale, CA 94089 US EMail: vishwas@ipinfusion.com